

Us and Them:

The Nature, Value, Limits, and Opportunities of Agreement and Disagreement

By Joel Jordan Wesley Yalland

Department of Philosophy, University of Kent

A thesis submitted in fulfilment of the degree of Doctor of Philosophy in Philosophy.

January 2024

102,463 words.

Go placidly amid the noise and the haste and remember what peace there may be in silence.

As far as possible, without surrender, be on good terms with all persons.

Speak your truth quietly and clearly; and listen to others, even to the dull and ignorant; they too have their story.

Max Ehrmann, *Desiderata*

It's the questions we can't answer that teach us the most. They teach us how to think. If you give a man an answer, all he gains is a little fact. But give him a question and he'll look for his own answers.

Patrick Rothfuss, *The Wise Man's Fear*

## Acknowledgements

This thesis has been a longstanding labour of love, albeit sometimes begrudging love. However, there are many people who have supported, encouraged, and challenged me through the process of writing it. The following are simply the brightest and most prominent figures, and many of them are presented as hypothetical speakers and agents in the examples I provide as tribute to them.

First and foremost, my supervisor, Professor Simon Kirchin, who has always been patient, caring, and understanding through seven years of my attempts to balance work, writing, and 'life', for which I cannot thank him enough. There have been at least a few times where he has made better sense of my ideas and arguments than I could ever hope to manage in my often-exhausted state but without his support this thesis simply would not have come together. It has been a genuine pleasure to work and engage with him over seven or so years, and he has my most sincere gratitude.

Second, various colleagues at the Universities of Kent and Glasgow. I thank my secondary supervisor, Dr Edward Kanterian, whose comments have all been constructive and thought-provoking. My original thesis project was utterly different from this one and late into my second year I was entirely uncertain of whether to change topics altogether. A conversation with Dr Graeme A. Forbes during an annual departmental walk brought me to the topic of disagreement, and without that timely discussion, this thesis would not likely have come about, for which he has my full thanks. Likewise, I am extremely grateful to Professor Mona Simion, for facilitating my membership within the Glasgow Philosophy department as a visiting PhD student, welcoming me into a network of outstanding professional philosophers and enabling me to lay the foundations for this thesis. I am also very grateful to my viva examiners, Professor Jon Williamson and Professor Adam Carter, whose questions made for thought-provoking and spirited discussions.

Moreover, various fellow PhD students and postdoctoral philosophers I have had the benefit of knowing through Kent Philosophy Graduate Society (PhiGS) and the COGITO Epistemology Research Centre at Glasgow.

At Kent these include Gavin Thomson, Mark Garron, Dr James Hootor, and Dr Alyx Robinson, and at Glasgow, Katharina Bernhard, Giorgia Foti, Dr Tim Kearn, and Dr Lilith Mace. Each of

these individuals have offered spirited and enrapturing conversations, often over post-talk or post-conference drinks, if not simply while sharing time and office spaces.

Third, the writing and editing of this thesis required two sabbatical periods in Glasgow and Los Angeles, respectively, for both of which I was out of work for three months and focused primarily on developing my thesis material. This meant I was entirely dependent on the generosity and hospitality of mutual contacts of my friends and family. As such, I am indebted to Maris Thomson and her cat Leo, and Ace Trump and his chief of security, Boomer the German Shepherd.

Fourth, my mother, Sarah, who coined my childhood nickname “professor” never deterring me from being inquisitive and who has always wanted nothing but the best for me. My father, Adrian, who has never doubted my thesis topic and has constantly championed its timeliness and ongoing relevance, as well as providing the occasional ~~distraction~~ discussion of music to occupy my time (especially regarding Pink Floyd, from whom the title of this thesis comes.)

My stepparents Richard and Chrissy, and my sister Amy, who likewise have always looked out for my best interests and offered clarity and practical, personal, and professional advice whenever needed. My extended family, especially my cousin Freya, with whom I have frequently vented about the process of writing and hope one day will be the second PhD in the family. Many of them have at one time or another struggled to understand the process or topic of this thesis but have nonetheless encouraged me and offered constant love, support, and prayers. Likewise, my dogs, Oscar and Bobby, a regular source of entertainment and comfort.

Fifth, various friends and colleagues, whether in Canterbury, London, or home in Devon. I regard many of them as my second family and they have often sustained me through various personal and professional struggles alongside writing this thesis, always pushing me to persist. My sincere thanks go to Ciara, Martin, Kiera, Lea, Nigel, Summer, Helen, Alex, Gary, James, Jake, Jon, George, and Dan.

This thesis is dedicated to the memories of my grandparents: June, Margaret, and Tony.

They all unfortunately passed away before I could see this thesis through to completion, but always had the genuine and sincere care for my wellbeing and progress, as well as the good sense to not enquire too much into my thesis topic or progress.

## Table of Contents

Acknowledgments	3
Dedication	5
<u>Introduction</u>	9
Philosophical Discussion of Agreement and Disagreement	11
Thesis Structure	13
Clarifications and Caveats	16
Lay Summary	19
<u>Part 1 – Agreement and Disagreement: Broad, Complex, and Indefinable</u>	20
<u>Chapter 1 – Simple and Doxastic Agreement and Disagreement</u>	21
1.1 – Agreement and Disagreement in Belief	26
1.2 – Inconsistency, Contradiction, and Contrariety	29
1.3 – States of Tension and Changes of Mind	33
1.4 – Superficial Disagreement?	36
1.5 – Agreement and Disagreement are Complex	37
1.6 – Credences and Epistemic Modals	40
1.7 – Judgment Suspension and Atheist/Agnostic Cases	43
1.8 – Recap	45
<u>Chapter 2 – Agreement, Disagreement, and Non-Doxastic States</u>	47
2.1 – Appraisal and Persuasion	49
2.2 – Motives	51
2.3 – Desires and Preferences	55
2.4 – Prescriptions and Recommendations	58
2.5 – Plans and Planning States	63
2.6 – Correctness and Normative States	66
2.7 – Disagreement as a Normative Judgment	69
<u>Chapter 3 – Respects, Speakers, and Ambiguity</u>	73
3.1 – The Simple View Revisited	74
3.2 – Satisfaction and Justification	77
3.3 – Joint Accuracy	79
3.4 – Perception and Incongruity	81
3.5 – Speaker Types and Devil’s Advocates	84
3.6 – Contrarians and Conformists	86

3.7 – Strangers	87
3.8 – Agreement and Disagreement as Real, Relevant, and Robust	89
3.9 – Concept Models and the Concepts of “Agreement” and “Disagreement”	95
3.10 – Family Resemblance and Ambiguity	98
 <u>Part 2 – Epistemic Peerage, Agreement, and Disagreement</u>	 102
<u>Chapter 4 – Three Debates Plot a Terrain</u>	105
4.1 – Uniqueness and Permissiveness of Evidence	109
4.2 – Uniqueness and Peer Disagreement	111
4.3 – What is Higher Order Evidence?	113
4.4 – The Import of Higher Order Defeat	115
4.5 – Peer Disagreement <i>as</i> Higher Order Evidence	118
 <u>Chapter 5 – Critiquing the Three Debates</u>	 121
5.1 – Uniqueness and Permissiveness Revisited	122
5.2 – Non-Extreme Uniqueness and Permissiveness	124
5.3 – Peer Disagreement and Non-Extreme Uniqueness	128
5.4 – Higher Order Evidence, Again	132
5.5 – Peer Disagreement and/as Higher Order Evidence	138
5.6 – Why Evidence of Peer Disagreement <i>is not</i> Higher Order Evidence	141
5.7 – Summary Thoughts	147
 <u>Chapter 6 – Peer Disagreement: Uncompelling, Inconclusive, and Insignificant</u>	 149
6.1 – Peer Disagreement is Uncompelling (in Four Parts)	153
6.2 – Equal Weight	154
6.3 – Extra Weight	158
6.4 – Right Reasons	161
6.5 – Total Evidence	165
6.6 – Peer Disagreement <i>Does Not</i> Advance Inquiry	168
6.7 – Peer Disagreement <i>is Insignificant</i>	173
 <u>Chapter 7 – Peer Disagreement is Idealised, Peer Agreement is Overlooked</u>	 179
7.1 – Epistemic Peerage is Unattainable	181
Fact or Perception?	183
Necessary or Sufficient?	185
Possible or Actual?	186
7.2 – The What, How, and Why of Idealisation	190

7.3 – Peer Disagreement is Idealised	194
7.4 – The Problem of Idealisation in Peer Disagreement	198
7.5 – Peer Disagreement Overlooks Peer Agreement	202
Total Evidence and Right Reasons	204
Equal and Extra Weight	207
 <u>Chapter 8 – Epistemic ‘Non-Peers’: Agreement, Disagreement, and Deference</u>	 211
8.1 – Epistemic Superiority	214
Experts	215
Authorities	217
Hybrid Accounts	217
8.2 – Epistemic Inferiority	218
Testimonial Disregard	219
Novices and Subordinates	220
Competent Inferiors	220
Incompetent Inferiors	221
8.3 – Disagreement with Epistemic Superiors	221
8.4 – Agreement with Epistemic Superiors	229
8.5 – Disagreement with Epistemic Inferiors	233
8.6 – Agreement with Epistemic Inferiors	238
8.7 – Closing Thoughts	242
 <u>Chapter 9 – An Afterthought on Agreement, Disagreement, and Testimony</u>	 244
9.1 – Acknowledgement and Purported Testimony	246
9.2 – Acknowledgement and Uptake	248
9.3 – Epistemological Implications	251
9.4 – Further Considerations	253
 <u>Bibliography</u>	 256



## Introduction

Agreement and disagreement are far harder to define or specify than we might think.

Perhaps this seems like an absurd claim, if not one that is plainly false, or perhaps it seems obvious or trivially true. Nonetheless it is an important point that requires elaboration, which is what this thesis, or at least the first part of it, aims to do. There might be many abstract and hypothetical cases of agreement and disagreement that we can think of as well as real life instances we have experienced, but I contend that there is not very much that they have in common except that we regard them all as agreements or disagreements. Likewise, we may have various concepts that we associate with agreement and disagreement, or use as technical or layfolk descriptors but, again, they will not capture all kinds of agreement and disagreement. As a result, many of the accounts of agreement and disagreement available can only really be described as *narrow*, whereas the account I offer is intentionally *broad*, albeit perhaps vague or imprecise. I contend that it is a mistake to think that we can describe or capture all cases of agreement and disagreement by prioritizing or emphasizing some specific factor, relation, or state, or a particular combination of these. No narrow account that attempts to explain or define agreement and disagreement in tightly specified ways or by focusing on specific aspects, factors, or types of agreement and disagreement, will succeed in describing *all* agreement or disagreement cases. Agreement and disagreement are, I contend, simply too elaborate, jumbled, and fluctuating to reduce to a single definition or factor that is common to all kinds and cases. In other words,

Agreement and disagreement are complex, multifaceted phenomena.

This might seem more obvious than the claim above, and again it may seem trivially true, but this would also be a mistake to believe. In technical philosophical terms, I regard them and argue that they are *ambiguous* concepts and phenomena; the intricacies of agreement and disagreement are greater than any single account can explain because, I contend, there are multiple factors that can vary and differ from one instance to another. The broad claim that agreement and disagreement are complex, multi-faceted phenomena is the most specific, all-encompassing claim we can make about agreement and disagreement *simpliciter*. However, while I provide and grant all the same details, I do not assume that any one account is exclusive and exhaustive, or not nearly as much as those who defend the accounts I consider. Because my broad claim is vague and imprecise, it might be unconvincing or insufficient for many people. In response, I note that while we can point to a range of situations and often say with confidence that “that’s a disagreement” or “this

looks like agreement”, we cannot always say what makes an agreement case an agreement case, or a disagreement case a disagreement case, by relying on a tight mould, format, or definition.

### Philosophical Discussion of Agreement and Disagreement

Within the philosophical discourse on agreement and disagreement, disagreement generally receives far greater attention than agreement, with many authors leaving agreement as an afterthought or considered sparingly and often merely as a contrasting phenomenon. I differ in this regard because I discuss agreement more explicitly and less as a function of disagreement. Nonetheless, disagreement tends to be discussed in three philosophical sub-disciplines:

In social epistemology, the main guiding question is what epistemic import or bearing (if any) does disagreement have on us? If we recognise that there is dissent from our opinion or attitude on some matter, should we be moved to change our attitude or factor in the dissent at all? Much of the discussion focuses on specific instances of disagreement between *epistemic peers* who are, roughly, epistemic equals, synonymous in their intellect, disposition, and capacities, with the prevailing thought being that our best means of determining what is epistemically significant about disagreement is to focus on idealised cases involving epistemic peers. In the second part of the thesis, I address the primary peer disagreement debate as well as underlying debates regarding evidential uniqueness and higher order defeat, but also argue that much of the peer disagreement discourse fails in its aim at demonstrating disagreement’s epistemic significance. I contend that it relies too much on idealisation, yields inconclusive and largely overdemanding replies, and neglects any prospect of peer agreement altogether, as well as the substantially more pervasive and typical instances of agreement and disagreement between non-peers.

Philosophers of language also have a longstanding debate between objectivists, relativists, contextualists, and occasionally expressivists about language, over the means of explaining certain (apparent) cases of disagreements involving claims that do not appear to have a straightforward truth value, such as differences of preferences and tastes. These look like disagreements, but it is far from clear what makes them disagreements as opposed to (mere) differences in preference or taste; claims that something is “tasty” could mean one of at least three different things. It could refer to an objective, perspective-independent property of tastiness which the speaker has identified. For others, “tasty” has a meaning and truth value relative to the speaker or a group to which they belong, or to the conversational context in which it is uttered. Others still maintain that it is simply one’s expression of their enjoyment in some way that would otherwise have been non-verbal. The objectivist can explain away disagreement

(or its appearance) as the result of mere differences of beliefs or speakers failing to comprehend the facts. However, those who endorse relativism, contextualism and expressivism cannot explain why differences of taste, humour, or moral appraisal, are or appear to be disagreements, which is typically called the problem of *lost disagreement*. I take no particular side on this debate, and largely do not engage with it except when doing so yields particular accounts of disagreement. My interest is not over whether we can reconcile the appearance of disagreement with underlying commitments about truth, knowledge, and language, but about why agreement and disagreement cannot and do not admit of any unified kind or common factor.

Finally, disagreement over moral matters has occupied much of the 20<sup>th</sup> century's discourse in moral philosophy and metaethics. Many accounts of what moral disagreement consist in can be attributed to C.L. Stevenson's distinction between disagreement in belief (over non-moral facts) and disagreement in attitudes, particularly moral appraisals. Subsequent writings have attempted to accommodate disagreement in other non-doxastic states and attitudes, suggesting that we can legitimately be regarded as disagreeing when we differ in desire-like states, plans, intentions, or prescriptions and recommendations. As with the debates in the philosophy of language, many of these accounts assume that at least some attitudes do not admit of objective truth-values and are of a non-cognitive, non-doxastic nature; subsequently, the relevant accounts of disagreement are regarded as separate from disagreements in belief. Mark Schroeder describes non-cognitivism in ethics as maintaining that "moral thoughts are not of the same kind as ordinary beliefs – at least ordinary *non*-moral beliefs", and that "moral thought and language are more intimately connected to motivation than non-moral thought and language".<sup>1</sup> Anyone who believes moral language and speech do not express beliefs or related states must reconcile this conviction with the fact that disagreement (including moral disagreement) typically seems to involve beliefs and related states. Again, while I discuss the various accounts of disagreement that have been offered within the moral and metaethical literature in the first part of this thesis, I am not aiming to offer a scholarly treatment of the moral disagreement literature specifically. This is not a thesis about metaethics, but about agreement and disagreement. I survey the various accounts of disagreement and show that they capture the cases they are concerned with but each fail to explain a wide variety of alternatives.

---

<sup>1</sup> Schroeder, M. 2010; 12-13.

## Thesis Structure

This thesis is divided into two parts, each dedicated to a different discussion regarding disagreement, that do not typically overlap or interact. In the first part, I argue against a vast community of scholars who themselves are not in accord or consensus overall. There is a great deal of disagreement about what disagreement is and is not, both generally and specifically; I contend that we can very easily make sense of specific cases or kinds of disagreement, but that does not help to clarify what disagreement is or involves *in general*. Many different explanations and accounts have been offered, and accounts of doxastic agreement and disagreement that focus on coarse-grained, full belief states, are often the first or chief examples. I contend that agreement and disagreement are far harder to define or specify, that they are complex and multifaceted phenomena, with several factors or components that shift and vary between cases. The upshot of this is that trying to determine necessary and sufficient conditions for the phenomena of agreement and disagreement is a fool's errand, and that we are better off thinking of them as ambiguous concepts. I dissect the various factors or components of agreement and disagreement that I believe fluctuate and differ between cases so much that it precludes any possibility of a precise, all-encompassing definition of agreement and disagreement.

Chapter 1 opens my discussion with a few (purportedly) simple views about what it is to disagree, which largely emphasise differences or incompatibility between doxastic states. That is, disagreement in belief. I discuss some important and widely discussed distinctions between the force and contents of speech acts, various kinds of inconsistency, and full belief states vs. degrees of confidence or credences, all while raising various counter examples to the simple views throughout. These various examples indicate that at least some agreement and disagreement cases involve non-doxastic states and attitudes, which is my focus in the majority of Chapter 2, with the latter sections concerning less obviously state or attitude-focused accounts of disagreement. Finally, Chapter 3 addresses the matter of what relations are involved in cases of agreement and disagreement, we might have in mind particular terms with layfolk or technical connotations (e.g., contradiction, synonymity, consistency), so I disambiguate and clarify what each relation involves. I also address the matter of which types of speakers or agent with whom we may engage in agreement and disagreement, because I believe this bears at least on how we may determine the significance or import of cases of agreement and disagreement, if not their very occurrence, and that there are subtle differences between some types of speakers that have been underestimated. Having clarified my claim that agreement and disagreement have various fluctuating and shifting factors, and that focusing on specific, narrow aspects overlooks and crowds out other plausible agreement and disagreement cases, I discuss

alternative models of concepts to necessary and sufficient conditions. I conclude that the concepts of agreement and disagreement are philosophically ambiguous: They pick out several different kinds and cases and thus have multiple meanings.

The second part of the thesis addresses the discourse on peer disagreement which is less concerned with what disagreement is or when it occurs, and chiefly with why it matters, epistemically. Chapter 4 is largely introductory and expository. I introduce my discussion of the peer disagreement discourse by summarising what I think it aims to do and where its focus lies. I regard the peer disagreement discourse as aiming to demonstrate what the epistemic bearing, significance, and import is, of disagreement, by focusing on hypothetical disagreement cases between epistemic peers, or speakers who are symmetrical in various regards that are relevant to their belief-forming. I also introduce two further debates which I consider fundamental to understanding the peer disagreement discourse. One concerns the uniqueness or permissiveness of evidence, whereas the other is about higher order evidence and evidential defeat. In Chapter 5, I discuss how these debates overlap and criticise some of the assumptions that are typically made in the peer disagreement discourse regarding evidence.

Over Chapters 6-8 I address the peer disagreement discourse directly. In Chapter 6, I present the most discussed viewpoints in the discourse, which are all chiefly claims about how we rationally ought to respond once it is established that we disagree with our epistemic peer (or with someone we perceived as a peer before establishing disagreement.) I contend that these arguments are largely unconvincing, and that they either lead to widespread dogmatism or suspended belief, or simply demand too much for rationality. The upshot of this is that the peer disagreement discourse fails to explain or establish what is epistemically significant about disagreement and overlooks at least one plausible account of what makes disagreement epistemically important. In Chapter 7, I consider the reliance on idealisation in the discourse and, ultimately, the disproportionate and insufficient consideration of peer agreement. The peer disagreement discourse is premised on an acknowledged idealisation regarding the prospects for epistemic peerage; there are many conditions for peerage which are rarely met collectively and are certainly a great challenge to meet individually, and most authors who write about peer disagreement likely think this but do not explicitly say it. I raise the concern that the overreliance on idealisation in peer disagreement does not prove especially fruitful, and in fact may generate more bad outcomes than good. I also contend that the discourse and the wider epistemology of disagreement fails to consider any prospect of peer agreement, with discussion of agreement confined to cases involving clear group consensus or expert opinion. I then extend this final discussion in Chapter 8, I argue that the focus on disagreement between epistemic peers is

excessive, and that realistic cases of agreement and disagreement between speakers who fall short, either slightly or substantially, of being epistemic peers, are overlooked. We will find that ‘non-peer’ and certain related descriptors are vague and imprecise terms, and that there are many kinds of epistemic asymmetry that interact with agreement and disagreement in different ways and raise unique implications and questions.

### Clarifications and Caveats

A few points are needed at the outset. First, I have in mind two narrow sorts of accounts of agreement and disagreement. The first are action-based definitions to the effect of “two agents agree if and only *iff* they  $\Phi$ ” and disagree “if and only *iff* they do not  $\Phi$ ” or “if and only *iff* they fall short of  $\Phi$ -ing.” The second are relational accounts that argue something to the effect of “two agents agree or disagree in virtue of having particular states or attitudes A and B which are related in some logical or discursive manner.” There are plenty of examples of both, and some accounts may fit into both categories, but no single account will accommodate all cases of agreement and disagreement. There may be some instances which fit with the two kinds of account I mention above but are nonetheless not obviously or intuitively agreement or disagreement. Many authors have either argued this or alluded to it, some more explicitly than others, but I make the following additional claims over the course of Part 1:

1. There are various changing factors or components that make it nearly impossible to identify common features or similarities between cases of agreement and disagreement.
2. Subsequently, it is not possible to adequately specify or discern necessary or sufficient conditions for agreement and disagreement, and the many accounts available are each too narrow to capture all cases.
3. Therefore, the phenomena of agreement and disagreement are best understood as philosophically *ambiguous* concepts, which admit of multiple definitions or examples.

Second, at times I will discuss concepts or notions that appear to be synonymous with agreement or disagreement, other times I will focus on particular factors that have been proffered as *conditions*.

For example, ‘incoherence’, ‘incompatibility’, and ‘contradiction’ can all be regarded as synonyms for disagreement if they are construed or conceived in the right way but also are treated and sometimes named as conditions for disagreement, whereas beliefs are only ever regarded as conditions. To see what I mean, consider the phrase “all disagreement is just an incompatibility of beliefs.” We can imagine someone picking out the concept of disagreement with the term incompatibility, but there needs to be a specification of what is incompatible which, in this case, are beliefs. Other candidate states are regarded as apt for disagreement, such as appraisals, tastes, or various desire-like states. Again, many of the accounts I consider aim (I believe) at identifying necessary and sufficient conditions for disagreement, and sometimes for



agreement also, but at best they only pick out conditions for some kinds of agreement and disagreement, not for all. If we cannot capture all cases, instances, or ideas of agreement and disagreement with at least one necessary condition, then we surely cannot define agreement and disagreement as unified kinds, hence why my account of agreement and disagreement is broad.

Third, much of the philosophical discussion of disagreement is confined to the perspective of certain areas or domains. In claiming that agreement and disagreement are not reducible to necessary and sufficient conditions, I have in mind agreement and disagreement in general, rather than in any specific, singular domain such as politics, ethics, religion, or science. One might regard what I am doing as a kind of conceptual engineering, but I make no claim about this. Moreover, I do not regard myself to be arguing for or against any robust facts about agreement or disagreement.<sup>2</sup> I do not maintain that we can know when agents, speakers, or parties agree or disagree in any particular circumstance, but neither am I advocating for scepticism or error theory about these phenomena. Ultimately, my aim is to try to make sense of disagreement more comprehensively than I believe the current discourse and literature has achieved so far. A comparison may be drawn with the value pluralism Isaiah Berlin presents in *The Pursuit of the Ideal*. For Berlin, it is a natural occurrence that humanity will be divided according to the values promoted and prioritised by different collectives or communities. We are free to critique and condemn each other for having different priorities to ourselves, but we ultimately cannot place one particular value, such as liberty or equality, as supreme and coerce all others into pursuing it:

What is clear is that values can clash – that is why civilisations are incompatible. They can be incompatible between cultures, or groups in the same culture, or between you and me. You believe in telling the truth, no matter what: I do not, because I believe that it can sometimes be too painful and too destructive. We can try to discuss each other's point of view, we can try to reach common ground, but in the end what you pursue may not be reconcilable with the ends to which I find that I have dedicated my life. Values may easily clash within the breast of a single individual; and it does not follow that, if they do, some must be true and others false.<sup>3</sup>

Berlin says no ideal value can ever outclass all others, and no priority can supersede all alternatives or provide an objective metric with which to judge communities, collectives, or individuals. I argue that no narrowly specified or simple account of agreement or disagreement

---

<sup>2</sup> See Bex-Priestley and Shemmer 2017; 198-9.

<sup>3</sup> Berlin, I. 1988; 10.

can be regarded as outclassing the various alternatives, let alone as capturing all kinds and cases of agreement or disagreement. Once again, the most general thing we can say is that they are complex, multifaceted, and ever shifting phenomena.

### Lay Summary

Agreement and disagreement are far harder to understand, define, and explain, than we (philosophers or layfolk) might think; they are complex, multifaceted, often shifting and varying phenomena that are not clearly reducible to a single condition or set of conditions. Moreover, agreement has received substantially less attention in philosophical discourse than disagreement.

Disagreement talk in philosophy is generally divided between sub-disciplines: moral philosophers dispute whether moral disagreement is an issue for moral realism and offer various accounts of disagreement in non-belief states. Philosophers of language question whether certain instances of disagreement appear legitimate despite involving no truth-evaluable states. Finally, the epistemology of disagreement, the most prominent and most recently developed discussion about disagreement, asks whether instances of disagreement have any epistemic import or bearing on our judgments and attitudes. The consensus is largely that the significance is discernible by focusing on idealised cases of *peer disagreement*.

This thesis is divided into two parts, the first part raises the prominent question of what *are* agreement and disagreement, both generally and specifically? I consider various notable and intuitive accounts and instances of agreement and disagreement, many of which are or involve technical concepts of states, relations, and speakers. I contend that the complexities of agreement and disagreement are too great and the variables too many to define them with a singular and *narrow* account or by emphasizing one particular aspect. I offer a comparatively *broad* account that is pluralist in nature; I do not rule out any one account as illegitimate, rather I embrace them all as parts of a whole, contending instead that agreement and disagreement are best regarded as ambiguous concepts.

The second part tackles the peer disagreement discourse, dissecting its premises and challenging many of the assumptions it has. I raise what I regard to be many issues for the discourse, including that it relies too heavily on idealisation, that it neglects agreement and disagreement between a variety of non-peers, as well as overlooking peer agreement altogether, but also that it simply fails to deliver its two aims. Focusing on peer disagreement tells us nothing about whether and why disagreement is epistemically significant, nor are any compelling replies offered within the discourse. I conclude by drawing together the two strands of thought and proposing an account of agreement and disagreement as possibly grounded in testimony.

Part 1: Agreement and Disagreement: Broad, Complex, and Indefinable

## Chapter 1: Simple and Doxastic Agreement and Disagreement

It is, perhaps, a prevailing assumption, or even intuitive, that disagreement typically looks like a case of two parties believing or assenting to conflicting propositions. That is, we disagree when, for example, you assent to a claim and I deny it, or when you believe something is so, is true, or is apt or fitting, and when I in turn believe that same thing is not the case, or that it is false, wrong, or irrelevant. If we each hold a belief that could not be held or affirmed simultaneously, we disagree. In other words, disagreement might be thought of as something like the following:

In the most straightforward sense, two individuals disagree about a proposition when they adopt different doxastic attitudes toward the same proposition (i.e., one believes it and one disbelieves it, or one believes it and one suspends judgment).<sup>4</sup>

A firm and decisive stance seems core to this and to other intuitive conceptions of agreement and disagreement, and while one makes no firm or decisive *positive* claim when one suspends judgment, that need not mean that no firm or defensible claim is made whatsoever. Judgment suspension might not involve belief or disbelief, and while it certainly appears different from (straightforward, intuitive) cases involving belief and disbelief, there are certainly plausible cases of agreement and disagreement in which one speaker suspends judgment. In contrast, the received view might be that two speakers *agree* when they share the same propositional attitude about some matter, i.e., when two agents share beliefs that are near identical or that otherwise mean the same thing.

We can point to a range of situations where we are conversing with a speaker, hearing a report of some third party's utterance, or reading an assertion, and say with some confidence that it is an instance of agreement or disagreement. But we cannot always say why something is an agreement or disagreement or, more specifically, what makes a case of agreement an agreement, and what makes a case of disagreement case a disagreement. I am not alone in making this observation:

In these discussions, theorists have been trying to capture a notion of disagreement that conforms to intuitions or everyday judgments as to what it takes for two people “really” or “genuinely” to disagree... I am not going to try to say what it is genuinely to disagree. I am unsure as to whether there is any clear, univocal pretheoretical notion of disagreement that it makes sense to try and analyse or account for in this way.<sup>5</sup>

---

<sup>4</sup> Frances, B. and Matheson, J. 2019; §1.

<sup>5</sup> Kölbel, M. 2014; 101-2.

It is typical, of those who try to explain what disagreement is or involves, that they start with an argument like the following, from John MacFarlane:

The Simple View of Disagreement. *To disagree with someone's belief that p is to have beliefs whose contents are jointly incompatible with p.*<sup>6</sup>

Other accounts of disagreement have been offered that are also regarded as “simple”, and I will address them as we go along. For now, there are two aspects of an exchange that MacFarlane emphasizes in this version of the Simple View. The first is the *contents* of two or more speakers’ beliefs, and the second is the fact that the contents of the relevant beliefs are *incompatible* with each other. Briefly, I read MacFarlane as understanding ‘contents’ to refer to the semantic or pragmatic content of one’s attitudes, or the meaning they attribute to their utterance and its constituent words. However, an important distinction can and should be made between the contents of an utterance or attitude and its *force*, which refers to the nature or intended use of that utterance or attitude. In other words, the contents of a speaker’s utterance indicate what the speaker intends for them to mean or say, and the force of a speaker’s utterance indicates what they intend to do with those words or that utterance (e.g., with the same words one may hope or desire, describe, assert, or demand.) On MacFarlane’s view, disagreement in belief occurs when the contents of one’s belief are jointly incompatible with the contents of another speaker’s belief. The relevant notion of force is already indicated because MacFarlane specifies disagreement *in belief*. The crux of this reading of the Simple View is that, for any two or more speakers, those speakers disagree with each other if none of them can adopt each other’s belief without dropping or changing a belief of their own. MacFarlane’s version of the Simple View is explicitly an account of *personal* disagreement (“I disagree *with you*”), but agreement and disagreement can also be (merely) propositional (i.e., “I agree *with your view*”); only the former requires that speakers be aware of each other or in dialogue. Moreover, agreement and disagreement accounts can also be differentiated according to whether they refer to the state of *being in* agreement or disagreement, and the activity of *coming to* agree or of having a disagreement. I am largely interested in the former.<sup>7</sup>

Already, it should be clear that there are different aspects to consider or weigh up. In this account alone, we have a specification of disagreement *in belief*, an emphasis on the ‘content’ of beliefs, and a vague condition that there be ‘joint incompatibility’ between those contents. Each of these concepts, factors, or elements are (purportedly) fundamental to understanding and

---

<sup>6</sup> Macfarlane, J. 2014; 121.

<sup>7</sup> See Cappelen, H. and Hawthorne, J. 2009; 60.

therefore to identifying (some) disagreement, let alone what agreement may require or imply. I will elaborate about the contents and force of beliefs shortly, but for now I will simply say that MacFarlane's version of the Simple View does not capture all instances of disagreement. Some intuitive disagreement cases involve a relation between states that aren't beliefs, and sometimes beliefs can have jointly incompatible contents without indicating disagreement. In this first chapter, I will introduce various other accounts of disagreement that have been described as "simple" and aim to show that they do not sufficiently capture all cases of disagreement, nor can they be adapted to explain or describe all cases of agreement. I will do so by suggesting various examples that contrast with the various accounts I discuss, along with certain distinctions and technical concepts.

What MacFarlane has in mind by beliefs having jointly incompatible contents is not a matter of whether either of the two beliefs in question is true or false, but that they cannot be held together at the same time. This differs from technical, logical relations about propositions and belief states such as contradiction and inconsistency, but it will be helpful to have at least an idea of these notions in mind as well, so the following can be regarded as strictly working definitions:

*Contradiction*: belief in something necessarily false or simultaneously believing and not believing.

*Inconsistency* (of propositions): in any given set of propositions, it is not possible for all of them to be true because one proposition contradicts one or more of the others.

And, for the purposes of envisaging agreement cases, I assume the following as corresponding:

*Sameness*: identity of beliefs, either in terms of their semantic or pragmatic contents.

*Consistency* (of propositions): in any given set of propositions, it is possible for all of them to be true at once, no one proposition contradicts any of the others.

While the relations of logical consistency and contradiction can help us to understand agreement and disagreement, truth is not a necessary condition or a determining factor for whether two speakers agree or disagree in belief. Agreement and disagreement are neither exclusively logical relations, nor are they exclusively discursive, two speakers can agree on some matter and share a belief that is false, and likewise there can be disagreement in states that are not truth-apt. Again,



personal agreement and disagreement relate to speakers' beliefs or attitudes, whereas propositional disagreement concerns the attitude itself or its contents.

I may have gone too far, too fast, so let us come back to MacFarlane's version of the Simple View, on which two speakers disagree in belief when they each have a belief or beliefs whose contents are jointly incompatible. One cannot hold both beliefs simultaneously and would have to revise or relinquish one of them. Something like this will seem very familiar, and it holds true for many cases of disagreement across a variety of domains and subjects, but it still needs to be deconstructed. Firstly, what do we mean by the contents of a belief? In the most plain, simple terms, the contents are what one says when one performs a speech-act: the "bare, forceless representations that are put forward in different ways".<sup>8</sup> They differ from the *force* of a speech act, which is essentially how something is said or the way in which one intends to convey something. Believing, wondering, and desiring are all very different things, but they could all be performed with the same shared contents in mind, such as "that it will snow later today." One further distinction is worth noting as well, since there is at least some discussion among philosophers of language over whether the distinction between semantic and pragmatic contents is legitimate, and what it refers to:

... we say that the semantics-pragmatics distinction is, for example, the distinction between what is said (illocutionary sense – in the cases traditionally focused upon: what is asserted) and what is meant (illocutionary sense – in the cases traditionally focused upon: what is communicated) ...<sup>9</sup>

This distinction between semantics as 'what is said' and pragmatics as 'what is meant' is a simple, but common account of the difference between the contents of a speech act, although other distinctions that could be said to apply are also available. For the sake of my discussion here, it need not matter all that much whether one endorses this over any other distinction, but it suffices to highlight another complication in establishing disagreement.

Suppose the contents of my belief that "*p*" and your belief that "not-*p*" differ, in such a way that it looks as though we disagree but in fact our beliefs have contents that facilitate agreement. Now suppose instead that we each make an assertion, and both say the same thing, but we each have in mind a different understanding of what we mean, or that we each associate our wording with different connotations. For example, if you and I were to both say, "It is sunny outside", one might think we agree because we have said the same thing. If, however, we do not

---

<sup>8</sup> Hanks, P.W. 2007; 142.

<sup>9</sup> Szabó, Z.G. 2005; 6.

mean the same thing or misunderstand each other (e.g., I mean it is incredibly warm today, you mean it is very bright) then either we agree in a very limited sense, or we are just talking past each other. There need not even be misunderstanding, miscommunication, or divergence in what we intend to say and mean. Perhaps we actually disagree over what it is to be “sunny outside” and this is unexpressed or unacknowledged.

This should start to make the differences clear. Insofar as agreement and disagreement differ from each other, it is not enough to say that disagreement is just “not-agreement”, and agreement is “not-disagreement.” As we have seen, sometimes what speakers say is neither agreement nor disagreement, but in fact leads to some terminological difference. Some might regard cases like this to be splitting hairs or playing semantics, which does not do much except trivialise the exchange and obscure the mere appearance of any agreement or disagreement. However, not all exchanges fall into one of these three categories: speakers can believe or assert things without any relation that indicates agreement or disagreement but are also not sufficiently relevant to each other to count as a terminological difference. If I were to say, “It is really hot today” and you were to respond by saying “There’s a cat on the roof over there” we are just making inane comments that are not relevant to each other, but do not clash either. Likewise, if two agents are in a conversation that results in them questioning the same subject or matter together, they do not clearly seem to be agreeing or disagreeing, but what they are doing is terminologically relevant and, presumably, sufficiently semantically and pragmatically clear to rule out irrelevance.

As such, there are a wide range of categories that conversation could fall into, with some looking like agreement and disagreement cases and others not. I have clarified the difference between the force and contents of a speech act and given what I consider to be enough of an explanation between semantic and pragmatic contents to show that sometimes we can appear to agree or disagree but not actually do so. However, this does not address MacFarlane’s iteration of the Simple View completely; I have not elaborated on the relation he emphasises (joint incompatibility, and to some extent non-cotenability). Moreover, I have not yet addressed the focus on doxastic states or full belief states, which I turn to now.

### 1.1: Agreement and Disagreement in Belief

So far, I have only considered a picture of agreement and disagreement that focuses on belief, disbelief, and perhaps suspension or withdrawal of belief. If it is not possible to hold two beliefs at the same time (i.e., if they are logically contradictory), then two speakers who hold one of those beliefs each are in disagreement. For example:

**Water A:**

“I’m quite thirsty.”

“Why not have some water? It’s important to drink water regularly.”

“You’re right, I’ll pour a glass now.”

**Water D:**

“I’m quite thirsty.”

“Why not have some water? It’s important to drink water regularly.”

“It isn’t that important to drink water.”

The latter case fits just fine with the version of the Simple View offered by MacFarlane, since one cannot both affirm and deny the importance of drinking water, nor can one suspend judgment about it while believing or disbelieving that it is important. But also, it looks like a case of disagreement, and I suspect very few people would say otherwise. Again, agreement and disagreement come apart from the truth from a proposition and can occur regardless of whether one view is ultimately factually true or correct. The point is that some instances of disagreement do look like this, where it would be contradictory to hold two belief states about the same proposition at the same time. However, not every instance of disagreement looks like this:

**Tea/Coffee:**

“I’m putting the kettle on. Would you like anything?”

“I like tea.”

“I prefer coffee.”

Earlier I noted that speech acts and language can have different contents, but also different *forces*; in other words, it has various functions and speakers use it to do different things. The example above is a case in which two speakers are merely reporting their preferences, indicating their preferred beverage without any attempt at judgment or persuasion, unlike the following version in which language is used *expressively*:

**Tea/Coffee 2:**

“I’m making a cup of tea, are you thirsty at all?”

“Ooh, I’ll pass on tea, thank you. But if you have any coffee, I would love some.”

In this instance, it is not the case that both speakers are reporting a preference or desire, instead the first speaker is making an enquiry, and the other is expressing a preference in response. We may infer that the two differ in preferences, but it is not explicit. Nonetheless, the two preferences for coffee and tea cannot be held simultaneously, one cannot prefer tea over coffee *and* coffee over tea at the same time, without being irrational, that is. However, the first speaker’s preference for tea does not prohibit or preclude the other speaker from preferring coffee nor from satisfying either (or both) preferences.

Language need not only be used expressively, but we also use it for *argumentative* purposes as well by justifying, defending, critiquing, and persuading, as is the case in the following version:

**Tea/Coffee 3:**

“I’m making a cup of coffee; would you like one?”

“Ooh, too much coffee is really bad for you, do you have any tea instead?”

We can infer disagreement in this case from the fact that one speaker desires a coffee and the other expresses an appraisal regarding the coffee consumption as well as a contrary preference in the same utterance. The disagreement is plausible, therefore, but not explicit. The two speakers might agree that consuming too much coffee is bad, and this may inform or come apart from one speaker preferring and/or desiring coffee over tea, which rules out disagreement in some sense. If they differ in preference *and* on the badness of excessive coffee, then MacFarlane’s version of the Simple View captures disagreement in one sense, but not in another. As we will see, I contend that different kinds of agreement and disagreement can overlap without being directly causally related. Because enquiries, appraisals, and preferences are all different kinds of states from each other and from beliefs, MacFarlane’s version of the Simple View may not fully explain cases like Tea/Coffee 2 and 3.

You might disagree with this analysis for some reason or another. Perhaps you are firmly committed to a strong cognitivist realism that posits negation or belief-incompatibility as essential for disagreement, and deny that differences in preferences, desires, appraisal, or taste count as real disagreements. We could reformulate examples such as the ones above so they are all disagreements in belief, but I do not believe it would be entirely clear what belief the disagreements would be over in such cases; the speakers in Tea/Coffee 2 are not obviously asserting beliefs. Perhaps one speaker could straightforwardly deny that the other prefers tea over

coffee, which would fit with MacFarlane's Simple View, but at the same time looks strange and unintuitive. My aim is not to offer a fully thought-out defence of the legitimacy of agreement and disagreement based around non-doxastic states, because doing so will divert attention away from accounts and factors of agreement and disagreement. I will assume that such states are apt for relevant sorts of agreement and disagreement, and I consider various accounts in the next chapter. Regardless, MacFarlane's version of the Simple View is, perhaps, too simple, and only accommodates some of a wide variety of plausible disagreement cases, without even considering agreement.

### 1.2: Inconsistency, Contradiction, and Contrariety

Other formulations of the Simple View or, rather, other accounts of disagreement that are named the "Simple View" are available. Yonatan Shemmer and Graham Bex-Priestley note the tendency to focus on disagreements in belief and the scarcity of concise accounts or formulations of disagreement:

Clear formulations of the simple view are rare. To the best of our understanding, this is in part because there are very few discussions of disagreement altogether, and in part because many philosophers took it for granted that something like the simple view is correct and therefore didn't need to provide an exact analysis of disagreement. On the simple view, disagreement is a relation that holds between beliefs or between people in virtue of their beliefs.<sup>10</sup>

They then argue that the simple view is better understood as a collection of accounts of disagreement which all present disagreement as some impossible combination of beliefs. In other words, any account of disagreement may be regarded as simple if it posits that disagreement involves agents holding or asserting beliefs that:

- i. Negate each other.
- ii. Have inconsistent contents.
- iii. Cannot both be true.

However, they also note that wherever some author has offered an account that fits within the family of simple views, acknowledge the possibility of disagreements in states that are not belief-

---

<sup>10</sup> Shemmer and Bex-Priestley 2021; 495.

based, and that many of them argue their simple views can accommodate these cases. Disagreements in taste or preference, such as the latter two versions of Tea/Coffee above, are derived from a disagreement in belief over what the right taste is, and the only true disagreements are belief-based.<sup>11</sup> I have discussed negation as a condition or basis for disagreement already, since it is just the same as affirmation or assent vs denial, belief vs disbelief, these can all be used interchangeably, the idea is just that speakers disagree *iff* one speaker believes that  $p$  and the other believes “not- $p$ ” (or some equivalent proposition that entails “not- $p$ ”).

The second factor named by Shemmer and Bex-Priestley is that the contents of beliefs must be inconsistent with each other to facilitate or signal disagreement. Importantly, there are different kinds of inconsistency relations that might be said to hold of a proposition or belief. They can be *factually* inconsistent, which occurs when they are inconsistent with the facts and demonstrably false. The other kind of inconsistency, which I believe is what Shemmer and Bex-Priestley are referring to, is *logical inconsistency*, which will require some elaboration. Inconsistency is typically said to hold of propositions together in sets or groups, in which at least one of the propositions within the group generates a contradiction. Most classical, conventional, or traditional systems of logic follow Aristotle’s Law of Non-Contradiction:

*For the same thing to hold good and not to hold good simultaneously of the same thing and in the same respect is impossible* (given any further specifications which might be added against the dialectical difficulties. This, then, is the firmest of all principles, for it fits the specification stated. For it is impossible for anyone to believe that the same thing is and is not.<sup>12</sup>

Aristotle maintains that it is impossible to believe a proposition and its negation at the same time, or to hold beliefs of the form “ $p$  and not- $p$ ” Those who reject this, or at least do not regard it as always true, allow that some contradictions can be true and that it is at least sometimes legitimate or permissible to hold beliefs of the form “ $p$  and not- $p$ ”. Ordinarily, any group or set of statements is *inconsistent* if it is not possible for them all to be true (or true in the same world, at the same time), or if the meaning of the statements together requires one of them to be false. For example, if I say, “everyone has left the building” and you say, “Charles is still inside the building”, our two statements are inconsistent with each other because they cannot be true at the same time. Moreover, our statements also indicate that we disagree in the second, ostensibly

---

<sup>11</sup> Ibid; 495.

<sup>12</sup> Aristotle, Christopher Kirwan (transl) 1993; Book  $\Gamma$  1000b18-22

simple sense that is picked out by Shemmer and Bex-Priestley. You do not need to explicitly deny what I say, nor I you, but we nonetheless believe things that are inconsistent and cannot be true simultaneously, so we disagree in that sense.<sup>13</sup>

However, the third sense of disagreement they believe Simple Views pick out is when beliefs cannot be true together, and this is where things can get muddy. I said earlier that Aristotle conceives of contradiction as the state of believing the same thing is and is not (at the same time). Aristotle maintains that there is something amiss, wrong, or otherwise impermissible about believing and denying any given proposition at the same time, and we use ‘contradiction’ or ‘contradictories’ to describe these sorts of propositions. It is sometimes all too easy to confuse inconsistency with contradiction; in the simplest possible terms, “ $p$  and not- $p$ ” is a contradiction, whereas “ $p$ ” and “not- $p$ ” are, in at least some cases, inconsistent. Out of these two propositions, “ $p$ ” and “not- $p$ ”, only one can ever be true in any possible time and world. However, since one might think that holding a conjunction of beliefs is the same as believing that conjunction, this is why it might be easy to mistakenly identify contradiction and inconsistency as synonymous. As John N. Williams notes:

If a man believes that  $S_1$ , believes that  $S_2$ , and believes that  $\sim (S_1 \& S_2)$ , then he does not hold contradictory beliefs, as he would were he to both believe that  $(S_1 \& S_2)$  and believe that  $\sim (S_1 \& S_2)$ . Nor does he hold a self-contradictory belief, as he would were he to believe  $[S_1 \& S_2 \& \sim (S_1 \& S_2)]$ . He does not hold contradictory beliefs, since there are no two beliefs which contradict each other. Moreover, the subtraction of any one of his beliefs has the result that it is now possible for all his remaining beliefs to be correct.<sup>14</sup>

Importantly, Shemmer and Bex-Priestley appear to have confused inconsistency with something else. They make no mention of contradiction in their presentation of Simple Views, so perhaps they have confused contradiction with inconsistency or believe that inconsistency picks out some relation other than beliefs not being true simultaneously. To be clear, we have contradictory beliefs when the truth of one belief necessarily entails that the other is false, or when we believe a proposition and its negation together at the same time. Recalling the earlier example, if you believe that Charles is in the building and I believe that he is not in the building, our beliefs are

---

<sup>13</sup> It is not specified whether inconsistency must be a fact of the matter, or if disagreement can arise simply because one or both speakers perceive their beliefs as inconsistent. My earlier point about terminological disputes and apparent disagreements might lend credence to saying inconsistency must hold for some disagreements, since we might otherwise be mistaken in thinking our beliefs or assertions are inconsistent.

<sup>14</sup> Williams, J.N. 1981; 600.

inconsistent because they cannot be true together, but they *contradict* each other because our respective beliefs have contents (propositions) that negate each other. One must be true and the other must be false.<sup>15</sup>

However, and importantly, this is not the same as when our beliefs are *contrary*, which is a weaker form of inconsistency that holds of beliefs that could both or all be false. If I believe everyone has left the building and you believe that Charles is still inside, both our beliefs could be false, because Charles might have not been in the building in the first place; they cannot both be true, so together they are inconsistent, but they could both be false, so they are contrary. Using the same example to illustrate these three different, but similar relations might be confusing, so let me offer another example. It is impossible for it to be Winter and Spring at the same time, but it could be neither Winter nor Spring, so “Winter” and “Spring” are contrary. “Winter” and “Not Winter” are contradictory, since it must be one or the other, it is impossible for it to be Winter and Not Winter at the same time, just as it is impossible for it to not be Winter and Not Winter at the same time.<sup>16</sup> In summary, we have the following technical definitions:

*Inconsistency* (of beliefs): cannot be true at the same time or together.

*Contradiction*: simultaneously believing and not believing a claim.

*Contrariety* (of beliefs): could be false at the same time or together.

If we take MacFarlane’s version of the Simple View as referring to disagreement as the inconsistency of beliefs in the above sense (i.e., the impossibility of two beliefs being true simultaneously), then cases like Tea/Coffee 2 still do not fit. However, we could more charitably read MacFarlane as interpreting inconsistency in a broader sense that encompasses contradiction and contrariety. Likewise, MacFarlane may not be an all-out advocate of disagreement exclusively being *in belief*. In this case, Tea/Coffee 2 still does not look straightforwardly like a disagreement in belief, but it does look like a disagreement in preference, and Tea/Coffee 3 could be considered a disagreement based on each speaker’s attitude being contraries.

---

<sup>15</sup> Other conceptions of contradiction include it applying to objects that are mutually exclusive and exhaustive, or to things that differ in both quality and quantity (e.g., something cannot be simultaneously complete and partial), or belief in something that is necessarily false or that one knows to be false.

<sup>16</sup> See Beziau, J.Y. 2016; 158.

To my mind, any sets of properties with ranges or scales and ambiguous midpoints make for interesting examples. One cannot reasonably be both old and young at the same time, but there is perhaps an ambiguous point in life where one is *neither*. We could likely say the same of any properties with ranges and ambiguous or grey areas (e.g., “tall” and “short”, “happy” and “sad”).



The Simple View that Shemmer and Bex-Priestley discuss makes no mention of what it is for agreement to occur, but we can suppose that there are corresponding simple views available. Perhaps speakers agree in belief when one belief does not negate the other, when beliefs are consistent in content, or when they can both be true together. The opposite of inconsistency, as I present it above, is the mere capacity to be true at the same time but, as I have noted, a degree of relevance is required since beliefs that can be true together can also be inane and unrelated, as we saw above with my observation that it is hot today and yours that a cat is on the roof. The same goes for belief pairs that lack any negation, since “*p*” and “not not-*p*” may not always mean the same thing, the latter could mean “possibly *p*” or “probably *p*”, or something to that effect.

### 1.3: States of Tension and Changes of Mind

So far, I have discussed two different accounts of disagreement that purport to be “simple” views, which have yielded three different respects in which beliefs and belief states can differ or clash in such a way that facilitates, enables, or signals disagreement. Disagreement, on these accounts, occurs when beliefs or their contents are incompatible or inconsistent in some way, when they otherwise cannot be held simultaneously. However, these simple views are incapable of ruling out the possibility of disagreements which involve various non-doxastic states, and likewise they fail at accommodating and explaining them. Subsequently, more needs to be said about disagreement; any plausible account should include intuitive cases of non-doxastic disagreement *and* be able to explain when and how they arise, and likewise with agreement. I note that, so far, I have said very little about agreement and much about disagreement. This is partly intentional because the iterations of the Simple View have proven to be multi-faceted and complicated, but also because there are no formal counterpart accounts of agreement offered by the authors that present them.

I will address agreement directly soon, but first I want to consider one final account of disagreement that is also provided by John MacFarlane, one that is not named as a version of the Simple View, but will nonetheless be helpful for my purposes here:

A disagreement, we might think, is a kind of conflict or dispute. To disagree with someone is not just to have a different attitude, but to be in a state of tension that can only be resolved by one or both parties *changing* their minds.<sup>17</sup>

---

<sup>17</sup> MacFarlane 2014; 123.

The versions of the Simple View above place a relation between doxastic states as primary or central to disagreement. This account differs because it focuses on the relation between attitudes, but it does not explicitly name any attitude as primary, central, or exclusive. This means it has the upshot of allowing for disagreement that involves beliefs, desires, appraisals, and so on. MacFarlane then goes on to note that it is not possible to disagree in this sense by *merely* having two attitudes that cannot be held together, using the example of two agents with desires for each other's company. Such attitudes differ from beliefs in the sense that the attitudes themselves cannot be true, although they can certainly be justified in the sense that beliefs can be justified (i.e., with reasons and/or evidence), and there can be true or false propositions that any given agent or speaker (for example) desires  $x$  or prefers  $y$  over some alternative. For desires, appraisals, or preferences to be true makes less sense than, say, for them to be sincere or genuine. For such attitudes to be satisfied, as I understand MacFarlane, is simply for there to be a state of affairs such that these attitudes are met. That is, there is a world-to-mind direction of fit. As Alex Gregory describes it, beliefs are states which we have reason to change given our actual state of affairs or world, but desires are states that give us reason to change the world.<sup>18</sup> It is not possible to desire to be with Jane more than anyone else while also desiring to be with Bob more than anyone else, the two desires cannot be held simultaneously and seem to be in tension with each other (as was the case with preferring tea over coffee and vice versa.) Importantly, however, if two parties were to hold one of these attitudes each, we might say that the tension can be resolved without either party changing their mind, or simply that there is no tension. If Jane would rather be with Bob than anyone else and Bob feels the same about Jane, their attitudes cannot be held at the same time, but they can be satisfied together. This is why there does not seem to be any disagreement; there would be tension for a single party to hold these two desires, but not for multiple parties to hold one desire each.

Because MacFarlane contends that we cannot meet this account of disagreement merely by stressing the non-cotenability of states or attitudes, he instead suggests that disagreement in (non-doxastic) *attitudes* should be explained by using a proposal from C.L. Stevenson:

The difference between the two senses of “disagreement” is essentially this: the first involves an opposition of beliefs, both of which cannot be true, and the second involves an opposition of attitudes, both of which cannot be satisfied.<sup>19</sup>

---

<sup>18</sup> See Gregory, A. 2012; 604.

<sup>19</sup> Stevenson, C.L. 1963; 2.

Disagreement involving non-doxastic, non-belief states occurs just when we each have attitudes that cannot be satisfied together, or when your attitude precludes mine being satisfied and vice versa. Stevenson offers his account as a way of accounting for disagreements over moral matters, and it is rooted in his expressivist reading of moral language. Recall the Tea/Coffee examples above, with which I noted the difference between merely reporting a preference and expressing it. Stevenson's point is that when we make moral utterances, we are expressing how we appraise some moral act, choice, or judgment, or more specifically, we are expressing whether we approve of it.

Moral judgments are not ordinary propositional attitudes of the sort that we have when, for example, we report that we have a particular dietary requirement or allergy:

**Vegan/Vegetarian:**

"Do you have any dietary requirements the caterers need to be informed about?"

"Yes, I'm a vegan."

"Interesting. I'm not a vegan myself but I am a vegetarian."

As in the first instance of the Tea/Coffee example, this only seems to be a case of reporting propositional attitudes, there is no tension, inconsistency, or negation between the two beliefs, since one can believe that they are a vegetarian and still believe that another person is a vegan. Moreover, both beliefs could be true or false, since both speakers might be speaking insincerely or exaggerating their assertions. The case above could also be rephrased so the states involved are not doxastic:

**Vegan/Vegetarian 2:**

"I'm a vegan because I believe there are strong moral reasons to favour a vegan lifestyle."

"I'm a vegetarian, but I concede that there are moral grounds for veganism also."

Importantly, this is ambiguous since the two speakers could again both simply be reporting that they believe there are moral reasons that justify veganism, in which case there is no disagreement in belief because, as in the first case, there is no tension, negation, denial, or inconsistency. The speakers can be said to agree that "there are moral reasons favouring veganism", even if only one of them is a committed vegan. What about disagreement in attitude? One might think that the difference in these speakers' convictions, dispositions, and actions indicates that they do in fact disagree, because only one has committed to veganism. As with the first version, if we were

to imagine a single agent holding the beliefs that “I am a vegan and I believe there are strong moral reasons for veganism” as well as “I am vegetarian and I believe there are strong moral reasons for veganism”, there would not seem to be any internal tension or incoherence. However, this is just a technicality since veganism is merely a stronger or extended form of the commitments of vegetarianism. All vegans are vegetarians, but not all vegetarians are vegan. Stevenson’s account gives the same results, the beliefs can both be true, and the attitudes can both be satisfied, so they agree in both senses.

#### 1.4: Superficial Disagreement?

It is too easy to say that the accounts I have just considered fail to offer an account of disagreement because they sometimes give us agreement cases instead. Nonetheless, what I have said so far suffices, I argue, to show that there are different instances of agreement or disagreement that can sometimes coincide and come apart. Just because speakers agree or disagree in one sense need not mean that they agree or disagree in all senses. For example:

##### **Vegan/Vegetarian 3:**

“I read Peter Singer’s moral arguments for veganism recently and found them persuasive enough that I have since become a vegan.”

“Fair enough, I’m still unconvinced that the moral reasons for veganism are strong enough.”

In this version of the case, it seems that there are again two different aspects to bear in mind. There is the fact that the first speaker is a vegan, which the second speaker could hypothetically dispute but it would seem a little odd. The crux of things here is whether the moral reasons for veganism, and/or any arguments from Peter Singer, are sufficiently compelling to persuade each speaker to embrace veganism. Importantly, this shows that the written word can be ambiguous in a way that vocalised assertions or speech acts are not; if we imagine two speakers uttering these in conversation, either or both speakers might be implicitly trying to persuade the other to change their mind. This is lost or ambiguous in the written quotations above because each speaker could be *merely* reporting their beliefs and convictions; we are left to infer what each speaker intends.

Regardless, what is clear is that each speaker differs in how persuaded they are by the moral reasons for veganism, so they differ in that sense. Must this be construed as disagreement, though, or even a tension? If one is a vegan but finds the moral reasons for veganism

uncompelling, this need not be odd or inconsistent, since there might be sufficiently persuasive non-moral reasons. Likewise, one could acknowledge the moral reasons for veganism are compelling while not embracing veganism for other reasons. The fact that one speaker is a vegan does not typically interfere with another speaker's vegetarianism, nor vice versa, so there is no tension or implication that one ought to revise their commitments. It would make no sense to say that one finds the moral reasons for veganism both persuasive and not persuasive, assuming one has in mind only the same group or set of moral reasons, but this is a fairly clear case of one agent's self-contradiction. Neither speaker in Vegan/Vegetarian 3 is straightforwardly denying that there are moral reasons for veganism, they simply differ in the degree to which they are persuaded by such reasons.

I contend they do disagree in attitude, but only in a rather superficial or weak sense. Not all differences of states are necessarily tensions, let alone disagreements, and neither do all disagreements necessarily require a change of mind; sometimes we can let disagreement stand and simply agree to disagree, in however weak or substantial a sense it is, without any incivility, hostility, or vitriol. As far as the third version of the Vegan/Vegetarian case goes, I contend that insofar as it is a disagreement, it is only based on a difference in one respect (that we know about). It would be a different, more consistent case of disagreement if one speaker wholeheartedly advocated for veganism on a range of grounds or reasons, and the other did just the opposite. The topics on which we disagree are frequently multi-faceted, complicated, and exist in networks which overlap and are sometimes far-reaching. For example, if we are discussing whether a state should provide aid in a particular foreign conflict, there are several long-standing debates, discussions, and questions that underwrite it. A few examples that come to mind are questions concerning our moral duties and responsibilities regarding others, political dilemmas concerning autonomy, and even legal considerations about foreign jurisdiction.

### 1.5: Agreement and Disagreement are Complex

Many who write about agreement and disagreement acknowledge that many cases involve topics or subjects which are deeply complicated and that are rarely reducible to simple questions with limited implications.<sup>20</sup> Some also note that agreement and disagreement are multi-faceted in nature as well, or that certain isolated accounts do not accommodate all cases or instances. Many who write about agreement or disagreement have in mind some narrowly constrained notion of what agreement and disagreement involve, and while this is fine for capturing or explaining any relevant cases, no single *narrow* account, I contend, can be regarded

---

<sup>20</sup> E.g., Adam Elga: "in messy real-world cases, the disputed issues are tangled in clusters of controversy." 2007; 493.

as all-encompassing. No account of disagreement or agreement will successfully accommodate all cases so explaining all kinds of agreement or disagreement with a specific definition or a set of necessary or sufficient conditions is simply not possible. In contrast to the established literature, my aim is to offer a *broader* account of agreement and disagreement as multi-faceted and complex.

What do I mean by this? Simply that there are multiple shifting factors and components, many moving parts – the states, relations, and agent (or agents) involved in any given case of agreement or disagreement are not uniform or consistent for all possible cases. As we will see further on, I believe the best technical way to understand the concepts and phenomena of agreement and disagreement is in terms of philosophical *ambiguity*, there are no features common to all cases or kinds, so no narrow or simple definition can adequately capture them all. Let us regard any simple account of agreement or disagreement as emphasizing a particular aspect, factor, or condition (e.g., the states involved or the relation between the relevant states.) My approach differs by not presenting one account or condition as ultimately significant, rather it is instead more pluralist, and acknowledges the inconsistency and lack of uniformity between cases. There are components of the way we express or utter our doxastic and non-doxastic states that are not fixed, and different combinations of these components will result in different kinds or cases of agreement, disagreement, if not sometimes mixed cases or non-cases (e.g., merely verbal disputes or terminological differences.) However, to call agreement and disagreement complex and multi-faceted is a rather general and perhaps vague claim. Importantly, there is the risk that an account could be *too broad* and result in misdiagnosing particular instances of commonality or difference as agreement or disagreement. For example, the fact that you and I are both in the state of being human need not mean that there is any concord between us which is relevant to agreement, nor need the fact that I am human but my dogs are not mean that we are in disagreement with each other.

Whether we have in mind cases that involve two agents conversing with each other or simply having their respective attitudes in isolation, or even just one agent comparing their attitude to that of someone else, I believe we can ask the following:

1. What things or *states* are involved?
2. How, or in what *respects*, are these states related (e.g., alike, different, consistent, etc.)?
3. With whom, or what kind of *speaker*, is one engaging or responding to, if any?
4. Are they *aware* of one's agreement or disagreement?

And, perhaps to a lesser extent,

5. What technical terms or vocabulary are available to capture these relations?

I do not believe agreement or disagreement are reducible to a single definition or account (e.g., agreement happens when speakers R and S  $\Phi$ , or when R and S  $\Phi$  in circumstances C; R and S disagree when they  $\Psi$  in C, or  $\Phi$  in circumstances other than C). Conditions of this sort would suffice, but I simply do not believe that they are easily attainable; if they are, they are not enough to capture all cases of agreement and disagreement, given the different factors I have noted.

Importantly, this should not be taken as the conjecture that there are never any real, genuine cases of agreement or disagreement, nor should my claim above about the superficiality of Vegan/Vegetarian 3 be confused for a denial that it is a real disagreement. As Bryan Frances says:

The lesson is this: *even if the disagreement concerns some vague, ambiguous matter that is open to several reasonable yet differing interpretations, there can be genuine disagreement and an absolutely true answer.* This shows how tricky things can get when attempting to discover genuine disagreement.<sup>21</sup>

This captures the spirit of what I am trying to argue, although I am less convinced that every dispute admits of an absolutely true answer. Sometimes the subject matter can be vague or complex, sometimes the agreement or disagreement can be obscure or difficult to determine. Sometimes these can both be the case, other times it can be neither. This does not mean that I will raise various accounts or cases only to dismiss them; my approach is pluralist; one might regard it as contextualist, but I will not explicitly describe it as such. I contend that we can make some broad or general claims about what agreement and disagreement appear to involve, or about what cases seem to look like, but these phenomena cannot be sufficiently captured or described with necessary or even sufficient conditions. Instead, I argue we are better off regarding them as ambiguous concepts, with some overlapping features but none that are universally or uniformly consistent, and therefore definitive. There is no baseline case, instance, or example of agreement or disagreement which we can extrapolate, generalise, or adapt to

---

<sup>21</sup> Frances, B. 2014; 17. Frances later (26) goes on to describe what he calls the “Disagreement Factors”, such as one’s data, evidence, background knowledge, and circumstances. He maintains these “lead people to divergent yet reasonable answers to a single question”, but it is not clear if he regards these as conditions for disagreement itself, or factors that merely influence or direct how one forms their beliefs.

generate other cases. Frances thinks that some agreement and disagreement cases can be vague or indeterminable; I want to extend this thought and maintain that agreement and disagreement in general are ambiguous.

### 1.6: Credal Disagreements and Epistemic Modals

Most of the accounts of disagreement I have discussed so far have prioritised doxastic disagreement where one party believes or assents to a proposition and the other denies or disbelieves it. Such accounts conceive of doxastic states as *coarsely grained*, as involving either full belief or disbelief, perhaps also the suspension or withdrawal of belief. I now turn to accounts that also include *fine grained* notions of belief that incorporate degrees of confidence as a quantifiable state, or *credences*:

Credences are more fine-grained than beliefs and are often given a value on the [0,1] interval, where 1 represents maximal credence  $p$  is true, and 0 represents maximal credence  $p$  is false... Unlike belief, there are (in principle) an infinite number of credences one can take toward a proposition. While the concept of credence grew out of work on subjective probability... many contemporary epistemologists posit a close connection between credence and the familiar notion of confidence.<sup>22</sup>

Disagreement, as it is discussed in the philosophy of language, tends to assume that beliefs are coarse-grained full states, that belief in  $p$  is the state of taking  $p$  to be true and accurate. The epistemology of disagreement embraces a quantitative, fine-grained conception of belief that includes both full doxastic states and credences, although not simultaneously. Propositional disagreements are only conceived as doxastic or credal at any one time. The thought is that two speakers disagree *iff* they each hold different credences regarding some proposition or hypothesis, with each credence corresponding to a different degree of confidence. The higher one's credence, the more confident one is in the truth of the relevant proposition.<sup>23</sup>

As far as disagreement goes, those who endorse credal disagreement as separate from doxastic disagreement maintain that any differences in credence or degrees of belief or confidence constitute a disagreement in that sense. For example, if Jake and James are discussing the weather and Jake has a moderately strong belief in it raining later (i.e., he believes that the likelihood of rain is 80%), and James is fully confident (100% likelihood), we might equate these to credences of 0.8 and 1, respectively, in the belief that it is going to rain later. Those who

---

<sup>22</sup> Jackson, E. 2020; 1-2.

<sup>23</sup> See Palmira, M. 2018; 179-80.



endorse or envisage credal disagreement maintain that this is a disagreement between them, since one cannot hold both credences simultaneously. I consider this disagreement in the same superficial sense as that I believe describes the difference between the vegan and vegetarian differ in Vegan/Vegetarian 3.

Ultimately, both speakers in each of these examples affirm the same respective belief, they simply differ in the extent to which they are confident in that belief, since James is more confident that it is likely to rain just as the vegan is more confident in the moral reasons justification for veganism. However, to be clear, I would regard a more substantial or stark difference between two speakers' credences as a more robust case of credal disagreement, say credences of 0.8 vs. 0.2. Along similar lines, Torfinn Huvenes discusses the worry that credal disagreement predicts too much disagreement:

It is true that insofar as the differences in credences are too small to be noticed, we are not likely to think that the parties disagree. However, that is something that we might be able to live with. After all, there might be all sorts of propositions that we disagree about without anyone recognising that fact. That is true regardless of whether we allow a difference in credences to constitute a disagreement. Furthermore, when the differences are noticeable, it seems natural to think that there is a sense in which the parties disagree. If one is still worried about predicting too much disagreement, a possibility is to say that we have disagreement only if there is a significant difference in credences. What counts as 'significant' will presumably depend on the context.<sup>24</sup>

Huvenes entertains the possibility of credal disagreement being attributed to any 'significant' difference in credences, with significance being context-dependent, but concludes that while this sort of solution is more realistic, it would be difficult to meet and makes it difficult to identify genuine cases of credal disagreement that do not also involve different propositional beliefs.

I am not convinced that minor differences in credence illustrate disagreement in the same way, or at least as substantial or robust a disagreement, as far greater asymmetries in credence. Likewise, two agents believing the same proposition with the same degree of confidence indicates agreement in a rather limited sense. Let me illustrate my point by instead imagining the rain example above in terms of modal propositions instead of credences or degrees of confidence:

---

<sup>24</sup> Huvenes, T.T. 2015; 993.

### **Rain:**

Jake: “It’s probably going to rain today”

James: “It’s definitely going to rain today.”

Put this way, I do not believe it would be self-contradictory if a speaker were to believe both claims, but it would still seem strange to do so, they might appear uncertain or doubtful, or they might be evoking something comparable to a distribution of credences (i.e., a credence between 0.7 and 1 in the likelihood of rain.)<sup>25</sup> Alternatively, one might think that contradiction and contradictoriness are all or nothing relations, so perhaps one can only ever believe either that it is probably going to rain *or* that it will definitely rain, or neither, but certainly not both. The exchange seems closer to disagreement, but closer still to a mere difference in expectations. The truth bearing of each attitude is irrelevant since neither can be right or wrong without the other following suit. The case could be modified such that there are stark differences and more explicit disagreement, but that would only tell us about cases in which one speaker says something to the effect of “probably *p*”, “highly likely *p*”, or “almost certainly *p*”, and another says “doubtful/unlikely/possibly *p*”. It would be inconsistent to believe that something is both highly likely and near impossible at the same time, so speakers who differ in their expectations to this extent are more clearly and substantially disagreeing. Huvenes’s aim is to make sense of the relation between disagreements and modal expressions, including terms such as “may”, “might”, “possibly”, and “probably”. He posits them as expressing different credences rather than outright propositional belief states, as well as embedded propositions which sometimes differ from the propositions that are expressed or uttered.<sup>26</sup>

Huvenes’s overall argument does not require that context-dependent significance be fundamental for credal disagreement, and while it does not necessarily aid the distinction between credal and doxastic disagreement, I regard it as helpful for understanding my point about the superficiality of minor differences in credence counting as a disagreement. I am sympathetic towards maintaining a distinction between credences as fine-grained doxastic states and belief, disbelief as full, coarse-grained doxastic states. However, I am simply not convinced that a credence difference of 0.05 and a credence difference of 0.75 should be regarded as equally robust, substantial, or significant cases of disagreement. I would regard the former, e.g., a 0.8 vs. 0.75 credence pair as closer to doxastic agreement. I now turn to cases involving suspended judgments, withdrawn beliefs, or mere absence of conviction; that is, cases in which at least one speaker has no strong belief state regarding the matter at hand whatsoever. I

---

<sup>25</sup> See Sturgeon, S. 2020.

<sup>26</sup> Ibid; 994-5.

consider cases involving these to be quite telling about the indeterminacy of certain cases with regards to identifying them as (explicit, robust) agreement or disagreement.

### 1.7: Judgment Suspension and Atheist/Agnostic Cases

I have so far focused only on cases of disagreement (and, sparingly, agreement) in which speakers on both sides have firm beliefs or propositional convictions. However, many accounts of doxastic and credal disagreement include suspension of judgment as a relevant state that is apt for disagreement, so two speakers can disagree if one holds a belief on some matter, and the other has suspended belief about that same matter. For example, Alvin Goldman says the following:

Instead of confining disagreement to cases of one person believing a proposition and another disbelieving it, let the term “disagreement” apply to any case of two people holding *contrary*, or *incompatible*, credal attitudes towards the same proposition. This includes one person believing the proposition and the other suspending judgment. And, instead of restricting the range of doxastic attitudes to the tripartite categories of belief, disbelief, and withholding, let us include graded beliefs or subjective probabilities among the set of categories... Finally, assume that not only belief but any doxastic attitude can exemplify the property of being justified or unjustified.<sup>27</sup>

It might be tempting to think that those who suspend judgment on a matter have no belief, no horse in the doxastic race, as it were, so there is simply no way to compare those who suspend judgment with those who (dis)believe. However, Seth Yalcin notes that characterising suspension as merely failing to agree with everyone reduces the suspender to a state of disagreeing with everyone who has a firm belief or disbelief, essentially trivialising disagreement and making it hard to characterise it “in any interesting sense”.<sup>28</sup>

We can make sense of agreement and disagreement between judgment suspenders if we follow Jane Friedman and Filippo Ferrari, who each posit suspending judgment as a *sui generis* state of commitment to neutrality and refraining from belief and disbelief. Ferrari contends that suspending judgment cannot be accounted for merely by reference to the nature of belief or disbelief.<sup>29</sup> Whereas Friedman contends that:

---

<sup>27</sup> Goldman 2010; 190.

<sup>28</sup> Yalcin, S. 2011; fn15.

<sup>29</sup> Ferrari, F. 2022; 529-30.

Suspension requires some sort of decision about or commitment with respect to the truth of *p*; it isn't a state that we are in in virtue of being opinionless, rather it is a state of opinion. It is in this sense that suspension is, or at least involves, a proper doxastic commitment about the truth of *p* on the part of the subject. The most natural way of understanding this commitment is as an attitude. What sort of attitude? A subject who suspends is effectively neutral or undecided about whether *p* is true. Her attitude then is one that represents or expresses or simply is her neutrality or indecision with respect to the relevant content.<sup>30</sup>

In other words, when we suspend judgment over a proposition, we have some conviction and are taking a firm stance with respect to that proposition, just not one of belief or disbelief. It is not a failure to believe, it is an active commitment to non-belief, and still represents a defensible stance. Even the withdrawal of belief may be thought of as a defensible stance, in the sense that one may have an established stance or position on some matter but then face sufficient reasons against that position that warrants retreating from it and adopting a position of agnosticism. Alternatively, one may simply be in a position of uninformed uncertainty, not being sufficiently aware of the matter or the relevant considerations to back any firm stance. The prominent case that comes to mind is that of theists, atheists, and religious agnostics. A theist and an atheist may disagree with each other in a sense that is clearer or more explicit than the sense in which either disagrees with an agnostic, but the agnostic need not be characterised as having *no belief* on the relevant matter. It makes no sense to say that agnostics agree with either the theist or the atheist, nor that theists and atheists agree with each other.<sup>31</sup> That does not mean we should say that only theists and atheists disagree with each other.

Importantly, theistic beliefs can admit of a range of degrees of confidence in the way I have described already; theists and atheists can be more or less convinced and thus vary in their confidence even amongst themselves, which echoes the discussion about epistemic modals. It would be strange to say that every speaker who is newly pursuing inquiries of personal faith are as committed or convinced as those who have longstanding religious beliefs such as the Pope or the Archbishop of Canterbury.<sup>32</sup> While I have said the agnostic need not be characterised as having no belief on the relevant matter, I contend they ought not all be regarded as suspending

---

<sup>30</sup> Friedman, J. 2013; 59.

<sup>31</sup> On at least some questions. There is prominent disagreement on the question of whether any gods or beings that are (e.g.) divine, supernatural, omniscient, omnipotent, responsible for all created things, etc. exists. But there may be other non-theistic commitments that theists, atheists, and agnostics may share.

<sup>32</sup> However, one's personal faith is largely a self-determined matter, so some theists may regard themselves as equally persuaded and strong in conviction as any prominent religious figurehead, and likewise for prominent public atheists such as Richard Dawkins or Christopher Hitchens.

judgment in the way I have discussed in this section. Even if one does not regard (theistic) belief as a quantifiable state, agnosticism about theistic matters surely admits of a range of judgments just as much as theistic and atheistic propositions. An agnostic need not be actively committed to non-belief and non-disbelief, they may be uncertain and at the same time actively pursuing an independent inquiry into the matter. Suspended judgment, then, is a legitimate doxastic state and one that can be defended or justified, and there is a real sense in which one can agree or disagree with speakers who have suspended judgment on a matter. That is, a pair of speakers who have each suspended judgment on a matter, either because they share the commitment to refrain from believing either way or, for example, because they each believe that the available evidence and arguments are inconclusive, can be regarded as agreeing. This is not officially endorsed by the likes of Ferrari or Friedman, but I believe it can be inferred from their claims regarding the nature of judgment suspension and the live possibility of disagreeing with a judgment suspender.

At this point, one might wonder if I regard these cases of agreement or disagreement as superficial in the same way I have described the other cases I have discussed. Recall that I regard the vegan and vegetarian who are persuaded by the (purported) moral justifications for veganism to different degrees, and cases involving speakers with minimally disparate credences to both indicate disagreement in a superficial sense, if not something more akin to agreement. I will not offer a detailed account of what I consider superficiality and robustness to consist of here. However, at a first pass, what I regard as a superficial disagreement refers only to the mere difference in states (or the fact that there is such a mere difference) whereas a robust disagreement involves the more realised, detailed activity or process. Robust disagreements involve reasoning, justification, evidence (both shareable and non-shareable), whereas mere differences in states such as belief, persuasion, confidence, and so on, are minimal, superficial kinds of disagreement. On this understanding, the mere fact that one speaker believes that  $p$ , another believes that not- $p$ , and/or a third speaker has suspended all their  $p$ -related beliefs, is akin to the mere minimal differences in states that I have already discussed. Such cases, like the Vegan/Vegetarian 3 case and minor differences in credence case, indicate a superficial disagreement.

### 1.8: Recap:

I will say more on this later on, but for now I take myself to have sufficiently laid the foundations for considering several ways in which we can agree or disagree with each other. I have considered various accounts of disagreement that purport to be “simple”, but which actually signal that mere doxastic disagreement does not capture all the respects in which we

agree or disagree. I have elaborated on distinctions between the semantic and pragmatic contents of a state or speech act, and the force of a speech act, as well as surveying certain possible kinds of incompatibility that may hold between speakers' states or their contents (both general and specific or technical). Finally, I have begun to consider accounts and examples of disagreement that do not simply involve full doxastic states of belief vs. disbelief, or assent vs. denial. Overall, I regard myself to have reached the following conclusions:

- 1) Purportedly "simple" accounts of disagreement are either too simple or not sufficient to capture a great many instances of examples we would consider intuitive and familiar.
- 2) The focus on doxastic disagreement proves to be too narrow and simplistic because it overlooks differences in pragmatic content, which might prove to be a barrier to disagreement or indicate a mere terminological difference.
- 3) There are plausible cases of disagreement in taste, preference, appraisal, and degrees of confidence or persuasion regarding certain propositions, all of which are not necessarily accounted for by mere doxastic accounts of disagreement.
- 4) Moreover, the simpler doxastic accounts admit of various respects in which states can differ or be incompatible, while not always delineating the differences between each respect.
- 5) Any given instance of agreement or disagreement has several aspects that vary or fluctuate; there appears to be no common, uniform, or prototypical instance of agreement or disagreement and the most general claim we can make is that they are complex and multi-faceted:

I elaborate further on these points, addressing various non-doxastic accounts of disagreement within the literature, as well as the question of whether having different types of speakers or interlocutor bears on whether and how we may classify an interaction. The implication of all this, I argue, is that trying to establish a fixed account of necessary or sufficient conditions that accommodates all possible cases of agreement and disagreement simply will not work. It is a fool's errand and a wild goose chase, and after considering some alternative concept models, I propose that agreement and disagreement are better understood as ambiguous.

## Chapter 2: Agreement, Disagreement, and Non-Doxastic States

In the previous chapter, I considered various accounts of disagreement that are purportedly ‘simple’ because they emphasize specific conditions for disagreement that are at least sufficient, if not purportedly necessary. I argued that such narrow accounts cannot generalise or explain all agreement and disagreement cases, and that while there may be several ways in which we can agree or disagree, no single factor, state, or relation can truly be regarded as common to them all, except the fact that we call them ‘agreement’ or ‘disagreement.’ My account is simple in that it describes agreement and disagreement in very straightforward, concise terms (i.e., they are complex and multi-faceted phenomena), but I stand apart from the previously considered simple accounts because I avoid emphasizing any given condition as necessary or non-negotiable. Indeed, my aim is to refrain from the pursuit of necessary or sufficient conditions for agreement and disagreement. If it were just the states or attitudes that vary, then we might posit both phenomena as a relation between held or expressed states. However, there are multiple states that are apt for agreement and disagreement, and multiple respects in which these states can be alike or different from each other, both generally and in specific ways, so the states involved cannot take sole focus. Each account might capture *some* instances of agreement and disagreement, but they fail to account for all cases because they are too narrow and specified; in other words, agreement and disagreement *must* be multi-faceted. Nonetheless, some of these accounts deserve consideration, which is my focus in what follows. As in the previous chapter, when referring to my more complex, borderline pluralist reading of agreement and disagreement, I will use the terms ‘concord’ and ‘tension’ as catch-all terms for a variety of relations that might occur in agreement and disagreement cases. If I believe a particular technical notion is relevant then I will use it, but otherwise we may, I argue, assume that there is either only one unspecified conception of likeness or consistency for agreement, and one unspecified conception of difference or inconsistency for disagreement, or multiple conceptions that coincide.

My focus in this chapter is just on the different kinds of states with which there can be agreement and disagreement. Most of these are non-doxastic, but I want to consider at least one account that is at least partially doxastic. I begin by reconsidering C.L. Stevenson’s account of moral disagreement as disagreements in non-doxastic attitudes of appraisal. Stevenson’s account places emphasis on two aspects, first of which is his moral emotivism: Stevenson regards moral appraisals as being naturally persuasive rather than doxastic. The second is that disagreement in attitude, on Stevenson’s account, involves either a desire-like state or the motive to ensure that one’s opponent change their attitude. I consider the possible agreement or disagreement that arises out of the prisoners’ dilemma and contend that it blurs the lines between agreement and



disagreement in isolated forms. I then discuss other accounts which involve desire-like states such as preferences, intentions, recommendations, and plans, before considering two final accounts which view agreement and disagreement not in terms of relations between states, but as normative claims or statuses. I regard each of these accounts as delivering legitimate cases or kinds of agreement or disagreement, but I also contend again that they are all too narrow to account for the multifaceted, jumbled nature of agreement or disagreement in general, and so they are insufficient as a complete, unified account. In Chapter 3 I will consider the relations that hold between these various types of state or attitude. I will discuss contradiction, coherence, and the like, as well as considering various types of speaker that I regard as hindering the prospects for identifying or engaging in substantial and significant agreement or disagreement, and as bolstering my claim that agreement and disagreement are better regarded as ambiguous concepts.

### 2.1: Appraisal and Persuasion

When discussing them in the previous chapter, I noted that many moral and aesthetic disagreements come about as the result of speakers having different appraisals for some object, act, or agent. I spoke about C.L. Stevenson's distinction between factual disagreements in belief and disagreements in (non-belief-based) attitudes, wherein speakers have differing attitudes that cannot both be satisfied. The kinds of attitudes that Stevenson has in mind appear to be those that admit of satisfaction, rather than truth or falsity, which is why they differ from beliefs, since beliefs can be true or justified but they cannot be satisfied. Some beliefs can be held together but not true simultaneously, suggesting an internal incoherence, but this need not be the same for non-doxastic states and attitudes. One can be both hungry and sleepy at the same time, and one can only satisfy one of these desires or needs at a time, but it need not be considered incoherent that one is both hungry and sleepy simultaneously. Regardless, Stevenson's conception of disagreement in attitude has developed through his writings and is largely the foundation from which several alternative accounts of disagreement take their lead. Let us suppose that disagreement in attitude occurs when speakers have different evaluative states regarding something.<sup>33</sup> Stevenson characterises this in two different ways:

Two men will be said to disagree in attitude when they have opposed attitudes to the same object – one approving of it, for instance, and the other disapproving of it – and

---

<sup>33</sup> We might even consider belief and disbelief evaluative states, if we take them to indicate an evaluation of something as sufficiently true or false, justified or not, etc.

when at least one of them has a motive for altering or calling into question the attitude of the other.<sup>34</sup>

This occurs when Mr. A has a favorable attitude to something, when Mr. B has an unfavorable or less favorable attitude to it, and when neither is content to let the other's attitude remain unchanged.<sup>35</sup>

The earlier account emphasizes both a difference between each speaker's appraisal regarding some object as well as each speaker possessing some motive or desire to change each other's attitude. The latter, more recent formulation appears to have weaker commitments on the part of each speaker; rather than each speaker being motivated to change the other's mind, they need only have a desire that their interlocutor not evaluate the relevant object in the way that they do. That is, speakers disagree when they have different object-relevant attitudes that somehow irk each other. This echoes the attempt I made to distinguish robust and superficial disagreements at the end of the previous chapter, with the latter account of disagreement in attitude describing a rather superficial notion of disagreement (*mere* difference of attitude, perhaps some minor irritation).

However, my focus here is not to draw comparisons or simply retread the claim that disagreement in attitude is set apart from any of the kinds of disagreement that the simple views aim for, or any other doxastic-exclusive accounts. More can be said about the changes between Stephenson's earlier account and later account, particularly when we consider the role of persuasion he emphasizes in moral language. He contends that moral discourse is used primarily "not to indicate facts, but to *create an influence*... They *recommend* an interest in an object rather than state that the interest already exists."<sup>36</sup> Moral language, for Stevenson, is persuasive as well as descriptive. However, persuasion need not only arise in disagreement, but it could also occur implicitly in agreement cases and benign conversations by using terms such as "well", "perhaps", or "maybe", or explicitly with the use of expressions such as "You're wrong!" or "That's simply not the case!" Stevenson's expressivism need not obscure the point about agreement and disagreement, which is that sometimes our non-doxastic states can be apt for these phenomena, and while there need not be any underlying intention to persuade or compel each other, sometimes this does happen.

---

<sup>34</sup> Stevenson, C.L. 1944; 3.

<sup>35</sup> Stevenson 1963; 1.

<sup>36</sup> Ibid; 16.

## 2.2: Motives

Stevenson changes his emphasis on motive between his earlier account of disagreement in attitude to his later account. In its earlier, stronger iteration, he specifies that there is neither agreement nor disagreement between speakers when they simply differ in their beliefs or attitudes without any motive to make them alike. In other words, if speakers lack the desire to compel or persuade each other to change their attitude or state, there is no disagreement, and because each speaker's attitude differs, neither is there any agreement. A clear counterexample to this, I contend, is the prisoners' dilemma case, which shows that speakers can have multiple desires at once that do not always align, but which also do not always indicate agreement or disagreement. In the classic setup, two suspected criminals (call them A and B) are being interviewed separately. Their crime can only be proven if at least one of them confesses or testifies, so each criminal faces a choice between testifying against their partner and pinning the blame on them alone or remaining silent and co-operating with their partner. This results in one of three possible outcomes:

1. If one testifies against the other and the other remains silent, then the one who testifies goes free and the one who is blamed receives substantial prison time.
2. If both testify, they each receive less time than if only one testifies, but a greater combined length of time imprisoned.
3. If both remain silent, they receive a minimal charge and the shortest time in prison.<sup>37</sup>

In cases such as this, both A and B have an incentive to testify regardless of what the other one does; if one of them has chosen to remain silent, the other can either testify against them and avoid jail time or remain silent also and receive the smallest possible sentence. If either one of the speakers has chosen to testify, the other's incentive is clearly also to testify. However, because both A and B have been isolated and are interviewed individually, they will have no inclination or knowledge of each other's decision until the outcome is revealed.

We can assume, I believe, that both A and B desire to avoid being imprisoned, although any philosopher of language or linguist who thinks their desires express indexical perspectives may maintain that they do not have the *exact* same desire. That is, A's desire might have the content "I want A to be free" or "I want A to have minimal jail time", and B's desire might have the content "I want B to be free" or "I want B to have minimal jail time". These cannot be reconciled or equivocated with both of them having a desire with the content "*I* want to be

---

<sup>37</sup> Some versions of a simple two-party setup present the participants as players in a game with corresponding moves and results. See, for example, Kuhn, S. 2019; §1.

free”, or “*I want to have minimal jail time*”. I digress. If they have desires with all the same non-indexical content, it seems fitting to regard them as having (roughly) the same desire, namely, to either avoid or limit any imprisonment. We can, I think, charitably assume that both suspects have just such a desire, but they may also have a second desire of a different kind – one that is either self-serving or altruistic. That is, either or both suspects might also desire either that their partner is imprisoned, or that they (also) go free. These two kinds of desires can and likely will occur concurrently. Regardless, if we can at least roughly say that both suspects share at least one of these desires, then they agree in some sense.

What remains unclear is whether they each have any desire or motive to change the other’s mind, since there are at least two different kinds of desires that each suspect may hold and may or may not share with their partner. One might read this as saying there is *no* common target attitude over which there may be agreement or disagreement, or one may read this as saying there are *multiple* possible target attitudes on which the suspects may agree or disagree. In other words, for any prisoner’s dilemma-style scenario, each relevant participant (a suspect or prisoner) may have at least one of the following desires:

1. The desire to go free *or* the desire to minimise imprisonment.
2. The desire for their partner to go free *or* the desire for their partner to go to prison.
3. The desire to save their partner *or* the desire to betray their partner.
4. The desire that their partner shares all the relevantly held attitudes above *or* the desire that their partner shares some but not all of the attitudes above.

Because the dilemma is merely a thought experiment, it can either be set up in such a way that there is a common attitude on which the suspects may agree or disagree, or in another way such that there is *no* common attitude, and therefore neither agreement nor disagreement. The prospect of agreement or disagreement will hinge in part on whether we assume that both suspects are equally loyal to their accomplice or selfish, or guilt-ridden about their activities or unrepentant. We can only determine how divergent or aligned the criminals are by playing out each possible situation where either or both criminals protect or betray each other.<sup>38</sup> Therefore, the prisoner’s dilemma neither explicitly fits with Stevenson’s account of disagreement in

---

<sup>38</sup> One might think that the latter state is more aptly described as an intention rather than a desire, but I do not think it matters all that much for the sake of the discussion here. At least in prisoner’s dilemma cases, I do not believe the differences between motives, intentions, and desires are so substantial that it warrants regarding them as distinct, separate states.

attitude, nor does it offer a clear counterexample; nonetheless, I think it can still be shown to present some issues, as we will see next.

Both suspects are interviewed separately so it is possible that neither will be informed of the other's decision until both have decided, meaning each of them can only speculate on their partner's intention, and thus their decision. Whether one thinks of motives, intentions, and desires as separate states entirely or as synonymous terms for the same kind of state, we can nonetheless say there is a sense in which each criminal's 'desire' (used here as a catch-all) to avoid imprisonment can come apart from the 'desire' to protect or betray their partner. Suppose that both suspects share the former 'desire' and (in isolation) agree that they want to avoid imprisonment. Unfortunately, neither of them can avoid imprisonment without betraying the other, and the mutually optimal outcome requires that each suspect trusts their partner and remain silent. Stevenson maintains that disagreement in attitude occurs at least when two speakers have different attitudes (with a shared target) and that either each agent desires that the other change their attitude, or that both are compelled to actively persuade the other to do so. As I have pointed out, the suspects in any prisoner's dilemma case could have at least one shared or common desire, if not as many as three at once. However, the nature of such cases is that they are not uniform and thus have no fixed outcome, both with respect to agreement or disagreement between the suspects and regarding each suspect's individual fate.

Firstly, let us imagine that both criminals decide to act differently, one choosing to stay silent and the other choosing to betray them. They might both share the desire to avoid imprisonment but they may have different degrees of regard for each other, with the latter suspect seemingly having less regard for their partner than their partner does for them. Already we can see that this is a messy case that Stevenson's account cannot necessarily make clear sense of, because it would be difficult to envisage either suspect desiring that the other change their mind about avoiding imprisonment. However, were the silent suspect to know in advance that their accomplice intends to betray them, they would surely then have a vested interest in changing the other's mind, if not a stronger interest in changing their own decision and ensuring they both go to jail for a shorter period of time. The criminal who testifies would not likely be interested in changing either their own mind or that of their partner, who by now will have been imprisoned. Is it fit to call this a case of disagreement, either from the perspective of Stevenson's account or otherwise? On the one hand, both suspects share the common desire (we may assume) to avoid imprisonment but have different attitudes regarding their respective loyalty to each other, and plausibly only one would desire to change the other's mind. However, recall the distinction I drew earlier between propositional and personal disagreement (disagreement with

an attitude vs. disagreement with a speaker/agent). It is surely possible for Stevenson's account to accommodate both kinds, allowing for disagreement with speakers who have no awareness of it (or speakers disagreeing with us, unbeknownst to us. So at least one version of the prisoner's dilemma is consistent as an example of disagreement in attitude on Stevenson's understanding. If we focus on each criminal's individual divergent attitude, held in isolation, the difference constitutes the minimal, inconsequential, and superficial sense of disagreement I have discussed. The mere difference in their 'desires' does not count for all that much, it is only once both prisoners have decided and either one or both are imprisoned that the weight of their decision counts becomes clear.<sup>39</sup>

The second outcome has both criminals anticipate that they cannot both go free and each independently opt to testify and betray the other. Again, we are left to assume that they each have the same desire to avoid imprisonment, and they subsequently share an overriding attitude ('desire') of self-interest but presumably each criminal will desire that the other changes their mind. This is because whichever criminal changes from betraying their partner to remaining silent would take the fall and allow their partner to go free. Both criminals share the same attitude, but desire that the other person changes their mind which, I argue, contrasts with Stevenson's account and delivers an outcome that is closer to agreement in attitude than to disagreement. Again, I am unsure that a mere commonality of attitude suffices for agreement in any robust sense. Consider how it compares with other merely shared states: we both believe that hot weather in the UK is a rarity, that aubergines are unappealing in taste or texture, and that renaming Twitter to "X" was unnecessary. We agree on these matters, but it causes me to ask, "So what?"

Much like the second outcome, the third turns out closer to agreement (in Stevenson's sense) than to disagreement. Even if we agree with the indexical advocate and say that A's and B's desires are not identical in this case, they are at least sufficiently similar to suggest they agree, only in this instance the agreement does not arise out of self-interest but out of loyalty. However, I contend that the third outcome differs from the second only because in case with the third outcome, it is unlikely that either A or B has any desire that the other change their mind. Choosing to act in a way that is not self-interested provides the optimal outcome for both A and B, and while each of them may only be able to speculate on their accomplice's decision, they will each plausibly desire that the other opt for loyalty and *not* change their mind. This stands in

---

<sup>39</sup> Further on in this thesis, I refer to a discussion in the epistemology of disagreement that compares the epistemic import or bearing of apparent and actual disagreement, and contend that, epistemically, actual disagreement has far more significance than merely possible or distant disagreement (and likewise for agreement). As far as *classifying* disagreement goes, the same judgment does not apply because that is the question of whether there is disagreement, not whether it matters or bears on our attitudes.

direct contrast to Stevenson's notion of disagreement in attitude at all ventures. Disagreement requires difference in attitude and some 'desire' that each speaker change their mind, but just the opposite is happening here, so not only do the speakers agree, but they agree because there is a consistency to their shared attitudes, which the first two outcomes both lack. The first outcome delivers mixed results and the second delivers a weaker notion of agreement, but in the third outcome, I argue there is commonality across a range of attitudes. The desire to avoid imprisonment, (plausibly) the recognition that they cannot both go free, the acknowledgment that loyalty delivers the next best option and the one that benefits them both the best, and the hope or desire that their partner shares all the same attitudes. I believe this is more substance to this than merely sharing one particular attitude or state; it makes little sense to say that two speakers in agree in virtue of the fact that they are both, for example, hungry or bored. Importantly, I have not discussed the possibility that agreement and disagreement may result from the satisfaction (or prevention thereof) of states such as desires, motives, or intentions, but that is simply because that relates to the question of the respect in which states can relate to each other. I address this question, and which relations are relevant and possible, later. For now, the discussion of prisoners' dilemma cases, which pervade several domains, should show that Stevenson's account of disagreement in attitude does not fit all instances of conflicting attitudes. More specifically, it shows that speakers can have multiple coincidental attitudes that do not always consistently align, but in fact sometimes separate or diverge, making agreement and disagreement either harder to identify, or simply less robust.

### 2.3: 'Desires' and Preferences

In the previous section, I used 'desire' as a catch-all term that refers to several states, including but not limited to desires, preferences, intentions, and motives. I believe that these states are sometimes, but not always, interchangeable, and that while the intricacies of how these states relate, overlap, or differ are interesting and worthy of consideration, it would diverge from my wider focus here. Nonetheless, at least some authors regard these states as different enough that they posit unique and distinct accounts of disagreement based on these various states. Simon Blackburn regards disagreement as based on preferences, Michael Ridge discusses disagreement in prescriptions or recommendations, and Allan Gibbard defends an account of disagreement in plans and planning states. I will consider these various accounts now, beginning with Blackburn.

While he acknowledges that not all instances of divergent preferences amount to a disagreement or imply it, he contends that at least some actual or implied instances of disagreement are possible:

If I permit smoking in my house, but you forbid it in yours, we do not necessarily disagree about anything... If I am minded to permit smoking in our house, and my wife is minded to forbid it, we do disagree. Only one of these practical attitudes can be implemented, and I am for one, and she is for the other... If the case is like that of separate houses we can sometimes ‘agree to differ’, and drop the conversation. But sometimes, even if we do not have to coordinate our actions, we cannot agree to differ, for serious enough differences cannot be tolerated.<sup>40</sup>

Again, I regard desires and preferences as at least occasionally separate states, for example I might desire to go out to eat but have no preference between eating pizza or sushi. Nonetheless, there are at least some occasions where these two states either coincide, are consistent with each other, or are otherwise synonymous. Blackburn is interested specifically with the satisfaction of preferences, desires, etc. He regards it as apt for disagreement if two speakers desire different things that cannot both be achieved, and while I am sympathetic to that as a possibility, I will discuss it later. Suppose the speakers merely have different desire-like states, or states with different contents; regardless of whether both states can be actualised, satisfied, achieved, met, and so on, do the speakers disagree? Once again, I contend that when these kinds of divergence do yield genuine disagreement cases, it is disagreement in the weakest or minimal sense but, importantly, it need not be the case that every time two speakers desire different things, they necessarily disagree. I have said as much with the Tea/Coffee examples in the previous chapter: if I prefer coffee and you prefer tea, it need not mean that we have any (substantial, robust) disagreement, nor that either of us is denying that the other’s preference is legitimate or real. Nor is it necessarily the case that all instances of desiring the same thing equates to agreement. I believe the sensation of hunger illustrates this well, and that it would make little sense to say that two speakers who both say “I’m hungry” agree over the matter, really it is just a shared feeling or desire.

I believe the same can be said of the smoking example, if you and I do not mind smoking inside our respective homes then we may be merely apathetic about it (not disliking it enough to forbid it, nor sufficiently approving to endorse it). The possibility of agreement and disagreement over suspended judgment, which I discussed in the previous section, is echoed here, since apathy might be akin to agnosticism or suspension of belief. The fact that I do not

---

<sup>40</sup> Blackburn, S. 1998; 69. Whether it is apt to call this a preference is, I contend, an open question. I believe permitting smoking inside one’s house need not be regarded an endorsement or encouragement of smoking, just as preferring that smoking not take place indoors can fall short of outright forbidding it.



mind people smoking in my home need not be regarded as my endorsement or encouragement of it. What about cases in which one has two active desires that conflict and cannot be reconciled such as, for example, desires involving winning a game or sports match? Opposing players in any competitive sport or game will likely all desire to win, suggesting a common desire, but if they are grouped in teams then it would not be possible to satisfy the desires of each player or team.

Two questions come to mind: the first being whether two teams of players that all desire to win the match agree because they share the same desire or disagree because their desires compete. The second question is whether it is impossible or inconsistent to desire at the same time for both teams to win. Taking the first question, it is clear that they have similar kinds of desire or state, because each team and its members desire to win, or that their team will win. Again, if we say that their desires express different indexical perspectives, then there is no common desire shared by *all* of the players, so they are of a sufficiently similar kind, but the content of those desires will differ. Players and supporters of Team A will likely have a desire with the content that “Team A will win the match”, whereas players and supporters of Team B will have desires with the content that “Team B will win the match.” As such, while the desires are roughly similar, they are not entirely identical in content, so at best this is another example of minimal or superficial agreement. However, given the context of it being a competitive game, the two desires with roughly similar (non-indexical) contents are still plainly in tension with each other, and thus there is disagreement. They cannot both be satisfied, even though they may be alike in content, but the desire for one team to win almost undeniably entails the desire for the opposing team to lose.<sup>41</sup> The second question asks whether it would be impossible or unusual at the same time to have the desire that Team A wins *and* the desire that Team B wins? I contend that it might be somewhat strange but if one has reason to support both teams there is no hindrance in the same way that it is impossible to prefer tea over coffee at the same time as preferring coffee over tea. Certainly, the opposite could be the case, wherein one has no strong desire that either team win, or where one is apathetic about the outcome. Indeed, one might say that it is within the remit of good sportsmanship to desire that both teams win, but this is by the by. I believe this reinforces the idea that agreement and disagreement depend on more than just how two states or expressions relate to each other; the members of both teams may share the desire to win, meaning they have desires that are alike in kind or in content, but they differ in extension. More specifically, each team will desire that they or, from their perspective, “we” win, but their agreement will remain limited or superficial.

---

<sup>41</sup> I say more on the notion of disagreement based on the preclusion of mutual desire satisfaction in Chapter 3.

#### 2.4: Prescriptions and Recommendations

Recall that Stevenson's conception of moral language as prescriptive formed the basis for his distinction between non-moral disagreement in belief and moral disagreement in attitude, and why he floated the caveat that moral disagreement involves at least some degree of desire for change. Moral language is at least partly prescriptive, so moral disagreement aims to influence audiences and hearers and steer them away from their current attitude. However, some speakers can be dogmatic and tunnel-visioned, denying that any alternative to their own viewpoint has any legitimacy or import, making them apathetic or dismissive of what anyone else feels or says. Intuitively, at least some agreement and disagreement cases will involve degrees of implicit or explicit persuasion; we might convey this with desires such as "Oh, I wish you hadn't said that..." or by offering advice or caution with phrases such as "Have you considered..."

For example:

##### **Bruised Leg:**

Amy: "How did you get that bruise on your leg?"

Joel: "I'm not all that sure, but I probably just knocked it badly against something."

Amy: "Well you should go and get it looked at, just to rule out a Deep Vein Thrombosis."

Joel: "I'm sure it is just an ordinary bruise; you know how clumsy I can be..."

In this instance, one agent is making an informed recommendation; on the face of it, the other agent downplays and disregards their advice as unnecessary and not especially urgent. While it is left implied rather than expressed, we may assume that the agent is acting as though they would advise themselves not to get their bruise examined. We can imagine versions of this case where this is more explicit, say if there were a third party advising not to go to hospital rather than just the bruised speaker downplaying the urgency. Likewise, if they were to accept that it is better to get their leg checked out than not and goes without any resistance, this would be closer to agreement than if they were to express hesitance initially.

Michael Ridge offers a broadly prescriptive account of disagreement in attitude, but he first grounds this account in a discussion of Stevenson's various accounts of disagreement in belief and in attitude. Ridge characterises accounts of disagreement in attitude as fitting into one of the following two categories:

*Stevenson's account:* disagreement in attitude requires not only attitudes with incompatible contents, but a desire by one party to change the other party's attitude on this score.<sup>42</sup>

*Stevensonian account:* disagreement is just the having of attitudes which cannot both be satisfied.<sup>43</sup>

Although Stevenson himself has endorsed both of these as describing disagreement in attitude at one point or another, Ridge regards Stevenson's considered, official account of disagreement in attitude to be the former of the two above views, which involves what I call the 'desire clause'. That is, speakers disagree in attitude not simply when they have or express attitudes with incompatible contents, but when at least one speaker has a desire to change or challenge the other speaker's attitude. The most natural way to read Stevenson interprets the opposition of attitudes and the desire clause as individually necessary and jointly sufficient conditions. However, Ridge also believes there are plausible examples of disagreement cases available which fit with the 'Stevensonian' view but not with 'Stevenson's view.'

One such case is a modification of an example Stevenson himself offers. The case involves two agents who are required to dine together but have different preferences about where to eat, the case is modified because both diners hate each other:

In this version of the example, neither party need have a motive to change the other party's motives. Indeed, because each of them thinks they will in the end prevail, each has a motive to ensure that the person's attitude remains the same. For if the other party's preference were to change then it would be impossible both to have the desired sort of dining experience *and* make the other party miserable! It hardly seems like adding to Stevenson's story that the two parties hate each other and wish them to suffer should make it *less* plausible that they disagree!<sup>44</sup>

This modified case fits with the Stevensonian view of disagreement in attitude, since the two diners cannot both satisfy their desire to eat where they want *and* for the other diner to have an unenjoyable time. Certainly, this is more consistent with our intuitions concerning disagreement in belief, which should ideally be parallel with disagreement in attitude so we can explain why we classify both the opposition of beliefs and the opposition of attitudes as disagreement. Speakers

---

<sup>42</sup> Ridge, M. 2013; 42.

<sup>43</sup> Ibid; 44.

<sup>44</sup> Ibid; 43.

who differ in belief should be regarded as disagreeing, even if neither speaker desires to change the other's belief or they are unaware of each other and their dissent.<sup>45</sup>

The Stevensonian view has been endorsed by at least some moral expressivists, and even by some non-expressivists as well. Simon Blackburn, as I have already discussed, believes that the agent who permits smoking in the house disagrees with their wife who would rather forbid it, because only one of their attitudes can be satisfied. Likewise, Frank Jackson (who is *not* an expressivist) endorses something along the lines of the Stevensonian approach when he describes a conflict in desires as a matter of two speakers desiring isolated states of affairs for which it is impossible to obtain together.<sup>46</sup> However, while the Stevensonian view of disagreement in attitude is consistent with disagreement in belief, we do not generalise the concept of there being contentful states with incompatible contents. There are at least some non-doxastic states for which there *can* be differences between agents regarding the same topic, matter, or subject, for which it would be improper to classify such differences as disagreement. Ridge notes that we would not regard differences in perception of a common object as disagreement, and likewise with fantasising or supposing propositions. States of the same mental type with incompatible contents suffices for disagreements in belief and desire, but not in these other states, suggesting either that beliefs and desires are unique in some way, and/or that agreement and disagreement in those states are unique.

Ridge responds by proposing an account of agreement and disagreement based around prescription and advice in circumstances in which each party is advising the other about how to act in a particular situation. Mere counterfactuals are insufficient because agents can have ulterior motives, causing them to advise acting in ways they think one ought not to do. Ridge contends that his account applies to scenarios which are idealised in two particular ways, both pertaining to the kind of advice speakers are motivated to offer. Firstly, a speaker must offer honest advice by meaning what they say, only offering advice they would act on if they were in the relevant circumstances. Secondly, the advice must be fully candid, any considerations that are relevant to the decision must be disclosed along with an all things considered judgment. With these in mind, Ridge says that:

Two people (or two stages of the same person at different times), A and B, agree in prescription about D's  $\Phi$ -ing in C just in case in circumstances of honesty, full candour,

---

<sup>45</sup>Ibid; 41-3.

<sup>46</sup> See Jackson, F. 1985; 105. See also Ridge 2013; 45.

and non-hypocrisy, A would advise D to  $\Phi$  in C, and B would advise D to  $\Psi$  in C where  $\Phi$ -ing and  $\Psi$ -ing are identical.<sup>47</sup>

Disagreement occurs just when the (provisional) items of advice that the two speakers or stages of the same speaker offer –  $\Phi$ -ing in C and  $\Psi$ -ing in C – are incompatible. Such an account allows for cases in which speakers can neither agree nor disagree, when the advice offered is not identical or incompatible, or when one speaker has no view about how to act in C and thus is not disposed to advise at all. Ridge believes that this account can explain why speakers disagree even when neither has any actual motive to change the other's attitude, since advising incompatible actions *if* one were to advise at all does not suppose any motive or reason to change the other's attitude. Stevenson's emphasis confuses an actual desire to change others' attitudes with a counterfactual desire.<sup>48</sup> Recall the Bruised Leg example from above; Stevenson's account of disagreement in attitude explains why there is disagreement between Amy and Joel because Amy and Joel have different desire-like attitudes or appraisals regarding the urgency or necessity for Joel to go to the hospital. Ridge's account also explains why they disagree because Joel would not advise (himself) going to the hospital, whereas Amy clearly does, and the two items of advice are incompatible – one cannot coherently advise going to the hospital and not going to the hospital. However, one might say that Joel disagrees only *iff* he would advise in good faith *not* to go to the hospital, that Amy has good intentions and a sense of care, advises going to the hospital in good faith and would advise anyone in that position, including herself. The thought is that Joel might not advise himself to go to the hospital because he values his wellbeing less than he would value that of other people in the same position, so despite having the same sense of care for others as Amy, he simply deigns not to extend it to himself. Still, I do not believe Ridge's account fails to explain why they disagree, since the focus in this case *is* on whether Joel should go to the hospital, and while Joel's prospective or hypothetical advice to himself is incompatible with Amy's actual advice, as well as being inconsistent with his own other-regarding character, he would agree with Amy that others should go to the hospital, but nonetheless disagrees about whether he himself should go. That Amy is better-intentioned and more caring for others than Joel is, I contend, a mere difference in character. Joel has little self-regard, but he and Amy may still prescribe the same advice for others regarding a given situation. If they were to have relevantly similar other-regarding characters, this would not (I contend) itself count as agreement, Ridge's account of agreement is based on their would-be items of advice, for which it is surely apt to count them as agreeing.

---

<sup>47</sup> Ridge 2013; 56.

<sup>48</sup> Ibid; 56-7.

Moreover, Ridge's account shows how belief and desires are unique states. One can advise something they would not desire to do because one would be offering their advice in good faith by taking oneself to be inclined to act as one advises if one were in just the same circumstances. The fact that an advice-giver is weak-willed will have no bearing on how they advise someone who is not weak-willed to act in C, because they are supposing that they would also not be weak-willed in C. Belief is unique from supposing, perceiving, and fantasizing because one can advise belief in some proposition either by offering some practical argument, or simply by asserting that proposition, whereas one cannot advise fantasy, perception or supposition in the same way. Belief aims at truth in a way that fantasy and supposition do not, so if one says that p, they represent p as being true unless they are being dishonest or insincere.<sup>49</sup> I take Ridge to be contending that the likes of fantasy, perception, and supposition can only be *reported*, they cannot be *expressed* in the way that beliefs and desires uniquely can, if we recall the distinction I drew in the previous chapter. These are plausible merits to Ridge's account, and while Ridge intends for it to only explain agreement and disagreement in circumstances that he believes are idealised – those in which speakers are candid and honest – candid and honest advice seems to be more of a norm or an aspiration than an idealised condition for agreement or disagreement. That is, it would not be implausible or unimaginable for speakers to advise candidly, honestly, and sincerely.

I raise two points in response. The first is that we can quite plausibly have agreement and disagreement without setting a minimum threshold for honesty and sincerity in any given speaker's conviction or assertion. Candour and honesty would seem to be norms or at most sufficient conditions for agreement and disagreement rather than necessary or fundamental conditions. However, in Chapter 3 I discuss various types of speakers who lack sincerity and authenticity in their conviction or attitude and contend that this very fact limits the prospect of agreement or disagreement in any robust, non-superficial sense. Secondly, it is an undeniable fact that circumstances are rarely, if ever, ideal, given that speakers can be insincere, dishonest, manipulative, and self-interested. I have argued that I regard some examples of agreement and disagreement as superficial and others as robust according to whether there is some argumentative or justificational depth to one's expressions or states. Ridge's account treats speakers being candid and honest as an ideal or hypothetical state, which is surely more common than Ridge believes, but it says little about agreement and disagreement in those equally typical cases where speakers are not candid or are dishonest. Ridge's aim is to offer an account of disagreement in normative domains for those who endorse expressivism in those domains. It

---

<sup>49</sup> Ibid; 57-8.

succeeds in this regard but may be less compelling for those who endorse cognitivism about (e.g.,) the moral, aesthetic, or practical.

## 2.5: Plans and Planning States

I have noted that the general state of ‘desires’ may be cashed out in preferences and in prescriptions, I now wish to discuss the potential characterisation of desire as plans or a planning state and the account of disagreement in plans offered by Allan Gibbard. Gibbard, much like Stevenson, revises his account over time. In his earlier *Wise Choices, Apt Feelings*, he considers any claim that something is rational to express endorsement or acceptance of norms that permit it.<sup>50</sup> In his later *Thinking How to Live*, Gibbard contends that normative judgments endorse a kind of plan, emphasising the need for practical reasoning, decision making, and contingency plans, which are all based on whatever information one has:

Once we distinguish permitting oneself an act from ruling it out, we accord contingency plans in the following structure. A contingency plan deals with *situations* one might be in. In a situation there is a set of *alternatives* that are open to one. A plan is *complete*, for a situation, if for each alternative in this set it either rules out that alternative or permits it.<sup>51</sup>

For Gibbard, some speech acts express beliefs, others express decisions, and others still express nothing at all. When we accost someone, we confront them, perhaps because we believe they are familiar, but we may change our mind when we realise that we do not actually know them. Merely saying something to the effect of “Hey Bob!” does not express either our belief or our decision. Like Ridge, Gibbard contends that agreement and disagreement depend on conditional attitudes, but the attitude in question is one’s conditional planning state: two speakers disagree just *iff* they plan to act in incompatible ways in any given circumstance or situation. If players on opposing sports teams all think that winning is the thing to (try to) do, then regardless of which team they are on, they i) agree that winning is the thing to (try to) do, and ii) they plan to win. Likewise, if Joel thinks he ought to go to the hospital to get his bruised leg examined, all things considered, he plans to go to the hospital. If he later reinforces his approval of that decision, he

---

<sup>50</sup> See Gibbard, A. 1990; 46. He later (49) clarifies that ‘rational’ is perhaps not the most appropriate term and suggests substituting in talk of what it “makes sense” to do, or what states of mind and acts are “warranted”, “well-grounded”, “apt”, “best”, or “guided by reasoning that is good”.

<sup>51</sup> Gibbard 2003; 56.

agrees with his past self, whereas he disagrees with his past self if he later disapproves of his past decision.<sup>52</sup>

Graham Bex-Priestley and Yonatan Shemmer contend that Gibbard's notion of disagreement fails to hold if we adapt the sports team example and suppose that one of the players desires that the opposing team wins. Because that player desires for the other team to win, she plans to lose, whereas everyone else plans to win, and Bex-Priestley and Shemmer regard it odd to classify this as disagreement in virtue of each planning different actions, and equally they consider it odd to believe opponents that both desire to win as disagreeing. They believe Gibbard would respond that the player in the adapted example plans for a different goal, whereas if they were in the position and mindset of one's opponent, they would plan to win and thus not disagree. They then contend that this places constraints on Gibbard's account, which I read as saying that speakers' mental states are fixed in advance; a player must desire to win in virtue of playing for their team, they cannot have any conflicting desires.<sup>53</sup>

Bex-Priestley and Shemmer also contend that divergent plans intuitively do not count as disagreements because in most cases speakers that have different plans do not hinder each other from acting, meaning there is no cause for concern or interest. Gibbard contends that we should care about what others think or intend because we are social, interdependent beings. Sometimes we outsource our thoughts, judgments, and decisions to each other because we need to do so, and occasionally relying on each other often improves our experience in the "game of life."<sup>54</sup> This is why we should treat differences in plans, even those rooted in deeply entrenched life-attitudes, as fundamental. I am not convinced that all divergences in plans constitute disagreements, at least not in the fundamental sense that Gibbard believes they do, for similar reasons to my doubts regarding Ridge's notion of disagreement in prescription. Recalling my bruised leg example, I contended that if Joel lacked care for others but would nonetheless advise anyone with a bruised leg to get it examined, he agrees in a weak or superficial sense with anyone who would issue the same advice from a sincere, caring motive.

I contend we can say similar things about Gibbard's notion of disagreement in planning states, as the following cases show:

**Friday:**

"What are you doing this Friday?"

"I'll be going for coffee in the afternoon, otherwise I have no plans."

---

<sup>52</sup> Ibid; 67-71.

<sup>53</sup> See Bex-Priestley, G. and Shemmer, Y. 2017; 192.

<sup>54</sup> Ibid; 70-1. See also Bex-Priestley and Shemmer 2017; 193.



**Friday 2:**

“I’m going shopping in town on Friday. Do you fancy joining me?”

“I would love to, but I have lots of work on, so I’ll be in a café for most of the day.”

**Friday 3:**

“Hey, are you on your way to the café?”

“I’m really sorry, things have really snowballed for me, so I’ll be an hour or two late!”

“Okay, no problem. I’ll still be here, let me know your ETA and I’ll order for you.”

The first case is innocent enough, one speaker is enquiring about the other’s plans, there is nothing to suggest the two have or will agree on their plans for Friday, nor that their plans will be divergent. The second version of the case more clearly fits with Gibbard’s account of disagreement, since one cannot shop (in the sense of purchasing items from physical shops) at the same time as working in a café. The third version, I contend, is not a case of disagreement in plan because both speakers plan to meet each other for coffee, but their plans have not been fulfilled due to factors (presumably) out of either agent’s control. Despite this, both speakers still have the intention to fulfil their plan, neither one has acted in any way that hinders the other or prevents them both from ultimately meeting for coffee and fulfilling their plans. But Gibbard’s account does not seem to be able to account for this because it is constrained to fixed circumstances, it does not seem to be able to account for changes of circumstances or of plans.

Recall the Bruised Leg example; suppose we grant Gibbard’s claim and accept that if Amy thinks a bruised leg should be examined, she plans to get her leg examined if it ever gets bruised. In the original case, Joel dismisses the urgency, so Gibbard would likely believe that this means he plans to not (or never) get his leg examined. If Bex-Priestley and Shemmer correctly predict Gibbard’s response to the adapted sports team case, and if that response were the same in the Bruised Leg case, then both Joel and Amy would plan consistently even if they were in the other’s position and mindset. Amy would plan to not get her bruised leg examined if she were adopting Joel’s mindset, and Joel would plan to get it checked over if he had Amy’s mindset. Once again, I take this to mean that Gibbard believes speakers’ underlying attitudes are locked in for life, so if Joel is uncaring for others and dismissive of personal health, he is *always* uncaring and dismissive. However, this surely precludes any possibility of personal changes of attitude; it is unrealistic to suppose that one’s life attitudes are permanent, that if Joel downplays the urgency of a bruised leg once, he will always downplay it. Because Gibbard seems to have in

mind individual cases of plan-divergence, his account rules out any possibility that one may change their plans; Joel might ultimately decide to go to hospital because (e.g.) Amy persuades him to do so, or because he suddenly has an epiphany. He might even decide to go arbitrarily on a mere whim but, again, Gibbard's account does not seem able to predict this.

However, I am not convinced that mere differences in plans, those which cause no hindrance or obstacles, indicate any substantial sense of disagreement. The second version of the Friday plans case above fits with Gibbard's definition of disagreement because it involves two plans which one cannot fulfil simultaneously. Insofar as the speakers might each desire to meet up with the other, their circumstances and plans are not mutually conducive to meeting this shared desire. However, I contend that it really indicates that each speaker is going in a different direction for that isolated situation, it need not reflect any deeply entrenched attitude regarding how to live or any universal conviction towards always acting in *this* way in these circumstances rather than *that* way. Just as certain beliefs can arise in complex clusters of topics which overlap, so too can one's plans for any given day simply be one of many, a small component in the lengthy process that is someone's lifetime. Not only do mere differences in plans indicate at most disagreement in a minor, superficial sense, to my mind Gibbard exaggerates the import or bearing that any given plan has on one's whole life. Some choices and plans have substantial implications for one's life, but not every choice or plan is equally important or influential.

## 2.6: Correctness and Normative Status

So far in this chapter, I have focused on accounts of agreement and disagreement that emphasize non-doxastic attitudes as a way of appealing to non-cognitivism within normative domains. I now change tack and focus on three accounts that do not focus on a specific kind of state or attitude and where they do not emphasize a particular kind of non-cognitive disagreement over all others. The first two accounts are from Max Kölbel, who distinguishes agreement and disagreement *in content* from agreement and disagreement *in normative status*. As a preface to his accounts, he states he is not interested in presenting either of his accounts as "*the* notion of disagreement or *genuine* disagreement", but instead his aim is to offer merely stipulative definitions of agreement and disagreement. Much like myself, he does not believe that there can be a fixed account of disagreement that is pre-theoretic (or intuitive) or that captures and accommodates every possible case or instance. He observes that agreement and disagreement come in degrees or admit of respects, rather than being all or nothing phenomena:

Agreement concerns a *respect* of agreement, just like equality and similarity concern respects of similarity or equality. For example, in grammar a noun and an adjective might agree in respect of gender, but not agree in respect of number. Two figures might be equal in shape but not in colour. Presumably disagreement comes in as many respects as agreement.<sup>55</sup>

He then goes on to clarify both that he is interested only in agreement and disagreement among entities in a highly restricted domain (i.e., insofar as they have a certain belief or are asserting something).

The accounts of agreement and disagreement that Kölbel proposes are relatively simple but require some preliminary understanding. The first account of agreement concerns the propositional contents assigned to a belief or an assertion by the semantics of a natural language and the context in which a belief is held, or an assertion is made. The corresponding notion of disagreement depends on semantic contradiction, the inability for two beliefs or assertions to be true at the same time, with truth being localised to a point of evaluation (which could be a possible world or state of the world) rather than any objective notion of truth.<sup>56</sup> His second account of agreement and disagreement relates to the normative relation that holds between a pair of beliefs or assertions, specifically regarding relations of *correctness*. Beliefs and assertions are subject to norms of ‘correctness.’ The correctness of any given belief may, but need not necessarily be, synonymous with truth, but it is related to the concepts involved in a belief, as they are used by the community of thinkers to which one belongs. An assertion is *sincere* when the speaker believes what they say, whereas an assertion is *correct* just in case it is correct for the asserter to believe the content they assert. With those concepts clarified, let me now specify the senses in which Kölbel believes speakers can agree or disagree with each other. Beliefs and assertions agree in the first sense *iff* they have a coincidence of content, that is if they have the same content.<sup>57</sup> Speakers disagree in the first sense when the contents of their beliefs or assertions are semantically contradictory, when they cannot both be true at once. Agreement and disagreement, in the second sense Kölbel offers, is determined according to the normative relation between a pair of beliefs or assertions. More specifically, two beliefs or assertions agree when the correctness of one belief or assertion determines the correctness of the other, and vice versa, and they disagree if the correctness of one determines that the other is not correct.

---

<sup>55</sup> Kölbel, M. 2014; 102.

<sup>56</sup> Ibid; 104-5.

<sup>57</sup> Ibid; 103.

Again, correctness of beliefs and assertions is ultimately grounded in the norms regarding the application of the concepts within those beliefs or assertions.

Kölbel's distinction between these accounts of agreement and disagreement shows that there can be ambiguous cases, in which neither agreement nor disagreement (in one sense or several) is clear or explicit. Such ambiguities are, I argue, potential barriers to agreement and disagreement. Recall that there can be terminological differences that preclude agreement and disagreement but present them as appearing, often either because of some misunderstanding or miscommunication about the semantic and pragmatic contents of an utterance. Using Kölbel's two notions of agreement, one could imagine a case in which speakers make utterances that are not contradictory but are not semantically identical, thus indicating neither agreement nor disagreement in content. Presumably, there could be cases in which agreement in content could come apart from disagreement in normative status; that is, cases where speakers say the same thing with the same words, but where only one makes an assertion that is correct. Kölbel foresees just such a possibility and discusses it with an example:

Suppose at 12 noon Greenwich Mean Time (GMT) in London in summer Alfred has a belief that he sincerely expresses by uttering "It's 1pm." Simultaneously, Bert is in Barcelona and also has a belief that he sincerely expresses by uttering "It's 1pm." Let's assume that the belief and assertion of each of them concerns their respective locations. Due to the difference in time-zone, then, only one of them can be right. In fact, Alfred happens to be right, while Bert is wrong (British summer time = GMT+1, while Spanish summer time = GMT+2). Do Alfred and Bert agree<sub>C</sub> in their respective beliefs?<sup>58</sup>

Kölbel canvasses two answers, with the first maintaining that they both believe and assert the same content, the content that it is 1pm at the time they have their beliefs, and so they do agree. On such an understanding, their beliefs and assertions concern different time zones so, despite their agreement, only Alfred's belief is correct.

On the second theory, they believe and assert different contents using the same words. Alfred believes 'the content that it is 1pm at that time *in London*' whereas Bert believes 'the content that it is 1pm *in Barcelona*' and that the difference in these contents explains why only Alfred's belief is correct. Kölbel then continues with the example:

---

<sup>58</sup> Ibid; 107. Kölbel uses "agree<sub>C</sub>" and "disagree<sub>C</sub>" to refer to agreement and disagreement in the content sense, and likewise "agree<sub>N</sub>" and "disagree<sub>N</sub>" to refer to the normative status sense.

Let's look again at Alfred's belief as 12 noon GMT that he expresses by saying "It's 1pm", and compare it with another belief, namely Carl's belief 1 h later, also concerning London, that he would express by saying "It was 1pm 1 hour ago". Do these beliefs agree?<sup>59</sup>

If one grants that Albert and Bert agree in content, then one must deny that Albert and Carl agree in that sense because one would regard Carl as believing what the content that it was 1pm an hour earlier. Regardless, the fact that they do not agree in content does not prevent Albert and Carl from agreeing in normative status, since both beliefs are correct, and the correctness of either belief entails the correctness of the other. However, facts that go beyond the mere content of their beliefs are required to explain this. The second theory of content, which maintains that Albert and Bert disagree would present Albert and Carl as agreeing in content and explain their agreement in normative status without requiring any further explanation. Sometimes agreement in content entails agreement in normative status, but only some explanations can explain facts about the correctness of beliefs only by reference to the contents of those beliefs.<sup>60</sup> Again, Kölbel's accounts are only intended to identify and explain a specific subset of agreement and disagreement cases, His suspicion, which I share, is that there is no way of uniformly accounting for all possible instances of agreement and disagreement, but I will nonetheless reiterate that I regard each account of agreement and disagreement as too narrowly specified. Some of the examples I have provided (e.g., Tea/Coffee, or Vegan/Vegetarian 2) are not captured by Kölbel's accounts because they do not involve any semantic consistency or synonymy, but neither do they involve any contradiction, and yet they may plausibly be construed as disagreements all the same. None of this is to say that Kölbel's claims are wrong or mistaken, just they are only able to capture certain kinds and cases of agreement and disagreement. Kölbel himself acknowledges this, clearly stating that his intention is for these accounts to only describe a small set of agreement and disagreement cases. I am merely extending the scope of 'agreement' and 'disagreement' in general so that they embrace Kölbel's accounts alongside the various others I have considered and will consider.

## 2.7: Disagreement as a Normative Judgment

The last account of disagreement I wish to consider in this chapter, unlike the previous accounts, does not emphasize a relation between particular doxastic or non-doxastic states. Instead, it eschews the very notion that disagreement is an entirely descriptive, factual, or condition-led

---

<sup>59</sup> Ibid; 107.

<sup>60</sup> Ibid; 108.

phenomenon. Graham Bex-Priestley and Yonatan Shemmer's account posits disagreement as a partly normative matter. They contend that if both their account of disagreement and any normative anti-realism holds, then there are never any robust facts of the matter about whether two given people disagree. Moreover, their account is centred on any claim that two speakers or agents disagree involves a normative judgment that at least one participant has some reason to change their view. One might regard their account as teleological, since they place emphasis not on a common attitude between speakers (they believe their account could thus satisfy cognitivists and non-cognitivists), but on common goals, aims, or projects. If two speakers have a common goal with standards for forming or retaining related attitudes, those standards are what determine whether speakers agree or disagree.<sup>61</sup>

On all former expressivist analyses of disagreement, judgments that people are in disagreement are *descriptive* judgments. They describe a certain relation between the noncognitive attitudes of the disputing parties. This descriptivist approach is at the heart of the failure of expressivist accounts of judgments of disagreement. As we have seen, no set of purely nonnormative conditions in the offing has been necessary and sufficient for disagreement. Expressivists must stop seeing judgments of disagreement as purely descriptive and recognise their normative component.<sup>62</sup>

Bex-Priestley and Shemmer use the example of painting a fence as an example; if one agent plans to paint the fence green all over and the other plans to paint it blue all over, their common project is reason for either or both to change their plans. Perhaps one has more or most reason to change plans, but the mere fact that their shared goal generates a conflict between their plans, and therefore some reason to divert, suffices to indicate that they disagree.

Consider a different example. Their account leaves open the possibility of speakers having reasons to change their views, whether knowingly or not, because of some common goal, without thereby disagreeing:

Suppose A believes London is the capital of Kenya, and B believes Paris is the capital of Spain. They have a common project of finding the truth, and at least one of them (in fact both) has *pro tanto* reason to change his or her belief (because it is false). But the reasons for changing their attitudes are not implied by the differences between the attitudes, so this does not count as a disagreement. On the other hand, if B believes

---

<sup>61</sup> See Bex-Priestley and Shemmer 2017; 197-8.

<sup>62</sup> Ibid; 198.

London is not the capital of Kenya, *this* implies that either A or B has a false belief, and so at least one of them has reason to change his or her mind – A and B then disagree about whether London is Kenya’s capital.

Intuitively, with the first setup of the case, A and B do not disagree because they are both making assertions that, granted, are both false, but are also irrelevant to each other. The latter version of the case makes the disagreement clear whether one endorses or denies Bex-Priestley and Shemmer’s account; there is both a common-project-relevant reason for A to change their belief, and a descriptive reason (i.e., their beliefs are contradictory).

One might wonder what is so different between this account and the ‘Stevensonian’ account that emphasizes some incompatible pair of attitudes. Bex-Priestley and Shemmer specify that what is pulling the strings is not the incompatibility of some non-cognitive attitude pair, nor the inability to satisfy some desire-like state or the frustration of their overlapping goals, but some notion of ‘reason.’ Let me offer an example of my own. If Tim and I both have a common goal of completing our respective workloads, and we each desire the last remaining workspace, then neither of us, plausibly, is unreasonable in desiring the workspace, but neither can we both have or use it at the same time. Given all that I have said so far, one might think that I would regard this as simply reverting to a mere competition between agents who have incommensurable desires that cannot be satisfied simultaneously, that any reason Tim or I have to revise our desires is just because of the circumstantial facts: there is only one workspace, we cannot both use it at the same time and we cannot both expect to get all our work done in the same spot and time-frame.

Bex-Priestley and Shemmer would say Tim and I disagree because there is something stronger than the fact that we cannot both get what we want. Our common goal implies there is (some unspecified) reason for us to change our desires. *Perhaps* norms of etiquette might have some role or involvement, or other factors such as my prediction that I will not work productively anywhere else, but they would regard these as (merely) extraneous considerations. I contend that their account is in fact stronger than the Stevensonian view, as they specify elsewhere:

Now, we could re-describe the case so that Andy and Bianca do not have the usual goals when it comes to painting fences. Maybe they are involved in a competition where the winner is whoever covers more of the fence in their own colour. In this case they no longer disagree. Even though they cannot both succeed, the standards of their

competition along with Andy and Bianca's divergence in plan do not imply any reason for either person to abandon their intention. We think this is the right result. What separates disagreeing intentions from non-disagreeing intentions isn't merely their incompatibility, as the weak Stevensonian views have it, but the implication of reasons to change. According to our theory, then, the same mental states may constitute disagreement in some contexts but not in others. This is a feature, not a bug.<sup>63</sup>

I take them to be anticipating the concern I raise regarding the robustness of agreement and disagreement cases according to the other accounts I have discussed. Bex-Priestley and Shemmer appear to be in accord with me that minor differences in states or attitudes do not (always) entail disagreement. We certainly agree that there are agreement and disagreement cases involving non-doxastic attitudes, and that relations between states or attitudes of the same type can facilitate agreement in some contexts, disagreements in others, and neither agreement nor disagreement in other contexts still. They maintain there are no robust facts about disagreement but still provide desiderata for an account of disagreement and conditions that purport to cover disagreement in descriptive and normative judgments if not other attitudes as well. Similarly, I am arguing that we can make very few firm (positive) claims about what disagreement is, given how pervasive and diverse cases, kinds, and instances of agreement and disagreement are.

Importantly, Bex-Priestley and Shemmer's account is neutral with the states involved; they do not specify that, for example, beliefs, desires, intentions, or appraisals, are more or most significant or exclusive for identifying agreement and disagreement cases. The focus is more on an external judgment, maintaining that disagreement is a judgment about the standing or relation *between* those states. This, I contend, identifies a further factor and degree of complexity to the phenomena of disagreement *and* agreement and I agree in part with them, but I also argue that their account is only sufficient to capture disagreement in a sense, or rather a certain subset of disagreement cases. I am interested in the analysis and understanding of agreement, about which Bex-Priestley and Shemmer's account says and implies little, but I am fundamentally interested in demonstrating the various ways agreement and disagreement are each complicated and multifaceted. It is not just that agreement and disagreement are relations that hold between some particular group or pair of attitudes (whether merely held or expressed), but it is also a matter of how, specifically, they relate. I will consider these various relations in the following chapter.

---

<sup>63</sup> Shemmer, Y. and Bex-Priestley, G. 2021; 505.



### Chapter 3: Respects, Speakers, and Ambiguity

Over the previous two chapters, I introduced my account of agreement and disagreement as multi-faceted and deeply complex, with multiple shifting factors. I focused primarily on accounts that emphasize a relation between specific states, such as beliefs, appraisals, or desire-like states. I argued that, for each of those accounts, they are adequate for capturing a limited set of agreement and disagreement cases, but not for accommodating, describing, or explaining *all* cases or instances. Belief-based accounts are too narrow to explain why we have disagreements in desires or tastes, and disagreements in plans or appraisals do not necessarily show why two speakers who have different credences or degrees of confidence in a proposition are clearly disagreeing. I now change direction, focusing instead on two other factors that shift and change between cases. First, I consider different relations that may hold between a pair of states that in turn facilitate agreement or disagreement. Later, I discuss various types of speakers with whom we may converse, but (I argue) with whom we may not necessarily agree or disagree, or with whom agreement and disagreement cases are substantially obscured. I close by reiterating my overall argument: that necessary and sufficient conditions for agreement and disagreement are insufficient and not a worthwhile pursuit and that we are instead better off regarding agreement and disagreement as ambiguous concepts.

### 3.1: The Simple View Revisited

Recall that the initial accounts of disagreement I considered were all purportedly simple or which emphasized simplistic relations between doxastic states. For example:

The Simple View of Disagreement. *To disagree with someone's belief that p is to have beliefs whose contents are jointly incompatible with p.*<sup>64</sup>

As we explained above, the simple view is better understood as a family of views that analyse disagreement in terms of some or other impossibility of a combination of beliefs. On some variants the relation of disagreement holds when one belief is the negation of the other, on another when their content is inconsistent, and on yet another when they cannot both be true.<sup>65</sup>

Between these two accounts, we have at least a few ideas of possible relations between beliefs; they might also be said to hold for non-doxastic states or expressions thereof. Taking the first of these, two speakers or agents disagree when their beliefs have contents which are jointly

---

<sup>64</sup> Macfarlane, J. 2014; 121.

<sup>65</sup> Shemmer and Bex-Priestley 2021; 495.

incompatible with each other. However, it is ambiguous. It could be read as indicating that disagreeing speakers have beliefs which are logically contradictory – at most only one of them is true. However, there is also the possibility that the two beliefs simply cannot be held together:

For example, suppose George believes that all bankers are rich, while Sally believes that Vern is a poor banker. Sally's belief is not cotenable with George's attitudes, because George could not coherently come to believe what Sally does – that Vern is a poor banker – without giving up his existing belief that all bankers are rich.<sup>66</sup>

Incompatibility in *this* sense may coincide with contradiction; George would be self-contradictory if he were to adopt Sally's belief without revising or dropping his original belief, and so too would Sally if she were to adopt George's belief.

Importantly, the example involving them certainly provides an example of inconsistent beliefs, because Sally's belief and George's belief cannot both be true at the same time. Nonetheless, we already have at least three relations that may hold between (doxastic/belief) states, some of which are familiar from earlier chapters. Let us suppose the following holds of two states, *a*, and *b*, or their contents, or expressions of them:

*Non-cotenability*: *a* and *b* cannot be held or expressed simultaneously.

*Inconsistency*: *a* precludes the truth of *b*, *a* and *b* cannot be true at the same time or together.

Contradiction applies to simultaneous belief and disbelief in a proposition, or arguably any combination of divergent doxastic attitudes (i.e., either belief *or* disbelief *and* suspended belief). For the sake of simplicity, I think of contradiction in terms of negation, that is:

*Contradiction*: *a* and *b* negate each other or are otherwise different doxastic states.

I regard two speakers as having contradictory beliefs just when their beliefs, or the contents of them, negate each other. In a very straightforward sense, speakers disagree when one believes that *p* and the other believes that not-*p*, or some equivalent belief that implies not-*p*. Likewise, suspending belief is contradictory in this sense since suspending belief is, by its very nature, to not believe and to not disbelieve.

---

<sup>66</sup> MacFarlane 2014; 121.

That covers all three of the simple relations Graham Bex-Priestley and Yonatan Shemmer discuss in the second quotation above. However, these far from exhaust all the possible relations that may hold between speakers' states or expressions to facilitate disagreement. Moreover, the above considerations offer nothing that indicates relations which facilitate agreement, so I contend the following are (tentative) corresponding possibilities:

*Co-tenability*:  $a$  and  $b$  can be held or expressed together.

*Consistency*:  $a$  does not preclude  $b$ 's truth, nor vice versa, nor are  $a$  and  $b$  unrelated.

Once again, contradiction proves to be a bit of a sticking point. After all, if contradictory states are just those that negate each other, then the opposite of this would be states that are *identical*, meaning the upshot is that we could say there is at least some sense of agreement that occurs when a pair or group of agents or speakers express states with identical (semantic or pragmatic) contents. Likewise, I have also stated that contradiction is sometimes just for some speakers to have fundamentally different doxastic states (regarding the same proposition), so speakers would agree in a corresponding sense when they have identical doxastic states. What might we call these two relations between states? Synonymity seems apt but vague, especially given that states can be synonymous in more respects than just their type and contents. I contend a more adequate term is *concurrence*, so there is a sense in which speakers agree just when their states, their contents, or the expressions of those states *concur*.

*Concurrence*:  $a$  and  $b$  are states either of an identical type or with identical contents.

There are other relations relevant to this consideration, since if contradiction and negation are related, perhaps concurrence has a parallel relationship with something like *entailment*. As we have seen, there is at least one sense of agreement (specifically, propositional beliefs or assertions) between states when the truth of one state entails the truth of another. Max Kölbel contends that agreement in this sense is dependent on the correctness of the states in question, determined by the norms governing the contents of these states. Two states agree when one being correct determines that the other is as well. Correspondingly, we may say that there is disagreement between two states when the truth of one state either precludes the truth of another or entails its falsity, or on Kölbel's account when the correctness of one state determines the incorrectness of another. However, these accounts are all largely restricted and narrowly

confined, there are other states which can be or not be co-tenable, and for which it makes little sense to speak of truth, perhaps even correctness.

### 3.2: Satisfaction and Justification

Recall the initial Tea/Coffee example from Chapter 1, in which I compared one speaker's *liking* of tea with another speaker's *preferring* coffee. Herein we have two states which are related but different in kind. Recalling the sense of agreement and disagreement in terms of co-tenability from earlier, it is entirely possible to like or enjoy two things simultaneously, but one cannot prefer two (related) things over each other simultaneously. That is, enjoyment of tea and enjoyment of coffee are co-tenable, whereas preferring tea over coffee and coffee over tea are not. These two states might differ in another way that facilitates disagreement in some sense, one that is illustrated by some of the other examples I have used such as competing for a limited resource or the prisoners' dilemma. That is the *satisfaction* of states such as desires or preferences. Two states are incompatible with each other and disagree or facilitate disagreement in a sense if they cannot be satisfied or actualised together. Likewise, sometimes, it may be that there is agreement in some sense if two states can be satisfied together or simultaneously; at the very least it may be considered a happy coincidence, so I contend that a degree of relevance is required. One cannot prefer tea over coffee at the same time as preferring coffee over tea, but whether a speaker who enjoys tea can be said to agree with another who enjoys coffee *merely* because both speakers can have their preferred beverage is an open matter. Other desire-like states such as intentions or planning states might conflict because they cannot be satisfied together.

We quite frequently regard certain beliefs and disbeliefs as supporting the truth of other beliefs or disbeliefs, in other words as justifying other claims we hold, which is an example of justification in a normative sense, related to reasons for believing, doing, or saying certain things. Think about the sorts of claims we typically regard as supporting each other. Ordinarily, we might say that my perception of a clock reading 12:30 is evidence for my belief that it is currently 12:30, or at least that the clock reads 12:30. Not all beliefs are supported by evidence which indicates that those beliefs are *true*, others provide evidence for the *justification* of a belief. Suppose an officer is investigating a murder case, has a suspect in mind and is also currently in need of a critical piece of evidence such as a missing murder weapon. We might say that, were the suspect in question to confess or be found with the missing weapon, this supports the officer's judgment and indicates that it is correct and true. Subsequently, the officer's decision to arrest and detain the subject is also warranted. Any evidence about the reliability or justification

for a belief is not first order evidence for the truth or correctness of a judgment, but *higher order evidence* about the reliability of one's faculties, reasoning, and judgment-forming, rather than about any given judgment itself. For example, if my eyesight is functioning well, it justifies my perceptual belief that the clock reads 12:30. When we are considering whether one's judgments or evaluations have sufficient, relevant, or adequate evidence, as well as more general discussions of what those involve, we are discussing higher order evidence. Consider another example, from Daniel Whiting:

... suppose that a butler is investigating a murder. The trail of blood leading under the door suggests that a clue is to be found in the ballroom (i). The detective's assistant tells them that the evidence suggests that the butler did it (ii). The resident epistemologist says that, in that case, the detective is required to believe that the butler did it, not merely permitted to do so (iii). Finally, the detective is aware that the butler is their friend, which disposes them to think that the butler is innocent (iv).<sup>67</sup>

A great deal of the literature about higher order evidence concerns whether higher order evidence acts as defeaters for our judgments by undermining its justification, and I discuss the relation between first and higher order evidence, particularly its role in the peer disagreement discourse, in the second part of this thesis.

For now, however, one might think there is a sense in which states agree or disagree (or facilitate agreement or disagreement) if they support or challenge each other's *justification*. We have seen accounts that posit agreement and disagreement as determined by entailment relations (i.e., two beliefs agree just *iff* one entails the truth or correctness of the other). But does one state agree or disagree with another if it is higher order confirming or defeating evidence for it? The detective's judgment is that the evidence points to the butler's guilt, and the epistemologist judges that the detective is required to believe their assistant's claim that the butler was the murderer. These claims are supposed to lend credence to the overall judgment that the butler was the murderer. We can suppose that they do, but does that mean that these two higher order judgments (that the evidence indicates the butler was the murderer, and that the assistant's judgment is required) *agree* with the first order judgment "the butler was the murderer"? I am not convinced that they do. I believe it is conflating two different propositions: the first order judgment comes apart from questions about whether there is sufficient and/or reliable evidence *for* that judgment, so it is a mere implication at most. For any proposition *p*, if I were to believe

---

<sup>67</sup> Whiting, D. 2020; 790.

that I have reliable evidence for  $p$  but do not believe  $p$  itself, I do not agree with any  $p$  advocates about the matter itself, although any  $p$  advocate might agree with me that we do have reliable evidence for  $p$ . Regardless, the two are different matters, so it would be a mistake to say that believing there is or is not reliable evidence for  $p$  is equivalent to agreement over  $p$  or not- $p$ .

### 3.3: Joint Accuracy

Not all states are truth-apt; we might regard expressions of taste, desire, or preference as the sort of things that can be evaluated, but it makes no sense to describe them as true or false. Instead, it makes more sense to think of the contents of such states as fitting or not fitting, and perhaps speakers' expressions of them as sincere or insincere. Along similar lines, John MacFarlane discusses the problem of 'lost disagreement', which is the concern that any relativist about truth cannot account for disagreement cases, particularly over matters of taste:

If the truth of my claim that a food is “tasty” depends on how it strikes *me*, while the truth of your claim that the same food is “not tasty” depends on how it strikes *you*, then our claims are compatible, and we do not disagree. But it seems that we *do* disagree – even if we are aware that the source of our disagreement is our differing tastes.<sup>68</sup>

Likewise, there is a similar lack of clarity over agreement cases over matters of taste. If two agents both report to each other that (e.g.) “raisins are tasty”, it is not clear whether they agree on the fact that raisins are tasty specifically for one agent or the other, or for both agents. For MacFarlane, the kind of relativism about truth he defends posits truth as ‘assessment-sensitive’, the truth of the use of an assessment-sensitive sentence depends on features of the context in which that sentence is assessed.<sup>69</sup> An implication of this view is that there are different varieties or “levels” of disagreement, and thus the problem of lost disagreement is overemphasized. It is not that relativists cannot account for disagreements of taste, but that there are different ways of making sense of such disagreements that may not revolve around the truth of any such sentence. Specifically, he proposes the concept of *accuracy* and suggests that there are two senses of disagreement that arise when the contextually assessed accuracy of one attitude ensures the inaccuracy of another:

---

<sup>68</sup> MacFarlane 2018; 8.

<sup>69</sup> Ibid; 24.

**Accuracy.** *An attitude or speech act occurring at  $C_1$  is accurate, as assessed from a context  $C_2$ , just in case its content is true as used at  $C_1$  and assessed from  $C_2$ .*<sup>70</sup>

**Preclusion of joint accuracy.** *The accuracy of my attitudes (as assessed from any context) precludes the accuracy of your attitude or speech act (as assessed from that same context).*<sup>71</sup>

**Preclusion of joint reflexive accuracy.** *The accuracy of my attitudes (as assessed from my context) precludes the accuracy of your attitude or speech act (as assessed from your context).*<sup>72</sup>

These accounts are distinct from whether states or attitudes are co-tenable or can be satisfied simultaneously. If I have simultaneous desires, respectively, to go out in the rain and to remain dry, my two attitudes are internally incoherent, I cannot hold them together and I cannot satisfy any more than one at the same time. The notion of accuracy that MacFarlane has in mind is whether the contents of those states have the right intension, from the perspective of whoever is making the relevant assessment. MacFarlane uses the following example:

Consider Jane, in this world (the one we call “actual”), and June, in another possible world. Jane believes that Mars has two moons, and June believes Mars has just one moon. Both of their beliefs are accurate, since in June’s world Mars does have just one moon. Does Jane disagree with this belief of June’s?<sup>73</sup>

The two beliefs are not co-tenable, so Jane disagrees with June because she cannot share the belief that Mars has one moon without giving up her current belief that Mars has two moons; however, Jane’s current belief concerns *her* world, and June’s current belief concerns *her* world, with both beliefs being accurate. June’s belief being accurate in June’s world has no bearing on the accuracy (in Jane’s world) of Jane’s belief, and *vice versa*. More importantly, because accuracy is a context-sensitive matter, both of their beliefs could be inaccurate in some separate possible world, concerning that possible world, while not affecting each other, resulting in no disagreement between Jane and June.

Reflexive accuracy is used to understand how accuracy applies to states which lack standard truth values. The belief that something is tasty may be accurate from one perspective, but it is incompatible with the belief that that same thing is not tasty, meaning that from the same perspective they cannot be jointly accurate – leading to disagreement in taste in *that* sense.

---

<sup>70</sup> Ibid; 127.

<sup>71</sup> Ibid; 129.

<sup>72</sup> Ibid; 130.

<sup>73</sup> Ibid; 128.



However, a relativist will maintain that two speakers who differ in taste may also differ in perspective, meaning that joint accuracy is prevented but *reflexive* accuracy is not. Whether the prevention or preclusion of reflexive accuracy captures disagreements over matters of taste, desire, preference, and so on, depends on whether one is an objectivist or a relativist (or neither) about such things. If one believes that there are objective matters or facts concerning, e.g., humour, then they will maintain that accuracy and reflexive accuracy are one and the same, whereas the relativist maintains that they are separate because perspectives and contexts of evaluation can come apart. Regardless, differences in such states will appear to be a disagreement to the objectivist, but not to the relativist, or at least not as obviously or strongly. They may simply regard them as a difference in evaluation, perception, or whatever they consider taste to be. As MacFarlane puts it, the relativist “can acknowledge that, in some respects, disagreement about taste is less robust than paradigm objective disagreements, which do preclude joint reflexive accuracy.”<sup>74</sup>

MacFarlane offers no corresponding notion of agreement, but I believe we can surmise that speakers agree in taste in these senses when the accuracy of the one attitude ensures the accuracy of the other. It cannot be a case of mere non-prevention or non-preclusion for similar reasons that we have seen with other accounts since attitudes could be coincidentally accurate from the same perspective or context of evaluation while also being utterly irrelevant. One might regard it as accurate to hold both that “blue cheese is tasty” and that “Ricky Gervais’ jokes are tasteless”, for example, but the two matters are so far removed that one cannot feasibly say that speakers who hold the former belief agree with those who hold the latter. The notion of norms of correctness and correct beliefs that Kölbel uses to characterise agreement in normative status is echoed. Kölbel contends that beliefs and assertions are correct just *iff* they are in line with the way a community of thinkers, speakers, agents, etc. uses the relevant concepts in them. MacFarlane’s notion of accuracy for an attitude is just the truth of its contents from the perspective of a context of assessment; an assessment, on MacFarlane’s account, could seemingly be either communal or individual. Nonetheless, I regard the two accounts as grasping at similar notions and capturing them in two different ways.

### 3.4: Perception and Incongruity

All of the accounts I have discussed so far have largely emphasized factual states or relations, wherein two speakers agree or disagree because certain conditions obtain or are met. While I argue against the possibility of such conditions sufficing to capture a broad range of agreement

---

<sup>74</sup> Ibid; 130.

and disagreement cases, I have said little about the difference between established cases and cases which are merely possible or *perceived*. I have argued that many of the accounts in question generate superficial or thin cases of agreement and disagreement, and that more is required to make these cases robust or substantial. There are plausible reasons to regard distant or merely perceived cases as real and possible, but this characterisation is not quite what I have in mind here; rather, the matter at hand is whether real cases of agreement or disagreement may obtain because one of the agents or parties involved merely perceives it to do so. Many authors grant disagreement as the result of a perception of error, and correspondingly we might believe that perceptions of correctness or accuracy (in the layperson's sense) are apt for agreement; however, like many of the other accounts, this only gives us *a sense* of agreement or disagreement. An agent or speaker can perceive that one of the relations I have considered holds between their state and that of another agent or speaker, but the prospect generates further questions, such as:

1. Must both or all parties concerned perceive agreement or disagreement for either (or whichever is relevant) to occur or obtain?
2. What, specifically, might a perception of agreement or disagreement consist in? Can such perceptions be shared, common, or mutual?
3. Is there any difference between merely perceived agreement and disagreement and actual, confirmed cases? What might this mean for merely possible or distant cases?
4. If one of a pair of speakers perceives their interaction as indicating agreement between them, but the other speaker perceives them as disagreeing, can we rule either possibility out?

To be clear, I contend that agreement and disagreement in this sense certainly *is* possible, but I am less convinced of its import, bearing, or significance. Degrees of geographical or personal separation are a factor. Genuinely agreeing or disagreeing with someone, as a phenomenon, typically counts for more or has a greater impact than *merely* supposing that someone, somewhere does so, or than our unconfirmed perception or suspicions about a particular agent. The fact that someone somewhere dissents or disagrees with a particular attitude one holds will likely strike one as less significant *iff* it is faraway or unfamiliar, and the same goes for any semblance of consensus or concord. This is not to say that faraway or merely possible agreement and disagreement cases are not real, but experiencing them as actual phenomena as well as their implications are (for good or not) a confrontation. They are more important in a way that is, I contend, obvious but not trivial. Moreover, the perception of agreement or disagreement is less

significant because perceptions can be mistaken or misidentified, meaning one might be wrong about whether a particular speaker agrees or disagrees with them, whether in a particular sense or in several. As should be clear by now, not all thinner or superficial cases of agreement or disagreement are clear cut, let alone the more robust cases. Agreement and disagreement can be obscured by pragmatic misunderstandings or merely verbal disputes, they can involve mixed cases of overlapping agreement in some senses and disagreements in others.

These points lead to another interesting concept which, I believe, has parallels with agreement and disagreement. Philosophical discourse concerning the notion of *incongruity* is generally focused on presenting it as an account of humour or comedy. The thought is that we find amusing or humorous that which subverts our expectations or what we regard as atypical, absurd, or astounding.<sup>75</sup> Arthur Schopenhauer offers an account of laughter as a reaction to and an expression of incongruity between a concept and real objects:

All laughter then is occasioned by a paradoxical, and hence unexpected, subsumption, it matters not whether this is expressed in words or in deeds. This in brief is the correct explanation of the ludicrous.<sup>76</sup>

A comic or joke-teller, in fulfilling this role, acts as a metaphorical chaperone and narrator, guiding their audience or hearers through the narrative of a joke. Sometimes the joke ends as one expects it to, but other times the direction changes suddenly and shockingly in such a way that only the joke-teller can anticipate. It need not always be the case that humour is to be found in the absurd or incongruous. Sometimes we laugh more when things go as we expect or when the outcome is clear and discernible far in advance. Puns and dad jokes can be as predictable as mishaps and disasters happening for sitcom characters such as Del Boy or Frasier Crane, or real-life experiences such as submissions to “You’ve Been Framed.” Sometimes, having our expectations subverted only elicits shock in an audience, raising questions about whether there are constraints on what a joke-teller or comic can say and do both on and off stage. As Michael Clark notes, we could witness the same behaviour or event and react in different ways each time:

... instances of incongruity which are humorous are not always found humorous just because of the incongruity. Why are we amused by the behaviour of our own children and not by the same behaviour in other people’s children? Presumably because our own

---

<sup>75</sup> Morreal, J. 2020; §4.

<sup>76</sup> Schopenhauer, A. 1966, vol. 1; §13.

children are more endearing to us. Many jokes are effective because of their topicality and are completely unfunny when their references cease to be topical.<sup>77</sup>

I contend that agreement and disagreement also occur in a sense according to whether one's expectations are met or subverted, and this applies to a range of states. We can perceive a speaker as agreeing or disagreeing based on something they say or imply, and sometimes these perceptions can turn out to be mistaken. Likewise, we can expect a speaker to have some particular attitude and later find that this really is not the case, either to our delight or our horror.

### 3.5: Speaker Types and Devil's Advocates

The import or significance of agreement and disagreement varies with its locality and realism, and thus far I have argued that all the accounts of disagreement surveyed and the possible relations between our various states and attitudes all facilitate legitimate kinds of disagreement. However, I have been careful to emphasize that, while I maintain that many or most of these disagreement cases are real and genuine, they are not all of them substantial or robust, and at least a few of these cases will be apparent or illusory. Over the next few sections, I discuss cases that involve particular kinds of speakers or agents, which preclude agreement or disagreement in any robust sense. Much of the focus in the disagreement literature is either on non-specific agents or the likes of experts and epistemic peers, and largely presupposes but does not specify what disagreement is understood to involve. I have in mind here the following types of speakers or agents:

Devil's Advocates

Contrarians and Conformists

Strangers

These do not exhaust every type of speaker with whom we can engage and ostensibly agree or disagree, but something to bear in mind is that we often place trust in other speakers, which may involve an expectation that they will display integrity and be honest or transparent about what they believe or feel. Speakers of these types are not naturally disposed towards sincerity in their speech acts or to not sufficiently demonstrate any sincerity. As such, any interaction with these speakers, I contend, cannot be ruled in as agreement or disagreement in any strong or substantial sense, they may not even present anything more than the illusion of cases. Again, while I do not

---

<sup>77</sup> Clark, M. 1970; 21.

believe we can make any unifying claim about what agreement or disagreement necessarily are or involve, that need not mean that we can make firm claims about what they are not.

Let me begin with cases involving speakers who are acting as devil's advocates. By this, I mean not the devil's advocate in its historical, formal sense.<sup>78</sup> Rather, I mean the contemporary understanding of speakers who adopt a hypothetical attitude and express it to offer an alternative perspective, largely "for the sake of argument", and to contribute to a discourse without presenting any attitude or opinion they genuinely hold or endorse. Some advocates will disclose their insincerity, others will leave it to be determined or not, but regardless they are acting in a kind of persona, presenting a viewpoint or perspective that may not otherwise be offered. A speaker may also raise an alternative opinion or attitude as if there were a hypothetical devil's advocate, again often for the sake of dialectical progress. Sometimes one may do this with a view or perspective they genuinely endorse, but I contend that if one who is acting as a devil's advocate does not openly disclose that fact, then agreement and disagreement cases are only possible in the same sense that distant, faraway cases are. That is, they are prospective or possible cases, but they are not substantial or significant. However, if one's advocacy is left obscured and an audience or hearer has no inclination about what a speaker believes or feels, it leaves them in a position of uncertainty and vulnerability. If we contrast this with a case that involves suspended belief or judgment, we know that the speaker is sincere in their attitude, but that their attitude is one of agnosticism and they simply have no firm conviction either way.

To be clear, devil's advocacy is no bad thing, I am not suggesting that anyone who plays devil's advocate lacks virtue or good intentions. Having someone act in the role of the advocate can help to develop or advance an argument, because they will ensure that alternative perspectives are given a fair hearing and discourage groupthink mentalities and dogmatism.<sup>79</sup> However, devil's advocates are still human agents and thus subject to a range of flaws or fallibilities, as Katharina Stevens and Daniel Cohen note:

Even when we enter arguments without any personal investment, perhaps recruited as devil's advocates, we end up identifying with our assigned standpoints... The mere fact that our contributions are *ours* means that we risk losing face by losing: our reputation as reasonable, intelligent, or knowledgeable is on the line. The will to win takes over, eclipsing the motivation to help.<sup>80</sup>

---

<sup>78</sup> "Devil's advocate or *advocatus diaboli* – person appointed by the Roman Catholic Congregation of Rites to contest the claims of those put forward for beatification or canonization (i.e., being recognized officially as saints)." Bowker, J. 2003.

<sup>79</sup> See Bailin, S. and Battersby, M. 2016; 10.

<sup>80</sup> Stevens, K. and Cohen, D. H. 2021; 906.

Insofar as aiding argument and rhetoric, there are benefits to having a speaker's acting as a devil's advocate, and therefore to their disclosing it. The same goes for ascertaining agreement and disagreement. Without disclosure, we cannot rule in agreement or disagreement, we simply have a case that is ambiguous or indeterminate, whereas when an advocate is open about their position, intention, or motive, whether we agree with them or disagree can be established as real, albeit thin. In this respect, I endorse something like a Gricean maxim which compels agents to co-operate in conversation and disclose their attitudes fully and concisely.<sup>81</sup> The devil's advocate need not be asserting something they believe to be false, they simply may not be uttering or asserting what they truly believe or feel, so this does not forbid or preclude anyone from acting as a devil's advocate, just from not being open about it.

### 3.6: Contrarians and Conformists

Devil's advocates express states or attitudes they do not sincerely endorse to either advance a discussion or offer an alternative perspective or consideration, which may be considered productive or compelling reasons, whereas other speakers will do the same for altogether unproductive and unconvincing reasons: the *contrarian* disagrees purely for the sake of disagreement alone, whereas a *conformist* displays sycophantic behaviour and endorses an opinion or attitude merely to gain favour or reputation. These speakers generate agreement and disagreement that is superficial in the literal sense of the term, it is purely surface-level and illusory, they neither affirm an opinion or attitude sincerely nor do they act from intentions or motives that can generally be regarded as intellectually virtuous. Recall that in the previous section I alluded to endorsing a Gricean conversational maxim; Grice views conversation and dialogue as fundamentally co-operative and informative and implores speakers to embrace the following principles or habits:<sup>82</sup>

- Be as informative as is necessary but not overly informative.
- Avoid saying things one believes to be false.
- Avoid saying things one lacks adequate evidence for.
- Avoid obscurity or ambiguity.

---

<sup>81</sup> Compare H.P. Grice's Co-operative Principle: "Make your contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged" with his first maxim of quality: "do not say what you believe to be false." Grice, H.P. 1989; 26-7.

<sup>82</sup> Ibid; 26-7.

We cannot and do not need to be ideally rational or communicative, but Grice certainly identifies some key qualities for any semblance of productive, clear discourse. Agreement and disagreement are core components of many discourses, and the vices Grice warns against are clear hindrances to them, with at least some of them being exhibited by contrarians and conformists. Those who conform to another's judgment or disagree for the sake of disagreement alone may utter things they actually believe to be false or for which they lack sufficient evidence. They obscure or misrepresent their true convictions and attitudes, they may not even make particularly informative claims since they can express agreement or disagreement with minimal content.

Quassim Cassam's notion of the *epistemically insouciant* agent is relevant here; for Cassam, an agent is epistemically insouciant when they exhibit indifference towards facts or "a casual lack of concern about whether one's beliefs have any basis in reality or are adequately supported by the best available evidence."<sup>83</sup> For example, it is typical for politicians to be disinclined to evaluate evidence and acknowledge the reality of many situations on which they comment. Those who are epistemically insouciant are not concerned with either the facts or the evidence, and they may eschew any desire to engage in debate or dialogue, perhaps treating alternative viewpoints and perspectives as mere inconvenience. Importantly, Cassam clarifies that this is not the same as being intellectually dishonest, the epistemically insouciant agent does not deliberately mislead or lie to others by making claims they know are false. Lying results from the intention to conceal the truth, whereas Cassam borrows the notion of *bullshit* from Harry Frankfurt, and posits this as the output of the indifference towards truth he is interested in. Frankfurt posits the bullshitter as someone whose "eye is not on the facts at all... except insofar as they may be pertinent to his interest in getting away with what he says."<sup>84</sup> I do not want to say that the contrarian or the conformist are *necessarily* epistemically insouciant, but at least some will exhibit such obnoxious and dogmatic tendencies, so the comparison is certainly useful. Regardless, I am thoroughly unconvinced that we can say we agree or disagree with such speakers in just the same way as we do with someone whose assertions or utterances genuinely reflect their convictions, attitudes, or beliefs.<sup>85</sup>

### 3.7: Strangers

---

<sup>83</sup> See Cassam, Q. 2018; 2.

<sup>84</sup> Frankfurt, H. 2005; 56.

<sup>85</sup> Relatedly, the idea of agreeing or disagreeing with someone who is *acting* or putting on a persona may seem limited, or at least in terms of its robustness. Although I drew comparisons earlier with comics and audiences as occasionally reflecting agreement or disagreement, I regard this as at least plausibly more robust than agreeing with a conformist or disagreeing with a contrarian.

Finally, I contend that there are very limited prospects for agreeing or disagreeing in any substantial sense with speakers about whom we know very little or with whom we have limited familiarity. I do not mean agents or speakers who are publicly or professionally well known, or whose lives, attitudes, or perspectives are well-documented and commonly known, such as celebrities, politicians, or academics. Rather, I mean a genuine stranger, someone about whom we know one thing and one thing only: what they have said about a given matter, event, or object. Consider the following case:

**Busker:**

Imagine that a total stranger and I are both observing a busker performing in a London Underground station. Neither of us know anything about the other, but we both assume that we each have fully functioning, uninhibited faculties, and that we each observe the busker in just the same way. Now suppose that I have the attitude that “the busker is playing badly” and I assert just this, and the stranger responds by saying “Yes, the busker *is* playing badly.”

On some of the narrower or simpler senses of agreement and disagreement I have considered, we are saying just the same thing and presumably have in mind just the same relevant notions of “busker”, “playing”, and “bad.” We have expressed no attitudes with contents that differ, neither of us have said anything that reflect attitudes which cannot be held together, and whether it is a matter of fact or our perception, if one of us is correct in our assertion then so is the other. By all accounts, the stranger and I are in agreement.

However, recall that I explicitly said neither of us knows anything about the other, I have no awareness of the stranger’s background beliefs or attitudes, I do not know of any specific pragmatic or colloquial uses they may have for the contents of their assertion, nor they for me. Presumably, the only information that we have is that we are witnessing the same event and we each perform the same or comparable speech acts, expressing the same attitude. I have said that I do not know of any unique pragmatic contents the stranger may have, and it is plausible with many such cases that there are none. However, we can envisage at least some cases where there are different understandings of the uttered phrase. For example, the stranger may perceive the busker’s musicianship as bad because he perceives it as out of tune, whereas I may regard it as bad because he is (from my perspective) out of time. Without this clarificatory detail, we cannot ascertain the extent to which we agree; even if we were (unknowingly) using our language in the



same way, we would not be in a position ascertain that we do in fact agree, based on this pair of utterances.

But this raises the question of why I do not regard interactions with strangers as robust agreement or disagreement but do include cases involving familiar speakers who also may have ambiguous understandings or intended meanings. The difference is that with familiar speakers we are in at least a somewhat better position to recognise or know their understanding of terms, we may be able to pick up on cues to recognise what they mean or why. The same goes for strangers whose views or attitudes are publicly known. For example, although I may not know Timothy Williamson well enough to tell how sincere he is when he makes an offhand comment about the merits of knowledge-first epistemology, there are sources including his publications, interviews, or the testimony of others who know him personally, which I can rely on to trust in his sincerity and confidently rule in or out that I agree with him. With a total stranger, it is a total guess. Perhaps we agree, or perhaps we disagree, the uncertainty and obscurity are too great to say for certain, and the answer lies in pressing for details. Regardless, the prospect of (not) agreeing or disagreeing with strangers generates a notable line of scepticism that runs against the notion of propositional and faraway agreement and disagreement. On the one hand we can plausibly grant that there will be unknown figures who share our attitudes, evidence, or reasoning, and thus agree with us (or who disagree in virtue of *not* sharing). We certainly talk of such cases and grant that they involve (unknown or unacknowledged) strangers. On the other hand, when we consider acknowledged or known cases that involve strangers in person, the prospect of (robust) agreement and disagreement appears muddy.

### 3.8: Agreement and Disagreement as Real, Relevant, and Robust

Over the past three chapters, I have tried to defend the claim that there are many shifting factors and elements of agreement and disagreement that make it incredibly difficult to precisely define either phenomenon. They are imprecise, multi-faceted, and sometimes indeterminate or obscured, and no single account that has been offered suffices to accommodate or explain every possible case. It is not enough to say that agreement and disagreement depend on, or are determined by, what attitudes are involved, or by how they relate to each other, or who (if anyone) we are agreeing or disagreeing with, because different cases of agreement and disagreement depend on different combinations of these factors. I have offered my own response to various “simple views” of disagreement which contends that disagreement occurs when there is at least one kind of tension with respect to one or multiple states, and likewise that agreement occurs when there is one kind of concord with respect to one or multiple states.

However, this only vaguely gestures at what may be involved, and may not even accurately explain certain cases – we might disagree without there being any overall tension and I have explained that there are several senses in which attitudes may be in accord or concord with each other.

Moreover, I have argued that many of the various accounts I have considered have left much to be desired. It is not just that they can be challenged with plausible counterexamples, cases we intuitively regard as agreement or disagreement, but which are not accurately described by any given account. Rather, I have argued that many of the accounts I have considered describe cases which are superficial and lack substance. This is by no means a sceptical dismissal of such cases as not real or legitimate, or the contention that such cases are merely apparent. Some authors have drawn distinctions between disagreements that are superficial and disagreements that are not, such as David Enoch:

When two physicists disagree – sometimes rather strongly – about the truth (or acceptability) of a scientific theory, they typically agree at least about what would settle their disagreement... So although there is often disagreement in physics it is somewhat *superficial*, and underneath it lies a deeper agreement, if not about the truth of theories, at least about what findings support what theories and (roughly) to what extent. But it is exactly this feature, the thought goes, that is missing in cases of moral disagreement. Moral disagreement runs much deeper than disagreement in the sciences, because typically, or at least often, in cases of moral disagreement there is no deeper agreement underlying it, no agreement about how to settle the more *superficial* disagreement.<sup>86</sup>

Enoch's remarks about the distinction between disagreements in science and moral disagreements notes that scientists tend to have more higher order agreement, more consensus on what topics, theories, and evidence are relevant and sufficient, whereas this is entirely absent in moral discourse. In the first instance, he seems to have in mind a notion of superficial disagreement as referring to any surface level disagreement over some first order judgment. However, in the second instance I take Enoch to imply that such surface level moral disagreement is philosophically mundane or uninteresting, and second order matters, in ethics at least, are both more rigorous and interesting, as well as exhibiting more noteworthy agreement and disagreement. Enoch notes that there is pervasive disagreement not just about first order moral matters but about second order (i.e., metaethical) matters, and indeed about philosophical

---

<sup>86</sup> Enoch, D. 2009; 34, my emphasis.

matters in general, all of which seem to be largely insoluble but at the same time involve participants who have no cognitive shortcomings.<sup>87</sup>

Sarah Stroud, on the other hand, does not explicitly identify any kind of disagreement as superficial, but she notes that any genuine instance of disagreement requires at least some underlying agreement and a common object:

Every genuine disagreement requires a *common object* of disagreement. In a propositional disagreement, the thing A affirms must be the very thing B denies... If the thing A is affirming is *not* the thing B is denying, or if the course of action A supports is *not* the one B rejects, we lose our disagreement. A and B's respective positions become *orthogonal* rather than *opposed*, since they are not in fact taking positions on the same issue.<sup>88</sup>

This is, in part, what I have alluded to in discussing the semantic and pragmatic contents of speech acts. It is not just enough to say the same thing with the same words to agree, nor to mean the same thing, perhaps. Rather, it may also be that speakers must be aware that each is referring to the same matter, event, object, and so on, to agree or disagree. Stroud makes the point quite clear when she says, "Ambiguity and equivocation, in short, threaten to destroy genuine disagreement."<sup>89</sup> Stroud contends that any moral dispute between, for example, a Kantian and a hedonic utilitarian is merely apparent unless the two have a common perspective or moral framework. It is a natural consequence of hedonic utilitarianism that if any act generates more pleasure than it does suffering, then such an act is morally acceptable. An advocate for Kant's moral system will object both to the underlying moral framework the utilitarian endorses *and* the judgments it delivers which, in this case, is that slavery would be permissible *iff* it was to generate more pleasure than suffering.

The thought is that the disagreement becomes real only *iff* there is a common framework, perhaps because it is easier to pinpoint some sort of error when two utilitarians or two Kantians draw divergent judgments, whereas this is only a merely apparent disagreement because the Kantian and the utilitarian are using different concepts. Stroud goes on to deny the judgment that the Kantian and the utilitarian are not *really* disagreeing, arguing instead that they are engaging in a *conceptual disagreement* – a disagreement over the evaluative matter of whether it is fine, ok, good, to think in terms of some given concept. Other examples include speakers debating (for example) whether something is *spicy* or whether someone is *tall*. It is a more general

---

<sup>87</sup> Ibid; 47.

<sup>88</sup> Stroud, S. 2019; 18.

<sup>89</sup> Ibid; 18.

practical dispute over how and when to use the relevant term or concept, what the term means, or what class of objects it refers to within particular contexts or conversations. What some might regard as a difference of perspective or of pragmatic meaning or understanding, Stroud maintains is a genuine sense of disagreement. This sense of disagreement is rooted in the fact that of a pair of speakers, one has the relevant concept in their toolkit but the other does not, or the fact that both speakers use the same concept in different ways altogether.<sup>90</sup>

Between Enoch and Stroud's accounts, and the various specific accounts of agreement and disagreement I have considered, I believe there are at least three ways we may distinguish agreement and disagreement cases as either robust or superficial. More specifically, I think "superficial" might be regarded as meaning at least the following things:

1. A disagreement is surface level or merely apparent, either obscuring some *other* kind of dispute or a wider agreement.
2. A disagreement is not chiefly or demonstrably about the relevant matter it claims to be about. Or it is not possible to accurately locate or identify what the disagreement is *really* about.
3. A disagreement is not deep or important, perhaps concerning some mundane matter.

We see the first and third of these in Enoch's discussion, and the second in Stroud's, but importantly, none of these three reflect my earlier, first pass attempt, offered in chapter 1. There I regarded the difference being that superficial disagreement refers to a mere difference in states or attitudes, including beliefs, degrees of confidence, appraisals, or perceptions. Accordingly, robustness of agreement and disagreement refers, I believe, to a kind of consistency in cases. It is not just that two speakers believe something with the same contents, but they agree in a robust sense *iff* they have sufficiently similar reasoning and evidence for their common belief, and perhaps have engaged in a dialogue over it. This is sufficient, I believe, to explain why certain faraway or hypothetical differences of attitude or opinion register as disagreements, but not in any particularly noteworthy way. Perhaps it also accommodates the notion of superficiality that Enoch has in mind – that superficial disagreements are surface level, not as noteworthy or high stakes, and have more detailed, complex underlying agreement or disagreement. I believe we can grant this. For Stroud, when a particular concept fits into *some* speakers' conceptual arsenal but not others, they disagree, not over which objects fall into the class of extensions for that concept but *because* only some of the parties use or recognise the concept in question:

---

<sup>90</sup> Ibid; 21.

In that sense, to use a concept *F* is to begin to create an “*F* demarcation line,” as we might term it. And when it comes to concepts like LEWD, BLASPHEMOUS, and KRAUT, this might be precisely what you object to. This is very different from disputing what side of that line a given object falls on. Your disagreement, if there is one, is not primarily about whether your interlocutor has mis-classified his particular subject, *S*, with respect to its *F*-ness.<sup>91</sup>

In other words, conceptual disagreements occur when propositional disagreements do not. If you believe that it is blasphemous to urinate on graves and I do not, we have a propositional disagreement. If one of us rejects the concept of blasphemy altogether (for whatever reason), we have conceptual disagreement. Where I have described mere differences in states, attitudes, contents, etc. as superficial, or perhaps *thin* notions of disagreement, Stroud may think of these instead as examples of where the distinction is drawn. Recall the original example I used to motivate my first attempt at the robust/superficial distinction:

**Vegan/Vegetarian 3:**

“I read Peter Singer’s moral arguments for veganism recently and found them persuasive enough that I have since become a vegan.”

“Fair enough, I’m still unconvinced that the moral reasons for veganism are strong enough.”

On the face of it, the disagreement is *either* over whether there are moral reasons for veganism, *or* whether those reasons are sufficiently persuasive or strong. On Stroud’s account, we may say that both are plausibly straightforward propositional disagreements, but we can also reimagine this such that it instead creates a conceptual disagreement regarding the notion of moral strength or persuasiveness with regards to pro-veganism arguments. Perhaps the latter speaker *is* a vegan, but for non-moral reasons, perhaps it is not that they are unmoved by the strength of the moral arguments for veganism, but that they are an error theorist about morality. It is possible that some of the accounts or examples of agreement and disagreement I regard as superficial might simply be conceptual disagreements in the sense Stroud proposes. It may also be the case that it simply does not matter, that there is no clash or conflict with Stroud’s account to suggest that

---

<sup>91</sup> Ibid; 21.

some minimal likenesses or differences in attitudes are superficial cases of agreement and disagreement, after all.

A second pass attempt at capturing the notion of robustness or superficiality aligns with the third of the accounts listed above, which reduces robustness and superficiality to the *significance* or *mundaneness* of the subject matter. My intuition is not merely that a superficial agreement or disagreement case is any case involving an innocuous or mundane topic. Rather, robust agreement and disagreement cases are simply those which we value higher, or which strike us as more important, interesting, or as having higher stakes, those which ‘flag up’ more in our reasoning or attitude-forming. We are disposed to pay attention to certain topics, subjects, or disputes more than others, and perhaps this is in part due to them fitting into our personal interests, but it might also be that those topics and the subsequent agreements and disagreements are richer and more complex. In this sense, some disagreements are superficial because they lack depth or personal significance and because they do not ‘flag up’, but this does not undermine the fact that they are still genuine, legitimate instances of agreement and disagreement.

However, even this seems to fail to fully capture the sense of robustness I am interested in. I might have a robust, clear disagreement with a speaker, and we may both fully comprehend the subject matter and the meanings we assign to our utterances. For example, we may disagree over how many grains of sand there are on the beaches of Whitstable. Such a disagreement may or may not be superficial in the first pass attempt I offered earlier, in that it might involve mere differences in attitude with no underlying difference of reasoning or evidence. It is not superficial in the first and second sense above, but it could quite plausibly be regarded as superficial in the sense that tracks importance or significance, since the question of how many grains of sand Whitstable has will likely only matter for a select number of individuals, it need not matter much for the parties involved in the disagreement. Regardless, neither of my attempts at differentiating robust and substantial agreement and disagreement suffice to tell us anything about the nature of agreement and disagreement itself. The labels of robust and superficial are qualifiers that we may ascribe to agreement and disagreement cases, independently from whether we can confidently call them agreement and disagreement cases and why we can or cannot. None of the accounts of robustness and superficiality I have offered or discussed should therefore be read as necessary or sufficient conditions for agreement or disagreement. However, I do believe that more can be made of the robust/superficial distinction, perhaps considering whether any plausible metric is attainable to demarcate or rank cases or kinds of agreement and disagreement.

### 3.9: Concept Models and the Concepts of “Agreement” and “Disagreement”

The accounts of agreement and disagreement that I have considered have not all explicitly presented agreement and disagreement in terms of necessary and sufficient conditions. However, for most, if not all of these narrow accounts of agreement and disagreement, there is at least an implicit assumption that the factor or state emphasized by a given account is regarded as at least sufficient for many cases and perhaps necessary for some. That is, it can be inferred from any of the accounts I have considered that they are attempting to provide necessary or sufficient conditions for agreement and/or disagreement. My broader reading of agreement and disagreement acknowledges these various attempts, but also contends that they are fundamentally flawed because agreement and disagreement are simply too complex, multifaceted, and varies to reduce to a simple and singular account.

Many concepts are vague, complex, or which take multiple forms and thus cannot be adequately understood in terms of necessary or sufficient conditions. We might envisage agreement and disagreement instead in terms of a *genus-species* relation, wherein a general concept is marked from various broadly related species of that concept. Aristotle illustrates the difference in his *Categories*, noting that we can distinguish between non-human animals (or beings that can perceive but are irrational) and humans who are both perceiving and rational beings).<sup>92</sup> More recently, Simon Kirchin has raised the question of whether this suffices to capture the difference between thick and thin evaluative concepts, the latter being those with both normative and descriptive components or implications.<sup>93</sup> Kirchin regards understanding a concept in terms of genus and species requires both knowledge and understanding of a concept that is logically an conceptually prior – the *genus* concept – from which we can derive species concepts if and only *iff* we have sufficient understanding of those concepts. Aristotle’s distinction between rational humans and irrational animals is insufficient because some non-human animals can also be rational, so we need a better understanding of the *differentia* – the uncommon characteristics or features – which often harder than it seems for many concepts. Similarly, the concept of *goodness* is vague and multi-faceted, but also appears to take several forms in kindness, compassion, wisdom, or beauty, for all of which we have clear ideas of what is involved or required. If we commit to the genus-species model, we must have a clear idea of what it is for something to be good, independent from these derived forms, as well as understanding what each species concept has in terms of unique differentia. Kirchin is not convinced that we can simply begin with a concept of ‘good’ or ‘goodness’ and add in particular differentia to generate new species, so the

---

<sup>92</sup> See Studtmann, P. 2021; §2.2.1.

<sup>93</sup> See Kirchin, S. 2017; 1.

model does not work for goodness.<sup>94</sup> I contend that things are much the same with agreement and disagreement. If each kind of agreement or disagreement differed only in terms of, for example, the attitudes involved, then we could clearly say something like ‘agreement and disagreement are phenomena of a type *F*’, and differentiate credal agreement and disagreement cases from doxastic, preference, intention, etc. The differences between case types are more than just attitudinal, or relational, or agential, they are too many for either a genus-species model or necessary and sufficient conditions to be adequate.<sup>95</sup>

An alternative way of understanding concepts is one that can be attributed to W.E. Johnson; similarly to the genus-species model, it links general and specific concepts, but it has no requirement for specific concepts to have separate *differentiae*. We might recognise that “red” is a particular kind of the concept “colour”, but we likely do not understand what it is for something to either have the property “colour” or the specific *differentiae* that mark something as “red”, nor for any other colour:

If the *genus-species* model is applicable, we should be able to claim that RED is the ‘X sort of colour’, or ‘the colour with X-ness’, say, where the uniquely identifying X is something independently intelligible from and conceptually prior to RED. But there seems to be nothing to fill this role...

We cannot pursue any strategy that explains why RED and BLUE are part of the same family in terms of a commonality or commonalities between them, and further say they are distinguished from each other by something unique to them individually, for this is simply a restatement of the *genus-species* model... What makes it the case that RED and BLUE are determinates of a common determinable is that both concepts cannot apply, or be instantiated if one prefers, at the same time in the same place (or at the same time by the same object).<sup>96</sup>

On this *determinable-determinate* model, nothing can instantiate the same determinate concepts simultaneously. Shape and colour are parallel concepts, but they are not sufficiently related that instantiating them generates an incoherence. “Red” differs enough from “scarlet” and “square” that something can be all three at the same time without it being problematic, but something

---

<sup>94</sup> Ibid; 44-46.

<sup>95</sup> This should not be taken as an endorsement of “agreement” and “disagreement” as thick evaluative concepts, although I’m not opposed to the idea.

<sup>96</sup> Kirchin 2017; 46-7.



cannot be “red” and “blue” at the same time in the same physical space or area, because they are parallel concepts in both kind and specificity.

This might work for agreement and disagreement, since there is certainly a similar setup in that the general concepts of “agreement” and “disagreement” might admit of different determinates, specific instances that are perhaps distinct in specificity and kind. This might even account for my argument that some instances or kinds of agreement and disagreement are more robust than others, although my intuition is still that robustness and superficiality are a matter of individual evaluation or of consistency within cases. As Kirchin notes, the determinable-determinate model of concepts treats concepts as unified wholes, such that everything that is “red” is also “coloured.” He contends that this approach to understanding concepts does not fit particularly well with evaluative concepts, that it would be implausible to say that every act, object, or thing which is honest, for example, is also good when surely the range of acts or objects that can be honest surely includes at least some bad cases of honesty (in some sense or other).<sup>97</sup> This essentially reiterates G.E. Moore’s argument against the definability of “goodness”, which is that goodness is too simple and basic to define in terms of other properties, and that we should distinguish the question of what goodness is from the question of what things are good. Moore notes that we can think of goodness in terms of many derived forms such as pleasure, and consequently ask whether some such pleasurable thing is good, to which the answer is not always or necessarily, so this account of goodness is neither fixed nor closed, and therefore not sufficient.<sup>98</sup>

I believe the same can be said of agreement and disagreement. Recall the various factors relating to agreement and disagreement cases which I have considered so far in this thesis; it is not as simple as merely identifying different levels or a taxonomy of agreement and disagreement cases in the way that the determinable-determinate model requires. I contend that they shift and vary too much to fit within this kind of conceptual framework, and indeed as a genus-species model. If we are seeking necessary and sufficient conditions, it seems that the best we can hope for is either a set of conditions that are sufficient, but not necessary, or some parallel situation to what Moore envisages for “goodness.” We might be able to offer up a situation involving a pair of speakers, with attitudes of a certain kind, that concern a clear and mutually recognised topic, and relate in some way I have considered – and we may not still confidently always regard it as “agreement” or “disagreement”. What, then, are the alternatives? One possibility can be found in Wittgenstein’s account of family resemblance.

---

<sup>97</sup> Ibid; 56-7.

<sup>98</sup> See Moore, G.E. 1903; 1:6.

### 3.10: Family Resemblance and Ambiguity

Wittgenstein recognises that there are several concepts about which we cannot identify some common element, essence, component, or defining factor, but which we might confidently say there are several overlapping examples. Two of the more prominent examples Wittgenstein discusses are “language” and “games”:

Instead of pointing out something common to all that we call language, I’m saying that these phenomena have no one thing in common in virtue of which we use the same word for all – but there are many different kinds of *affinity* between them. And on account of this affinity, or these affinities, we call them all “languages”.<sup>99</sup>

Similarly, there are multiple examples of “games” available; some use cards or tabletop boards, some use balls, bats, or racquets, and others require electronic video game consoles. None of these have anything in common beyond the fact that they are regarded or labelled as games, meaning there is no plausible necessary or sufficient condition that will capture all examples of “game.”<sup>100</sup> We can unite general and specific concepts in a way that allows for a degree of flexibility and variety between specific concepts, but we lose any exhaustive defining feature:

Consider for example the activities that we call “games” I mean board-games, card-games, athletic games, and so on. What is common to them all? – Don’t say: “They *must* have something in common, or they would not be called ‘games’” – but *look and see* whether there is anything common to all. – For if you look at them, you won’t see something that is common to *all*, but similarities, affinities, and a whole series of them at that. To repeat: don’t think, but look!<sup>101</sup>

I can think of no better expression to characterize these similarities than ‘family resemblances’; for the various resemblances between members of a family – build, features, colour of eyes, gait, temperament, and so on and so forth – overlap and criss-cross in the same way. – And I shall say: ‘games’ form a family.<sup>102</sup>

---

<sup>99</sup> Wittgenstein, L. 2009 [1953]; §65.

<sup>100</sup> Ibid, §§66-7, 75.

<sup>101</sup> Ibid; §66.

<sup>102</sup> Ibid; §67.

Michael Forster argues that Wittgenstein is not denying the possibility of any family resemblance concept having a singular definitive or common feature, but rather that any such feature will justify every possible application of a concept to its many instances or examples.<sup>103</sup> In other words, we cannot always offer definitions of notions, concepts, or terms through identifying a common feature or element, sometimes all we can do is agree to apply these terms or concepts in the same general ways, giving us a rough and very general idea of what they refer to. Importantly, this is not the same as maintaining that family resemblance concepts are polysemic terms; rather, the difference lies in the fact that family resemblance concepts have multiple different meanings which gradually converge or combine, whereas polysemic concepts have several distinct meanings.<sup>104</sup> One might worry that it is because of the various instances of games, language, et al. bearing some shared features that Wittgenstein thinks they are a family, when he does not in fact explicitly say this, as Heather Gert notes. Rather, “family-making relations aren’t necessarily relations of resemblance”, it is the connections and relations between different instances or cases that make then a family, not the fact that they may each look alike in some respects.<sup>105</sup> Members of a family have their shared features in virtue of being part of the same family, not the other way around...

While this provides a better way of understanding agreement and disagreement than trying to discern necessary conditions, there are many concepts for which we could just as easily find such patterns; perhaps concepts such as justice or knowledge are family resemblance concepts. As such, I believe that agreement and disagreement are best understood as general, but *ambiguous* concepts. Importantly, this is not the same as saying that they are *vague* concepts since for a concept to be vague, it cannot be fully specified but it must admit of borderline cases (i.e., discernible instances where it is unclear where on the border it falls.) We have clear ideas of when there is agreement or disagreement and when there is not, but I contend that we cannot clearly define all of them with simple or narrow definitions because there are too many shifting factors. Rather than picking out a singular object, agreement and disagreement seem to pick out multiple objects or cases, and thus they are (or at least appear to be) *general* concepts. Agreement and disagreement are also *ambiguous* concepts because they have (or appear to pick out) multiple definitions. Just as a ‘bank’ can refer to both a financial institution and a riverside, agreement and disagreement can seem to refer to various relations between attitudes or expressions of attitudes.<sup>106</sup> Some cases involve distinct likeness, difference, or polarisation, and some involve

---

<sup>103</sup> See Forster, M. 2010; 69.

<sup>104</sup> Ibid; 70.

<sup>105</sup> Gert, H. 1995; 180.

<sup>106</sup> Sennet, A. 2023.

hypothetical or actual agents. We cannot capture all agreement and disagreement cases with narrow definitions that use simple terms or ideas, but we can say that ‘agreement’ and ‘disagreement’ pick out multiple relations, interactions, expressions, and so on. I contend this is the best way to understand agreement and disagreement. We have at least a rough idea of what is going in most circumstances in which we agree or disagree with some proposal or person, and sometimes there are similarities between specific cases that can be identified, but these similarities will only run so deep. In other words, no narrowly specified and unified kind or *genus* of agreement or disagreement can be discerned or identified.<sup>107</sup> At most, we can have a broad spectrum of kinds, cases, and instances, some or many of which will overlap in various respects, but perhaps not all.

Perhaps I am being overly pessimistic or permissive. It might be that my argument includes some cases which are not especially intuitive, or perhaps there are clear ways to differentiate between certain cases such that there *are* species of agreement and disagreement. I believe I can shore up my position a little more. I contend that necessity and sufficiency should both be understood in terms of a relation of entailment, albeit in different directions. A necessary condition is a non-negotiable element of a concept, object, or phenomenon, one which ensures whether one has it or not. A sufficient condition is that which is ensured, either individually or collectively, by the thing for which it is sufficient: Going to a train station is necessary for taking the train to London, whereas taking a train is sufficient but not necessary for getting to London because there are many alternative optional means of getting to London. The factors I have discussed might be sufficient for agreement and disagreement in general, if not necessary for specific cases or kinds; indeed, some of them may be synonymous for some folk notions of agreement or disagreement. Ultimately, however, I do not believe there are any plausible candidates for a necessary condition for agreement or disagreement, at least not one that explains or accommodates all possible kinds and cases, without which we cannot say for certain that there is agreement or disagreement. Wittgenstein thought that for certain other concepts this was also the case, and without any common feature that unifies all examples, it makes no sense to think of them as determined by necessary conditions or as a genus concept, but instead as mere family resemblance concepts. Different examples or cases will have some common factors, but no single element or characteristic can be attributed to *every* possible case.

---

<sup>107</sup> As far as I am aware, there is no significant discussion of *genus-species* or determinable-determinate in the literature on discussion. Both Mike Ridge’s 2013 and Graham Bex-Priestley and Yonatan Shemmer’s 2017 make passing references to there being a *genus* of disagreement, or of various kinds of disagreement being species of that *genus*, but neither offers any clarified or detailed defence of this claim.

For all the consideration that moral philosophers, philosophers of language, and social epistemologists have given disagreement (and agreement, but less so), little understanding of what agreement and disagreement *simpliciter* consist in has been achieved, I contend, for just this reason. Indeed, much of the discussion of disagreement in social epistemology has seemingly overlooked the matter altogether. Sometimes a given conversation might look more like agreement than disagreement, other times it will look more like disagreement than agreement. Some cases will not even involve conversation but can be attributed to the mere pair of attitudes or states that are in accord or tension with each other. Finally, sometimes agreement and disagreement cases will be indeterminate or obscure, and other times they will involve multiple overlapping senses of agreement and disagreement simultaneously. Nonetheless, trying to posit a single, uniform, all-encompassing definition that is concise and specific is admirable, but futile and largely a waste of effort. As I have noted over the previous chapters, we can offer speculative or narrowly specified accounts of agreement and disagreement, but these will do no better than to capture some small subsets of cases, certainly not all of them. Likewise, the phenomena of agreement and disagreement are too broad to posit them as *sui generis* concepts, as Moore suggests for goodness. Instead, I return to the two descriptions I offered at the outset: Agreement and disagreement are far harder to define or specify than we might think, *because* they are complex, multi-faceted phenomena.

## Part 2: Epistemic Peerage, Agreement, and Disagreement

The previous part of this thesis was concerned with discussing the nature of agreement and disagreement at length. I take myself to have shown simply that it is incredibly difficult to capture *all* kinds or examples of agreement and disagreement with a single account or set of conditions. It is easy enough to point to certain cases, examples, or kinds, but there is simply no common feature or element that unites them all, a vague picture that accommodates various exchange-types or state-pairs are the best that we can get. While I have judged that isolating the essence or nature of agreement and disagreement proves somewhat inconclusive, this need not mean there is *nothing* philosophically substantial or interesting that can be said about them.

Over the second part of this thesis, I address a specific type of disagreement, one which has received substantial attention in social epistemology. I contend that it would not be improper to equivocate the epistemology of disagreement (or certainly a vast majority of it) with the discourse I am interested in – that concerning disagreement between epistemic peers or *peer disagreement*. Put simply, the debate is tasked primarily with identifying or demonstrating what is epistemically significant about disagreement. Those who write on peerage and peer disagreement think we can do this by establishing what response is rational or rationally required once epistemic peers acknowledge that they disagree with each other. An epistemic peer is any speaker who is, or is perceived as, in an equal and symmetric epistemic position to us, and various metrics have been offered for ascertaining equality. Generally, these include equality in some or all the following: physiological and intellectual capacities, evidence, epistemic virtues and freedom from biases or vices.

Opinions are generally divided between *steadfast* views that advocate retaining one's opinion, attitude, or degree of confidence, and *conciliatory* views that favour concession, compromise, or suspension of belief. While the literature and discourse purports to demonstrate and discuss the epistemic significance of disagreement, I contend that it largely fails to do this. Rather, for the most part, the discussion is chiefly about epistemic rationality in niche circumstances that are heavily idealised and unrealistic, failing to adequately reflect real-life human capacities, resources, or experiences and, indeed, disagreement cases. The peer disagreement discourse has preconceived notions of epistemic significance which I argue are mistaken and it misrepresents disagreement in an entirely implausible way to fit those preconceptions. There is much talk about ideal cases of disagreement, but we learn very little from it about real life cases and, subsequently, why disagreement is epistemically significant. Furthermore, while the focus is on disagreement's epistemic significance, agreement is largely sidelined and constrained to cases involving groups who agree with each other. While we have

seen that *some* accounts of disagreement with counterparts for agreement, the peer disagreement discourse does not reflect this tendency.

Given all this, it is not clear to me *why* peer disagreement has garnered so much attention. Much is assumed about what it takes for speakers to be epistemic peers, and how easily they are identified or encountered, so the significance of peer disagreement is overestimated and exaggerated. Epistemic peerage is not especially common because the various conditions for peerage are nearly impossible to meet collectively or simultaneously. Even taken individually, they are certainly hard to meet very often. Peer disagreement is neither the route that the literature maintains it is to establish why disagreement epistemically matters, nor is it as interesting and fundamentally significant as it seems. Peer disagreement cases ensure stimulating discussions as hypothetical cases, but it is misleading to present them as in any way realistic, attainable, or radically more important than any non-peer disagreement, which are certainly more common and more easily imagined.<sup>108</sup> What I imagine the discourse is doing is *attempting* to steelman – use the strongest version of an argument to boost its credibility – and I argue it is incapable of succeeding. Moreover, two other important and interesting debates and concepts are subsumed by the peer disagreement discourse to *try* to strengthen claims about how to respond to peer disagreement, which I will introduce in what follows.

---

<sup>108</sup> I.e., disagreement with an ‘epistemic inferior’ or ‘epistemic superior’.



#### Chapter 4: Three Debates Plot a Terrain

Let me begin with a typical example from the peer disagreement literature, which will no doubt be a scenario that is relatable and familiar, at least in principle:

**Restaurant Check:** Suppose that five of us go out to dinner. It's time to pay the check, so the question we're interested in is how much we each owe. We can all see the bill total clearly, we all agree to give a 20 percent tip, and we further agree to split the whole cost evenly, not worrying over who asked for imported water, or skipped dessert, or drank more of the wine. I do the math in my head and become highly confident that our shares are \$43 each. Meanwhile, my friend does the math in her head and becomes highly confident that our shares are \$45 each.<sup>109</sup>

This is arguably the most cited and discussed example in the literature on peer disagreement. Supposedly, the two friends in this example are epistemic peers, which broadly means they are equals along some or all of the following metrics and respects:

*Evidence* – Evidential equality, relative to the question whether  $p$ , occurs when two agents share mutual evidence or are (otherwise) familiar with the same relevant evidence.<sup>110</sup>

*Cognitive* – Cognitive equality occurs when two agents are equally competent and intelligent and have the same faculties for processing evidence (e.g., perception).<sup>111</sup>

*Dispositional* – Two agents are of equal disposition if they are both equally disposed to respond to evidence in the same way, or in an epistemically appropriate way.<sup>112</sup>

Alternatively, disposition may be conceived in terms of a shared distribution of intellectual virtues and vices, such as honesty and freedom from bias.<sup>113</sup>

*Accuracy* – two agents are equal, relative to the question whether  $p$ , if both are equally likely to be accurate in their respective belief about  $p$ .<sup>114</sup>

Speakers are assumed, by virtue of being peers, to be equally likely accurate in their judgment, equally acquainted with the evidence that is relevant to the matter at hand, and equally capable of

---

<sup>109</sup> Christensen, D. 2007; 193.

<sup>110</sup> See Kelly, T. 2010; 112. See also Christensen, D. 2009; 757.

It's worth noting that access to evidence need not entail extensive familiarity with it. Two speakers could *just now* have been presented with all their evidence for the first time and therefore still be equals in this sense.

<sup>111</sup> Kelly 2010; 113. Oppy, G. 2010; 185.

<sup>112</sup> See King, N.L. 2012; 251.

<sup>113</sup> See Kelly 2005; 174-5.

<sup>114</sup> See Conee, E. 2009; 313. See also Kelly 2010; 112, and Enoch 2010; 956.

grasping it and reasoning well based on that evidence.<sup>115</sup> It is probably easier to imagine this for something like the Restaurant Check example because there is only a limited amount of evidence that relates to the share of the group's bill – the bill itself and whether anyone's capacity for calculating is notably compromised for whatever reason. The thought is that because epistemic peers are epistemic equals, each peer should consider the fact that they disagree with each other as a strong challenge to their respective attitude and to their rationality.

At least one of the diners in the Restaurant Check example must have miscalculated, but they could both be wrong. They are considered peers typically because among other reasons they have equal exposure to the relevant evidence, are equally capable and likely to have calculated correctly, and perhaps have a track record of doing so:

As I will use the term, the class of epistemic peers with respect to a given question are equals, not only with respect to their possession of... general epistemic virtues... but also with respect to their exposure to evidence and arguments which bear on the question at issue.<sup>116</sup>

Suppose that you and your friend independently evaluate the same factual claim... Each of you have access to the same crime statistics, sociological reports, and so on, and has no other relevant evidence. Furthermore, you count your friend as an **epistemic peer** – as being as good as you at evaluating such claims.<sup>117</sup>

Peerhood is sometimes presented in terms of an equiprobability of making true assertions regarding a given domain. More often, it includes some specific reference to the possession of the same evidence, and equality of reasoning capacities relative to that evidence.<sup>118</sup>

The discourse focuses on agents, speakers, and participants who are equals in various respects and who happen to draw opposing or contradictory judgments about a matter. Importantly, epistemic peers are not just equally intelligent in general, they are equally well-informed about the matter at hand, equally intellectually virtuous and equally likely to be accurate. The conditions

---

<sup>115</sup> As Adam Elga notes (2007; footnote 14), it is more often assumed that speakers *count* or *regard* each other as epistemic peers, which is notably different from cases where speakers factually, undeniably *are* peers. Elga thinks that it is okay to assume that peerage is a matter of perception or regard, even if one's evidence for regarding someone as a peer is misleading and causes one's assessment to turn out inaccurate.

<sup>116</sup> Kelly 2005; footnote 2.

<sup>117</sup> Elga 2007; 484.

<sup>118</sup> Kenyon, T. 2020; 70.

for peerage have expanded as the discourse has developed and gained interest, so more recent accounts or discussions specify further conditions for peerage. However, at least some acknowledge that having a greater number of conditions makes peerage a somewhat rarer relation.<sup>119</sup>

There are other, more complex examples of potential peer disagreements besides the Restaurant Check. Two meteorologists who have had the same training may have different expectations about the likelihood of rain, or two economists may differ over whether interests will increase or decrease.<sup>120</sup> Where discussion of disagreement in ethics might be tied to moral realism and anti-realism, or to linguistic and alethic pluralism in the philosophy of language, the peer disagreement discourse aims to highlight why disagreement matters *epistemically*. I argue that it is more concerned with the epistemic rationality of peers, and expectations regarding their judgments, attitudes, or degrees of belief once peer disagreement is acknowledged. That is, for any purported disagreement between epistemic peers A and B over a matter or proposition p, the chief question in the peer disagreement debate is what I call the “response question”: how would it be rational or rationally required for peers A and B to respond, once it is revealed that they disagree over p or once they discover their disagreement over p? Importantly, there are different dimensions to the response question, since the peers in question might be literal peers or may simply regard each other as peers, they may perceive their disagreement but not genuinely disagree (or their disagreement may not be apparent or revealed), and the question could be asking what is rationally optimal, as opposed to merely what is rationally required.<sup>121</sup> Because peerage is a relation of equality between speakers with respect to their epistemic capacities and resources, their epistemic situation is regarded as symmetric; however, this need not mean peers are regarded as each other’s mirror or clone, *because* they have differing and conflicting judgments or attitudes. All their evidence is mutual or shared, bracketing any inequalities or unique individual traits that might otherwise favour one speaker’s judgment over the other’s.

The response question is the chief question in the peer disagreement discourse, but the main answers that are often suggested and discussed draw on other debates that do not directly concern peer disagreement for motivation and support. Different authors have offered defences either of retaining one’s attitude, conviction, or belief (*steadfastness*) or of deferring to one’s opponent, compromising, or suspending judgment (*conciliation*). I argue later that the arguments for various responses are unconvincing and each face various common flaws; I also contend that

---

<sup>119</sup> See, for example, Lackey, J. 2010, and King 2012.

<sup>120</sup> The former is also from Christensen 2007; 194. The latter is from Feldman, R. and Warfield, T. 2010; 2.

<sup>121</sup> See Kirchin, S. and Pethick, S. (ms)

two other debates in epistemology are subsumed or crowded out by peer disagreement. In what remains of this chapter, I will summarise these two debates.

#### 4.1: Uniqueness and Permissiveness of Evidence

Part of the interest in peer disagreement can be attributed to the debate about the *uniqueness* of evidence, which concerns what it is rational to believe based either on individual items of evidence or a collective evidence body. In short, the division is between those who think that only one attitude or judgment follows from any piece or body of evidence (uniqueness), and those who think evidence permits more than one attitude (permissiveness). We need an idea of whether evidence is *unique* or *permissive* to make sense of how we may be rationally permitted to respond to peer disagreement, I contend that the main position in the response question assumes a strict reading of uniqueness. This is only a broad and general overview; the two contender positions can be endorsed to different degrees and there are other questions that can be asked which affect what one might say, such as:

- Are there plausibly any circumstances in which it is rationally permissible, given one's available evidence, to both believe and disbelieve a particular claim?
- Are there plausibly any circumstances wherein one faces a rational dilemma because *no* available attitude is rationally permissible or required?
- Does the justification of an attitude entail (or merely imply) that an agent is justified in holding that attitude?
- Are attitudes understood or characterized in terms of whole belief-states or degrees of belief?
- Do individual pieces of evidence permit a given response or responses, or does justification only follow from collections or bodies of evidence?
- Does the evidence that provides justification refer to all possible evidence, or just that which is available to an agent at a particular time?
- Does the justification provided by evidence apply only to a given agent at any one time, or to all agents?
- Is the justification for an attitude limited only to a given time, or does it extend over many or all possible times?

The most general formulations of uniqueness and permissiveness will not specify much about the type or quantity of attitude, nor about which agents, time slices, or evidence is involved.

However, a detailed examination of the debate is not necessary for my purposes here, so a general formulation of the two replies will suffice for our understanding:

**Uniqueness:** Given one's evidence, there is a unique rational attitude that one can take to a given proposition.

**Permissiveness:** Given one's total evidence, multiple rational doxastic attitudes are permitted with regards to a given proposition.

Again, these are general glosses of the debate, although more specific claims can be attributed to the earlier literature. For example, Richard Feldman presents Uniqueness as:

the idea that a body of evidence justifies at *most one* proposition out of a competing set of propositions... and it justifies at most one attitude toward any particular proposition.<sup>122</sup>

In contrast, Roger White contends that those who endorse Uniqueness are saying that:

given one's *total* evidence, there is a unique rational *doxastic* attitude that one can take to any proposition.<sup>123</sup>

Thinking again about the questions above, we can see that Feldman's version is broader than White's because it is concerned with both the justification for a proposition *and* the attitudes held regarding them. White's version of Uniqueness is just about the unique rationality of people taking an attitude, rather than about the attitudes themselves. Other readings that are less stringent or more flexible around particular domains or topics are also available. David Enoch contends that:

The Uniqueness Thesis is the claim that, with given evidence, there is a unique degree of belief that is maximally rational. So in order to deny this claim, one has to assert that there is no one degree of belief that is uniquely maximally rational, how about the view according to which there *is* one degree of belief that is *maximally* rational, but some other

---

<sup>122</sup> Feldman, R. 2006; 205. My emphasis.

<sup>123</sup> White, R. 2005; 455. My emphasis.

degrees of belief – though less than maximally rational – are still rationally permissible?<sup>124125</sup>

And, likewise, there might be more radical readings of Permissiveness available, as White notes:

**Extreme Permissivism:** There are possible cases in which you can rationally believe *P*, yet it is consistent with your being fully rational and possessing your current evidence that you believe not-*P* instead.<sup>126</sup>

We have a very general question, then: what can we rationally believe, given our relevant evidence? While this may appear simple, the questions above show that it rests on various factors. Permissiveness in its most extreme formulation will grant pluralism about any matter at any time, permitting one to be entirely self-contradictory without conceding any inconsistency. Weaker versions will instead maintain that there are some occasions where it is permissible read evidence in more than one way and remain rationally consistent, but that not all cases permit this. In opposition to the most extreme form of permissiveness, the strictest version of uniqueness is committed to objectivity about rationality, such that the attitude to be drawn from each piece or body of evidence is (perhaps) a factual matter. Individually, the pieces of evidence *E*, *E*<sub>1</sub>, and *E*<sub>2</sub> might each warrant one attitude *A*, but they might be overridden or defeated by a different item *E*<sub>3</sub>, so as a combined body of evidence together they warrant a different attitude *B*.

#### 4.2: Uniqueness and Peer Disagreement

With a rough understanding of the uniqueness and peer disagreement debates in hand, let me now explain how they relate and overlap. The response question, we may suppose, can be reinterpreted as asking whether either or both speakers are equally rational in having inconsistent or contradictory beliefs. The relevant notion of rationality comes in part from considerations of uniqueness, it is about whether a belief is evidentially supported or justified, and not merely true or correct.<sup>127</sup> We might think it is redundant when we disagree with a speaker who we either do not know well, or do not know at all, whereas we might think something is amiss when we realise

---

<sup>124</sup> Enoch 2010; footnote 9.

<sup>125</sup> See also Jonathan Matheson's description of the Underdetermination Thesis, according to which 'for any body of evidence there are multiple competitor theories that are all equally well confirmed by that evidence.' 2015; 52-53.

<sup>126</sup> White 2005; 447.

<sup>127</sup> Even this is only one way of interpreting a somewhat vague and unspecified concept. I read it in terms of the justification for a belief or reasons to have a relevant attitude here, rather than its truth, accuracy, or fittingness. See Greco, D. and Hedden, B. 2016.

that we disagree with a speaker who we trust and regard as both intellectually considerate and more intelligent. This will hold for many agents, but the peer disagreement discourse is premised on the thought that there is something specifically important about disagreeing with a speaker who is our equal in all the relevant ways, with someone who has no discernible or disclosed advantage or disadvantage.

The link between peer disagreement and uniqueness is most apparent by asking if both speakers are rational in having or maintaining their inconsistent attitudes from the same shared or common evidence. Gideon Rosen contends that that several longstanding disputes are like this, since many disagreements in history or science have occurred between speakers who are equally intelligent and thoughtful and who are all still regarded as rational despite having different views.<sup>128</sup>

Sometimes disagreement causes us to doubt the accuracy and rationality of a judgment or an attitude, other times it has no such impact at all. If there is a disagreement over a matter that is not open-ended, indeterminate, or pluralistic, based on evidence that is both unique and commonly available, at least one speaker has either made a mistake. That is, they have failed to be as rational as they *could* have been, and they lack the attitude that is either most rational or rationally required.

However, if evidence is permissive and allows multiple responses, peer disagreement is less worrisome and the need for anyone to reconsider their judgment is weakened *because* evidence is permissive. The uniqueness debate therefore determines what response to peer disagreement is rationally permitted or required, and whether the peers involved are rational after establishing their disagreement. In a truly epistemically symmetric situation, no speaker will have any reason to favour their own position that their opponent has not already defeated or does not also regard as supporting their own contrary position. Again, peer disagreements are set up such that all the propositional evidence is shared or disclosed, but this fact could just as easily suggest that the evidence is inconclusive. With neither speaker having any evidence that exclusively favours their view, it seems that they both must reduce their confidence or conviction and shift towards some form of compromise.<sup>129</sup>

This view – the Equal Weight View – is one of various conciliatory or *conciliationist* positions in the peer disagreement discourse, which are arguably the most popular or centrally discussed views in the discourse. A conciliatory view simply maintains that, based purely on the discovery or revelation of their disagreement, epistemic peers ought to change their beliefs or reduce the confidence or conviction they each have in their respective view. Equal Weight

---

<sup>128</sup> See Rosen, G. 2001; 71-2.

<sup>129</sup> See Matheson, J. 2015; 54.



maintains that speakers in a peer disagreement should each regard the other's view as equally important and likely true as their own; for many, the upshot is that they should split the difference between the two views, which generally means suspending belief indefinitely:

According to the equal weight view, one should give the same weight to one's own assessments as one gives to the assessments of those one counts as one's epistemic peers. If the view is right for the case of one peer, it surely is also right for the case of many peers.<sup>130</sup>

I argue that the pressure to conciliate in the Equal Weight view comes largely from granting evidential Uniqueness.<sup>131</sup> Many who accept one are inclined to agree with the other, and to think that disagreeing peers should revise their belief or confidence to an identical degree. If a less stringent version of uniqueness (such as Enoch's) is granted or true, then the pressure to conciliate and/or suspend judgment is reduced. Again, if the shared evidence is permissive and multiple attitudes are deemed either equally rational or merely rational in some non-problematic way, then the matter is moot because disagreement of most sorts will seem benign.

While there are few defences of permissiveness within the peer disagreement literature, the uniqueness discourse has plenty of advocates on both sides. Roger White worries that any advocate of permissiveness will think of belief-formation and rationalisation as no different from arbitrary methods such as taking a belief-inducing pill.<sup>132</sup> In response, Miriam Schoenfield claims that this sort of argument misinterprets permissiveness, maintaining that the uniqueness debate is instead about what beliefs are reasonable given one's epistemic standards. She claims that we can vindicate permissiveness by noting that different agents will have different standards that will at least sometimes warrant responding to the same evidence in different ways.<sup>133</sup> I will offer my thoughts on the uniqueness debate and its presence in the peer disagreement discourse in the next chapter, for now I turn to the other debate I regard as tied up in the peer disagreement discourse.

#### 4.3: What is Higher Order Evidence?

Much of the discussion so far has been about first order propositional evidence, which is any evidence that supports a propositional attitude such as belief or disbelief or, perhaps, credences

---

<sup>130</sup> Elga, A. 2007; 484.

<sup>131</sup> This has been endorsed by Thomas Kelly, although he contends that both Uniqueness and Equal Weight are "an extremely strong and unobvious claim." See his 2020; 121. See also Ballantyne, N. and Coffman, E.J. 2012.

<sup>132</sup> See White, R. 2005; 451.

<sup>133</sup> See Schoenfield, M. 2014; 200.

and degrees of confidence. Imagine that a trusted friend, relation, or a colleague, suggests that your first order evidence is unreliable or was acquired through questionable means, or just that your reasoning was hindered or impeded in some way. Evidence of this sort is typically called *higher order evidence* and relates to how reliable one's attitudes or the way one acquires or forms them is.<sup>134</sup> It might consist of, for example, facts about the circumstances within which one formed a belief, or about what constitutes sufficiently reliable and relevant first order evidence.<sup>135</sup> Higher order and first order evidence differ not just in what type of evidence comprises them, but also the way in which they facilitate defeat of a particular belief or attitude. A belief is defeated by first order evidence when its truth or fittingness is challenged or undermined; that is, when there is overwhelming evidence in favour of denying the belief or in favour of suspending belief. Higher order evidence tackles the justification for a belief and defeats the belief if it "rationalises a change of belief precisely because it indicates... evidence of my own rational failure."<sup>136</sup> This kind of evidence can take various forms, as Maria Lasonen-Aarnio notes:

It may be evidence that one is subject to a deep but undetectable cognitive malfunction; that one has made a simple calculation error; that one has failed to appreciate the import of one's evidence; or even that the epistemic rules one follows are incorrect.<sup>137</sup>

Imagine I perceive my friend driving a car that appears to be a BMW; I may take that as relevant first order evidence for the belief that my friend owns a BMW. But I might later recall other important information or facts that are relevant, such as:

- i. The fact that I was not wearing my prescription glasses,
- ii. The fact that I am not especially knowledgeable regarding car makes and models, and
- iii. The fact that I have worked several long days in a row and am rather tired at this time.

---

<sup>134</sup> Some (e.g., David Christensen 2010; 212-3) maintain it is a toxic type of evidence that almost always compels us to violate some rational epistemic ideal and reminds us that we are rarely in a position to know for certain that we are meeting these rational ideals.

<sup>135</sup> There is a further understanding of higher order evidence that relates to the (implicit or explicit) endorsement or platforming of speakers. The thought is that if I grant someone an audience or allow them to voice their views, I ostensibly provide higher order evidence that their views are worth hearing out (at least). Neil Levy discusses this reading and offers an insightful defence of no-platforming based on it in his 2019.

<sup>136</sup> Christensen 2010; 185.

<sup>137</sup> Lasonen-Aarnio, M. 2014; 315.

These may all undermine how reliable my belief that my friend owns a BMW is, as well as other beliefs such as my belief last night that she was driving a BMW, or more specific beliefs such as the type of BMW I believe she drives/was driving.

This should sufficiently clarify how the two kinds of evidence differ, but one might wonder why higher order evidence and higher order defeat are so important. As fallible beings, we all are prone to relying on each other, sometimes we quite doggedly compare other people's judgments and reasoning with our own. Higher order evidence matters because it demonstrates a person's reliability and generates a burden of proof. If we cannot show that an attitude has been reliably formed or well considered, we cannot and should not hold that attitude, regardless of how confident we are that it is true. More importantly, we cannot rely on the agent who formed it either. This holds true as much for our own attitudes as for that of other people. If we have reason to believe that our capacities and faculties are functioning poorly or insufficiently, or that our evidence is flimsy or misleading, and we cannot prove otherwise, then no matter how well we think we have reasoned or judged, or how confident we are in ourselves, we should adopt a different attitude. It is essential that beliefs in high stakes situations should be well-reasoned and reliably formed or acquired as much as possible and, likewise, when one's available first order evidence is limited, they should scrutinise it strongly. We should avoid accepting a judgment from an agent whose evidence is compromised or whose reasoning and faculties clearly lack integrity without first giving it due consideration.

#### 4.4: The Import of Higher Order Defeat

The higher order evidence discourse is *chiefly* concerned with the circumstances in which higher order evidence bears on our judgments or attitudes by defeating our reasoning and showing our attitudes to be unreliable. We are often more likely than not to find reasons to doubt our judgment and higher order uncertainty about what attitudes to hold or count as rational, so we cannot be self-assured or confident in all our attitudes, perhaps not even in many of them.<sup>138139</sup> Given the apparent pervasiveness of higher order defeating evidence, then, the two primary questions are as follows:

1. How strong should we regard higher order evidence as a defeater, compared to first order defeaters?

---

<sup>138</sup> Dorst, K. 2019; 35.

<sup>139</sup> This perhaps paints a picture of things that is somewhat bleak; we can certainly say that higher order supporting or confirming evidence exists as well, and even that it might be quite common. But much of the focus concerns higher order defeat instead.

2. How should we respond when we do acquire or attain higher order defeating evidence?

A third consideration is whether higher order evidence is common enough that it prevents actions that are rationally or epistemically *akratic*, meaning whether agents are prevented from knowingly performing actions or holding attitudes that are irrational. If higher order supporting evidence is as prevalent and strong as first order evidence in its defeating capacity, then we can't ever feasibly be epistemically akratic. Some have argued that akratic action is at least partially inevitable and that it is senseless to think that there could be a plausible obligation to avoid it, let alone a *prima facie* norm.<sup>140</sup> Consider some more examples:

**Book:**

Imagine I am reading about the changes in birth and death rates during the later Middle Ages, from a book that has been written by a renowned historian. Unbeknownst to me, most of the book was fabricated by the author and very little of it is in any way historically accurate.

**Taxi:**

I am trying to recall how I got home last night when I was somewhat intoxicated. Ordinarily I have a strong memory, and I recall phoning for a taxi last night rather than wanting to walk or drive home.

In the first case, we might think that the author's reputation and prestige as a historical writer makes it reasonable for me to believe based on what I've read from their book. While it might have provided a lot of evidence that supports the beliefs that I drew from reading it, the fact that the book's content is unfounded and misleading means that, while my reasoning is sound, ultimately it is unreliable. In the second case, if I remember the next day that I was drunk the night before, it will likely undermine my memory and give me reason to doubt that I actually took a taxi home. At the very least, it means I have no more reason to believe that over anything else.<sup>141</sup> This doubt could in turn be reinforced by my friend's later assertion that she witnessed me choosing to walk home rather than take the cab.

---

<sup>140</sup> See Lord, E. and Sylvan, K. 2021; 117.

<sup>141</sup> As Thomas Kelly points out, higher order considerations regarding one's state of competence when engaging in *practical* reasoning are not themselves reasons for or against choosing one act over another. As will become apparent, I'm not convinced that such considerations actually give us reason to form particular beliefs or attitudes either. See his 2010; 140.

Ultimately, facts such as the book being fabricated, me being drunk, and others like them can call into question how reliable my beliefs are, although there are certainly other cases in which higher order evidence confirms or supports my beliefs rather than defeating them. Finally, the following example comes from the higher order evidence literature itself and is discussed often:

### ***Drugs***

I'm asked to be a subject in an experiment. Subjects are given a drug, and then asked to draw conclusions about simply logical puzzles. The drug has been shown to degrade people's performance in just this type of task quite sharply. In fact, the 80% of people who are susceptible to the drug can understand the parameters of the puzzle correctly, but their logic-puzzle reasoning is so impaired that they almost invariably come up with the wrong answers. Interestingly, the drug leaves people feeling quite normal, and they don't notice any impairment.<sup>142</sup>

The example goes on to describe my being given coffee that has been laced with the drug before then doing a practice puzzle. I do not know that I've been given the drug straight away, but I nonetheless proceed through the logic puzzle and become extremely confident in my answer. I am only informed afterwards about the drug being administered, which then causes me to reduce my confidence. Because it appears no more or less likely that my reasoning has been impaired than not, it is arguably more reasonable to believe that there has been an epistemic malfunction on my part, and that any sound judgments are fortunate at best.

Examples like this and the ones above generate similar questions to peer disagreement, the chief question being one of response: what is it rational to believe once one receives or acquires higher order defeating or undermining evidence? For David Christensen, the answer to the puzzle in the Drugs example is not affected by the drug being slipped into my coffee or my being informed of this fact. If I were immune to the drug, it would not be obvious why my total evidence, which now includes the fact that I have been given the drug, no longer supports my original belief less than my initial evidence. I can still see that there is a relation of entailment between my evidence in the logic puzzle and my first conclusion, so whatever evidence that remains which is not in doubt will still support my initial answer. However, Christensen notes that it would be absurd and irrational for me to presume my immunity with no supporting evidence and instead argues that it is rational to lower the confidence I have in my belief about

---

<sup>142</sup> Christensen 2010; 187.

the correct answer to the logic puzzle. The lack of evidence for my immunity overrides my evidence in favour of my answer to the puzzle, so regardless of whether I am susceptible to the drug or immune to it, and even if my initial answer to the logic puzzle was correct, I have reason to doubt the reliability of my initial judgment. Christensen and many others in the literature contend that the higher order evidence provided by being told about being slipped the drugs compels me to bracket my initial evidence and reasoning. Consequently, I need additional and independent evidence to permissibly hold my original attitude against the higher order defeat.<sup>143</sup>

#### 4.5: Peer Disagreement as Higher Order Evidence

The higher order evidence and peer disagreement discourses share a question of response, with both asking what agents are required or expected to do by way of altering or updating their attitudes when presented with information that is or appears to be higher order defeating evidence.<sup>144</sup> In the Drugs case, I face a tension between two states: I can either follow my initial reasoning which, admittedly, was sound, but involves a dogmatic and unsupported assumption that I am immune to the reasoning-inhibiting drug. Alternatively, I can reduce my confidence and go against my original reasoning, instead having a judgment that is (for all I know) less rational and justified. However, many authors do not just regard higher order defeat and peer disagreement as having comparable impact, they contend that learning about peer disagreement is itself higher order evidence, and therefore defeats each peer's beliefs in just the same way.

Suppose instead of being told that I'm under the influence of a drug that inhibits my reasoning capacities, I'm in conversation with a good friend who is an epistemic peer. It turns out that we have drawn different conclusions about some matter (e.g., the Restaurant Check example) and we are trying to make sense of this disagreement. We both have a history of making roughly equally accurate and considered calculations, and neither of us have any strong reason to think that either one of us is any more or less inhibited. Plausibly, I now have some reason to think that I did not form or acquire my belief about a matter reliably, as in the earlier examples, and I should now reconsider my reasoning or change my attitude, or at least think I have a reason to do so. However, the same can be said for my peer, so neither of us has any reason to prioritise or privilege our own reasoning and judgment, or to think that the other *must* be wrong. After all, we both had the same bill as shared evidence. Those who regard it as irrational or problematic for me to maintain my belief in the drugs example, will say that it is also irrational or problematic to think my belief is more likely to be correct in cases of peer

---

<sup>143</sup> Ibid; 194-6.

<sup>144</sup> Again, similar questions may apply for higher order confirming evidence, but much of the focus is on higher order defeat.

disagreement. The upshot is the same, to permissibly retain my belief, I require evidence or reasons that I did not initially factor in or acknowledge; until I acquire this, I must bracket my original evidence and revise my belief or suspend it.

Thomas Kelly grants that our first order evidence counts towards determining what is reasonable for epistemic peers to believe once they recognise or acknowledge that they disagree. However, he argues that the Equal Weight View fails to capture the weight of higher order evidence in such situations. He contends that the reasonable response to peer disagreement is instead determined by the combination of the initial propositional evidence *and* the facts that each peer believes as they do, and calls this the Total Evidence View:

Initially, you and I have access to the same substantial body of evidence E, evidence that in fact strongly favours H over not-H; you respond reasonably and so are quite confident that H is true; I on the other hand respond unreasonably and am equally confident that H is false. Once we compare notes, our new total evidence consists of E\*:

- (1) Our original evidence E
- (2) The fact that you are quite confident that H is true
- (3) The fact that I am quite confident that H is false

Given that you and I are peers, it is plausible to suppose that the two pieces of higher order psychological evidence ((2) and (3)) are more or less equally strong pieces of evidence that point in opposite directions. All else being equal, then, one would expect E\* to favour H over not-H inasmuch as it is composed of a substantial body of evidence that strongly favours H over not-H, supplemented by two additional pieces of evidence of approximately equal strength, one that tends to confirm H, another that tends to disconfirm H.<sup>145</sup>

Kelly is careful to avoid saying that one kind of evidence matters more than the other and contends that whether this is the case is not a good question. He discusses something in the spirit of this view in his earlier writing on peer disagreement, but only discusses it as a named position in his later writings. In his earlier writing, Kelly argues that there is no reason to think the combined body of evidence mandates suspending judgment or agnosticism about H. If the original first order evidence supported believing H, then adding 2 and 3 and granting them equal weight will not affect this outcome at all: “Our original evidence E does not simply vanish or

---

<sup>145</sup> Kelly, T. 2010; 202.

become irrelevant once we learn what the other person believes on the basis of that evidence.”<sup>146</sup> More recently, Kelly has argued that there is something unusual about counting higher order evidence as additional first order evidence for an attitude, but that it does count as evidence about the nature of one’s attitude. Otherwise, it seems to be a prevailing assumption that the fact that epistemic peers disagree is or generates higher order evidence for each speaker.<sup>147</sup>

In this introductory chapter, I have offered a general summary of three important and interesting debates in epistemology. These are the debates concerning the uniqueness of evidence, the import of higher order evidence, and the peer disagreement discourse. I contend that the first two are foundational in the third, but also that many authors in the peer disagreement literature rely excessively on these other two interesting and important debates, essentially crowding them out. Most claims about how to respond to peer disagreement assume i) that evidence is unique to some degree, ii) that facts about peer disagreement generate higher order defeating or undermining evidence, or both. In the next chapter, I begin to critique the wider peer disagreement discourse, beginning by tackling both the uniqueness and higher order evidence debates, as well as their place in the peer disagreement discourse. I regard both extremes of the uniqueness debate to be unconvincing, and I am unconvinced that peer disagreement is genuine higher order evidence, even if it appears to have comparable import or bearing.

---

<sup>146</sup> Kelly 2005; 190.

<sup>147</sup> There is also Adam Elga’s suggested alternative, the Right Reasons View, which “says that in the face of disagreement, one should adopt whatever view one’s original evidence in fact supports... evidence that one had before finding out about anyone else’s conclusions.” This view essentially bypasses the question of whether peer disagreement constitutes higher order evidence and treats the fact that there is peer disagreement as inconsequential. I address this view in Chapter 6.



## Chapter 5: Critiquing the Three Debates

In the previous chapter, I introduced three debates within epistemology that stand apart but also overlap, such that two of the three are typically subsumed by the third. The two debates that I regard as subsumed and crowded out concern the nature of evidence and its bearing on our beliefs and attitudes. The latter debate concerns how to respond when disagreement arises between two speakers who are epistemic peers with roughly similar capacities, dispositions, and evidence. In discussing these three debates in the previous chapter, I made the following claims:

1. The peer disagreement discourse purports to be about identifying the epistemic significance of disagreement, but much of the discourse turns out to be about the rational or rationally required way to respond to disagreement in rather specified circumstances.
2. More specifically, the question is chiefly: *what is it rational for speaker A (and B) to do about their different attitudes regarding p, given that A and B are (or regard themselves) epistemic peers about p?*
3. The uniqueness of evidence is a key motivator for what is arguably the main view in the peer disagreement debate: the Equal Weight View. If evidence is not strictly or singularly unique then there is far less reason to conciliate or worry about peer disagreement.
4. Likewise, the question of responding to peer disagreement bears great similarity to one concerning responses to higher order defeating evidence. Opinions are divided about whether the fact that there is peer disagreement constitutes higher order evidence itself.

I will raise and discuss several concerns I have with the peer disagreement discourse over the chapters that follow this one, beyond what I regard as excessive reliance on, or assumption of claims about uniqueness and higher order evidence. In this chapter, I elaborate on my concern about this reliance or dependence. I argue that we have little reason to think that evidence is strictly unique, but neither do we have any reason to endorse permissive views about evidence. I believe the answer instead lies in between, granting some flexibility but not a free for all concerning rational attitudes. I also argue that peer disagreement is not defeating higher order evidence. Against the concern that this renders peer disagreement benign, I argue that it does not need to generate higher order defeat to matter. Nonetheless, if peer disagreement is a form of higher order defeating evidence, it neutralizes and leaves disagreeing peers in an identical situation to that which precedes their disagreement.

### 5.1: Uniqueness and Permissiveness Revisited

In the previous chapter, I presented the uniqueness debate as chiefly being about whether any given piece or body of evidence requires a particular attitude or if multiple attitudes are

rationally permitted. However, I also noted that one of the main views in the peer disagreement debate (Equal Weight) presumes that a stricter form of uniqueness is true, or at least is heavily reliant on it. Several readings of uniqueness and permissiveness are available that vary according to how one replies to various questions about the scope and focus of evidence and the agents and times involved, which only add complexity and depth to the uniqueness debate. Uniqueness bears on particular aspects of rationality. It is a truism that with tighter restrictions on agent rationality comes ease of identifying when one fails to act rationally, because there are clear (or clearer) facts of the matter. Why, then, should we think that Uniqueness, or at least a stronger version of it, is true? Extreme uniqueness makes rationality too stringent, even inaccessible and overdemanding. Extreme permissiveness is simply implausible and, more importantly, it makes rationality *trivial*. I contend the answer is to endorse a moderate or liberal version of Uniqueness that maintains some plurality of rational attitudes and avoids either a single rational attitude or a free for all.

I begin by considering some of the motives for and against stronger, more stringent readings of Uniqueness. For Roger White, denying Uniqueness brings about greater costs than accepting it, since permissiveness creates a tension between rationality and arbitrariness. There are no circumstances in which it is possible to be fully rational while believing both *p* and not-*p* from the same evidence. He contends that if we grant that more than one attitude is permissible given some evidence, it degrades the value of belief-forming and reasoning so that it is no different from taking a pill that induces an arbitrary, artificial belief.<sup>148</sup> If there is simultaneous evidence for and against a proposition, the rational response is to suspend judgment about the matter. Even if evidence is shared or common, our judgments will be affected by more than just what evidence we have at the time because our background beliefs and experiences will also bear on it. Again, the interest or motivation for the Equal Weight View comes from endorsing or presuming a strict reading of uniqueness, and the uniqueness debate seems to be *chiefly* about how to attribute rationality or justification to speakers and their attitudes.

Jon Matheson notes that extreme permissiveness – the claim that there are no rational limits on a body of evidence and that any given evidence can justify belief, disbelief, and suspension of belief – clearly supports remaining steadfast in response to peer disagreement. Indeed, any permissive view or non-stringent reading of uniqueness can be tied to steadfastness. However, Matheson argues that extreme permissiveness is utterly implausible, and that there must be some rational constraints on belief states. He notes that a moderate permissiveness, one

---

<sup>148</sup> See White 2005; 452.

which denies the extremes of both uniqueness and permissiveness is plausible but that it does not motivate steadfastness.<sup>149</sup>

Bryan Frances initially contends that the most extreme forms of permissiveness permit any attitude to be rational, downgrading the reliability of evidence as a guide to truth, which Frances regards as essential to the nature of evidence.<sup>150</sup> However, he proceeds to raise various issues for several variants of Uniqueness and tentatively concludes that Uniqueness cannot be true. He notes that the truth of a belief or proposition can come apart from whether available evidence makes belief in it reasonable or justified. Moreover, he makes a similar argument to Miriam Schoenfield by acknowledging that speakers can have different epistemic standards (or *starting points*, in his words). A speaker's starting point determines whether they are epistemically risk-averse or adventurous, and thus what methods for believing they regard as fitting or apt. If we assume that two or more speakers have different starting points but that they are both reasonable, it follows that their answers will likewise be both different and reasonable. Importantly, this assumption marks the difference between what each author is aiming for with their respective arguments. Schoenfield aims to use differences in starting points to vindicate permissiveness, whereas Frances merely wants to make sense of how there can be rational disagreement about a matter between two speakers who are epistemic peers. He is adamant that "it is clearly the case that not just *any* starting point will be reasonable."<sup>151</sup>

## 5.2: Non-Extreme Uniqueness and Permissiveness

The most stringent form of uniqueness maintains that there is an attitude that is either maximally justified or uniquely rationally required (albeit not necessarily discernible) and ought to be held given some relevant evidence set. But there are forms of uniqueness available which are less stringent and more flexible or liberal. Roger White raises one possibility, and David Enoch proposes another:

**Uniqueness:** If an agent whose total evidence is E is fully rational in taking doxastic attitude D to P, then necessarily, any subject with total evidence E who takes a different attitude to P is less than fully rational.<sup>152</sup>

---

<sup>149</sup> Matheson, J. 2015; 60. "*moderate permissivism* is the claim that... while a body of evidence does place some rational limits as to what you can justifiably believe on its basis, it is possible that you nonetheless have some options."

<sup>150</sup> See Frances, B. 2014; 179.

<sup>151</sup> Ibid; 184-5. See also Schoenfield, M. 2012; 199.

<sup>152</sup> White, R. 2014; 312.

The Uniqueness Thesis is the claim that, with given evidence, there is a unique degree of belief that is maximally rational. So in order to deny this claim, one has to assert that there is no one degree of belief that is uniquely maximally rational, how about the view according to which there *is* one degree of belief that is *maximally* rational, but some other degrees of belief – though less than maximally rational – are still rationally permissible?<sup>153</sup>

I do not find strict or stringent forms of Uniqueness compelling. The most extreme reading posits that a single, rationally required attitude is discernible from one's evidence, which restricts rationality to a kind of circumstantial epistemic luck. Either one has the necessary resources or evidence and can reason from that evidence to the rational attitude, or one cannot, and thus cannot be rational. The demands or constraints on rationality are therefore unattainable and overdemanding. Even if one believes they have things right, evidence can be misleading or inconclusive and agents can misinterpret it. More liberal readings of uniqueness such as Enoch's grant a degree of flexibility, there are higher and lower boundaries that mean some attitudes can be rational, others more rational, and others still most rational. More importantly, it allows for rational disagreement between agents which can be explained using non-evidential reasons.

Moreover, there are certain subjects or domains which intuitively admit of pluralism concerning values, preferences, and beliefs. Extreme uniqueness stands in direct contrast to this. Perhaps permissiveness will be more compelling if it is interpreted in a way that is less radical or extreme, since we can grant a relation between evidence and varying attitudes without maintaining that all attitudes are (equally) rational. White himself discusses the following possibility:

**Non-Extreme Permissivism:** there are possible cases in which believing P and believing not-P are both *well supported* by one's evidence and hence likely to be true.<sup>154</sup>

Recall that the most extreme reading of permissiveness grants that any attitude can be rationally permitted on the basis of any relevant evidence. White's non-extreme reading does not treat all attitudes as equally rational; they are all merely in a position to be "well supported". Although I have labelled it as a form of permissiveness, White does not because he is unsure that it *is* a form of permissiveness, even though he thinks it can face the same issues as its more radical counterparts.

---

<sup>153</sup> Enoch, D. 2010; footnote 9.

<sup>154</sup> White 2005; 451. My emphasis.

I contend that the most compelling answer to the uniqueness debate is either a less radical form of permissiveness or a less stringent form of uniqueness. Approaching rationality in either of these ways is more realistic because rationality is not out of reach, nor is it contingent on possessing the relevant evidence and knowledge of how it points. Moreover, there are various domains of thought which have multiple values and disagreement over how to prioritise them which are fundamentally insoluble. Isaiah Berlin discusses this quite prominently, positing that many moral and political values are incommensurable, that it is impossible to rank them universally or objectively, and that judging other cultures by the standards of one's own demonstrates a lack of understanding.<sup>155</sup> Elsewhere he notes that cultural differences cannot ever be set aside or overruled, and that it would be suffocating to force this:

But if we are to have as many types of perfection as there are types of culture, each with its ideal constellation of virtues, then the very notion of a single perfect society is logically incoherent. This, I think, is the beginning of the modern attack on the notion of Utopia, Utopia as such.<sup>156</sup>

Berlin's point is that certain value-laden domains, such ethics and politics, perhaps even aesthetics or epistemology, with different goods and values, are genuinely insoluble and that there cannot be a true ranking or supremacy of values in such domains. Many robust realists in such domains will disagree, in which case they may endorse or at least implicitly accept a more stringent reading of uniqueness. If, however, one agrees with Berlin that certain domains do not admit of an ultimate value or good, then they may want to say instead that multiple claims are permissible and justified. Regardless of whether one believes peer disagreements are genuinely possible in such domains, then, a value pluralist or non-realist may endorse a more permissive view of evidence, and therefore be unfazed by peer disagreements about moral, political, or aesthetic matters.

Disagreements over matters of taste or preference are fitting examples: we can grant that there are reasons that relate to (for example) the quality of each Star Wars trilogy, but not that any of that evidence will indicate that one of the trilogies is objectively or factually superior. Speakers will differ over which trilogy they prefer or rate as best and, in the grand scheme of things, there doesn't seem to be anything problematic or worrisome about such differences. While many might disagree over the matter of which trilogy is best (or even which trilogy is fit for consideration), at least some participants will think it absurd to say that all the evidence only

---

<sup>155</sup> See Berlin, I. [1978]

<sup>156</sup> Ibid; 42

ever points one way. Likewise, they will likely accept the differences in preference without claiming that those who differ are not rational. Aesthetic disputes and disagreements in taste have attracted plenty of discussion, but it is often framed in terms of whether 1) there is genuine disagreement between speakers or 2) whether either speaker has made any error or fault by holding and expressing their respective attitude.<sup>157</sup> In other words, uniqueness bears on disagreements and disagreement discourse beyond peer disagreements, but there has been little to no consideration of it.

The strictest forms of uniqueness do not hold for various kinds of disputes. Even if we think there is some degree of objectivity in those domains, it is not always clear which attitudes are correct or rational, and which signal an error in reasoning, so uniqueness bears on them but not perhaps in the way it is presented in the uniqueness or peer disagreement discourses. Those domains that have unclear objectivity lend themselves at least to a liberal form of uniqueness or a less extreme version of permissiveness, and it is an open question which of them is more fitting. However, it is also still an open matter whether a stringent reading of uniqueness is more fitting than liberal readings for domains with much clearer objectivity; that is, for matters of fact. I have argued that stricter readings of uniqueness generate excessive and unattainable standards for rationality. Along similar lines, Matthew Kopec has argued that a plausible counterexample to Uniqueness can be discerned if we suppose that one can know before forming an attitude that that very attitude will be accurate or true regardless of what it is (either belief or disbelief):

If it is correct that our subject can rationally decide to believe P, but would be just as rational if she instead disbelieved P, then we have a counterexample to Uniqueness. Either way she goes, she has the same evidence before she chooses that path. Since either choice is rationally justified, it is rational to believe P given her evidence E, while it is equally rational for her to disbelieve P given E. Since there is more than one doxastic attitude that it is rational for our agent to hold toward P given E, Uniqueness is false.<sup>158</sup>

Kopec assumes that one knows in advance that forming a belief in P is i) highly reliable, ii) just as reliable as any other method of attaining belief, and iii) as reliable as disbelieving P. Against White's concerns, this does not make belief-forming arbitrary, but there are no substantial differences between belief and disbelief. The only attitude we lack knowledge about the reliability of is belief suspension, but Kopec contends that it is implausible to think that the circumstances in his example require the subject to suspend belief. Epistemic rationality would

---

<sup>157</sup> See for example, Timothy Sundell's 2011.

<sup>158</sup> Kopec, M. 2015; 406-7.

be very strange if it required agents to suspend belief in circumstances where they could form beliefs that are highly likely to be true. I agree with this, although I do not think there are no circumstances at all in which we cannot know suspending belief is the most rational (or rationally required) attitude. Plausibly, if the evidence is inconclusive and there is an even balance between reasons to believe and reasons not to believe that P, then the answer is that we ought to suspend judgment.

None of this is to say that Extreme Permissiveness is off the hook though. Roger White thinks that it degrades belief-forming and makes it arbitrary, which is certainly a worrying possibility. For example, I might not think that it is especially likely to rain but I may have inconclusive evidence, which leads me to regard taking an umbrella as a sensible, and not a particularly troublesome precaution. The stakes and costs of taking an umbrella are relatively low, and it does not seem to be wrong or akratic to believe it is not likely to rain but act as though it is. In other words, there seems to be leeway for more than one rational attitude. But the safeguard I believe is present in the umbrella case does not occur in every case, and certainly not in those that have higher stakes. In some cases, we are undoubtedly far less justified in having inconsistency between our beliefs and actions, or for believing there is a safeguard when we make our decisions.

Juror cases are a clear example; we would expect any diligent and competent juror to think carefully about their vote and base it on all the evidence they have available. If we were to learn that a juror made a judgment of guilt or innocence based on an indifferent or “just in case” attitude, we would most likely be horrified (and rightly so), especially if they turned out to be mistaken. The most common standard in English criminal law is proof of guilt ‘beyond a reasonable doubt’, meaning the evidence needs to be substantial to support a conviction. This is more in keeping with the strictest reading of uniqueness, whereas anything even remotely permissive or pluralist would be more akin to a ‘balance of probabilities’ standard, in which one outcome need only be demonstrated as most likely or more than 51% likely. At any rate, Extreme Permissiveness has little justification. Even if one doesn’t accept White’s overall argument in favour of Uniqueness, his worries about arbitrariness still have *some* weight and relevance. Radical versions of permissiveness reduce rationality to a free for all and make it too easy, in contrast to stringent or extreme uniqueness which makes rationality overdemanding and nearly unattainable. If our evidence is inconclusive then we should suspend judgment or try to discern whether there is any other evidence that has been overlooked.

### 5.3: Peer Disagreement and Non-Extreme Uniqueness



I have considered some of the arguments for and against uniqueness in the discourse, and I have argued that extreme versions of both positions are unconvincing and have worrying implications. Extreme uniqueness places excessive demands on rationality by requiring evidential foreknowledge, whereas extreme permissiveness makes rationality trivial by granting a free for all for rationally permissible beliefs and actions. I have noted that some discussions and domains imply a degree of plurality or permissiveness more often, I now want to address the reliance of peer disagreement, and particularly the Equal Weight View, on evidential Uniqueness. If the demands of Extreme Uniqueness are as excessive as I believe, then this means that any view about peer disagreement which assumes or endorses it will also be overdemanding. It limits rationality to single attitudes and rules out rationality in circumstances where evidence or agents' capacities for reasoning are limited. If Extreme Uniqueness holds, and only one attitude is justified from a given piece or body of evidence, then rationality would require us to 1) have the relevant evidence, 2) know what that attitude is, and 3) be disposed to holding that attitude. In other words, it requires a healthy amount of circumstantial evidential luck, when evidence and reasoning abilities are arguably scarce, or at least unequally distributed. Even in cases where agents seem to have roughly similar and well-developed capacities and evidence, there may still be difference-making factors that render differences in judgment rational. However, as far as peer disagreement goes, it would be highly unlikely for two agents to have these traits simultaneously, let alone to also reach different attitudes based on the same evidence with no clear error on either side.

For example, in my Foul Shot example, Helen and Summer are playing a game of pool and come to disagree over whether Summer's most recent shot was a foul. Neither player is irrational in having different beliefs over the matter. Helen has perceptual evidence confirming that it was a foul, but Summer lacks this evidence so, from her perspective and based on all the evidence she has, her shot was not a foul. Summer may not be as rational as Helen, *all things considered*, but that need not mean that she is irrational. Pool tables are often found in crowded, busy pubs or bars in which it is far from easy to maintain perfect focus or always have clear view of the game. There may be limited relevant evidence available, but that does not guarantee that everyone will always successfully acquire it, and we can only be as rational as our evidence allows us to be.<sup>159</sup> Few agents will ever be ideally rational in the sense that Extreme Uniqueness seems to demand, and most of us will not even come close since we are fallible beings. Extreme

---

<sup>159</sup> This draws out an interesting question about whether the status of peerage requires that speakers believe they each have the same evidence, regardless of whether they in fact do. I also believe the foul shot example to be an interesting case where the speakers are very close to being peers, but because of the way the example is set up, they *just* fall short.

Uniqueness is not merely implausible, but unrealistic and overdemanding, and any argument that assumes it can be regarded as relying on an idealised notion of human reasoning, including Equal Weight.

Liberal or moderate versions of uniqueness and permissiveness are less stringent but still ensure that there are some boundaries or constraints on rationality. Importantly, however, liberal versions of uniqueness and permissiveness need not be one and the same; there is still a qualified difference between them. Even using the descriptors ‘liberal’, ‘moderate’, or ‘less extreme/radical’ makes it difficult to pinpoint specific views, so I will compare the version of Uniqueness I’ve already mentioned from David Enoch with the non-extreme *possible* permissiveness Roger White suggests. As a reminder, these are, respectively:

Non-Extreme Uniqueness: There are possible cases in which more than one doxastic attitude is rationally permissible, given some evidence, but there is one attitude which is maximally rational.<sup>160</sup>

Non-Extreme Permissivism: there are possible cases in which believing P and believing not-P are both well supported by one’s evidence and hence likely to be true.

Importantly, these are two very similar claims. Non-Extreme Permissivism says that there are possible cases in which one could have good reason (i.e., evidence) to believe P while also having good reason not to believe P. It does not specify that there must be some dominating or overriding evidence. Non-Extreme Uniqueness contends that there will be at least some cases where this is the case, so it makes just the same claim as Non-Extreme Permissivism, while also claiming that there are some possible cases where one will have stronger reasons to hold one attitude over another. After all, if more than one attitude is rationally permissible, given some evidence, then it goes without saying that all the relevant attitudes are well supported by that evidence. What Non-Extreme Permissivism is *not* doing is blanket endorsing belief suspension in some or all cases, because i) that would imply one’s evidence directs one to suspending belief as the only appropriate attitude, which ii) would reduce the view to a stringent form of uniqueness.

I am careful to say all the *relevant* attitudes because granting that more than one attitude rationally follows from evidence does not mean rationality is a free for all, otherwise this is just Extreme Permissiveness. I am more persuaded by the view I refer to as Non-Extreme Uniqueness because it balances allowing more than one rational attitude with keeping rationality

---

<sup>160</sup> The label of Non-Extreme is my own, rather than Enoch’s, and is offered deliberately to make for an easy comparison with Non-Extreme Permissivism.

within the realm of possibility, while also not granting that just any attitude can be rational, given some evidence. It fits with cases that are plausibly permissive, such as value- or taste-based disagreements, and avoids specifying that some attitude or other must be *exclusively* rational for each matter and evidence. What, then, does this mean for peer disagreement? Granting that some but not all attitudes are rational to some degree weakens the pressure to conciliate or compromise with one's peer, but it need not rule out the Equal Weight View and its competitors. We might still say that we should grant that a disagreeing peer's attitude is equally weighty to our own, without conceding that we should each compromise. This brings us to the position offered by Adam Elga known as the Extra Weight View that "you should be pulled a *bit* in the peer's direction – but not halfway. You should still think it more likely that you are right than that the peer is."<sup>161</sup> Endorsing Non-Extreme Uniqueness does not push towards either steadfastness or conciliation, it is simply saying that sometimes we can allow for a plurality of rational attitudes without any problem. Some attitudes can still be less rational or rationally erroneous, meaning some rational compulsion to reconsider or change an attitude is still required if one acts or believes irrationally. Disagreement (merely) suggests that the evidence appears to point more than one way, there may be some cases where reinspecting or reconsidering fails to get disagreeing speakers to the right answers.

However, endorsing non-extreme uniqueness does not guarantee that either party in a two-way disagreement has things right. We can see this with some examples from the literature. David Christensen offers an unrealistic version of the Restaurant Check case, in which the friend calculates everyone's share of the bill as being ridiculously high, higher than the bill itself, in fact. In cases such as this, it is still entirely possible that both my own calculation and my friend's calculation are mistaken. Christensen thinks that cases like these do not generate any pressure to conciliate, and I agree that this is so, since it would be absurd to think that one diner's share of (e.g.,) a \$225 bill is \$450!<sup>162</sup> Secondly, Thomas Kelly offers some examples (in terms of credences) with the aim of presenting a case for rejecting the Equal Weight View:

*Case 4.* Despite having access to the same substantial body of evidence E, you and I arrive at very different opinions about some hypothesis H: while I am quite confident that H is true, you are quite confident that it is false. Indeed, at time t<sub>0</sub>, immediately

---

<sup>161</sup> Elga 2007; 485. See also Thomas Kelly, who phrases the view as stating that: "each party to the dispute is permitted to give some special, presumptive weight to his or her own judgment." 2010; 161.

<sup>162</sup> See Christensen, D. 2007; 199.

before encountering one another, my credence for H stands at .8 while your credence stands at .2. At time  $t_1$ , you and I meet and compare notes.<sup>163</sup>

Kelly is sceptical that Equal Weight follows from cases such as this because nothing is specified about the relationship between E and H, and it's implausible how confident anyone should be in that situation. He then continues the case to show that Equal Weight sometimes fail to deliver a rationally optimal or compelling outcome:

*Case 4, continued.* In fact, hypothesis H is quite unlikely on evidence E. Your giving credence .2 to H is the reasonable response to that evidence. Moreover, you respond in this way precisely because you recognise that H is quite unlikely on E. On the other hand, my giving credence .8 to H is an unreasonable response and reflects the fact that I have significantly overestimated the probative force of E with respect to H.<sup>164</sup>

Christensen's extreme Restaurant Check case shows a clear error on one side and that we need not always think there is pressure to conciliate, regardless of whether we endorse Equal Weight or not. Kelly's case shows not only that Equal Weight can deliver unintuitive and undesirable responses, but also that it can force an erroneous response by compelling an otherwise rational agent to change their attitude for a rationally sub-optimal alternative. If Equal Weight holds then both peers ought to split the difference and adopt a .5 credence, but this does not follow from E. Kelly contends that Equal Weight and Uniqueness (in its strictest form) are tied together, and that Uniqueness is only plausible if we focus on a coarse-grained doxastic attitudes (i.e., whole belief states, not credences.) Nonetheless, he rejects Equal Weight because it can generate undesirable consequences. I argue these worries can be avoided by endorsing Non-Extreme Uniqueness instead. Equal Weight need not follow from Non-Extreme Uniqueness, and Non-Extreme Uniqueness is neither too permissive or too demanding, so it avoids the issue I have raised for both extremes of uniqueness and permissiveness.

#### 5.4: Higher Order Evidence, Again.

In the previous chapter, I introduced the popular distinction between first order propositional evidence (which is any reason to hold an attitude or belief) and higher order evidence about how reliably one acquires evidence or forms an attitude. Within the relevant discourse, first order

---

<sup>163</sup> Kelly, T. 2010; 190.

<sup>164</sup> Ibid; 191.

evidence relates to the truth of a belief or the fittingness of an attitude, whereas higher order evidence is about its justification. The discourse has the following two questions as its basis:

1. Does higher order evidence exist, either independently from or as a proper subset of first order evidence?
2. If higher order evidence *does* exist in some capacity, what is its bearing on our first- and higher order attitudes and judgments, and on our justification for those attitudes?

However, the chief focus is on higher order defeat and our response to it, i.e., the following:

3. If higher order evidence has any bearing on our lower-level attitudes (either defeating or confirming), should it cause or compel us to revise our attitudes? If so, then how?
4. In light of the possibility of higher order defeat, can or should we embrace or reject the prospect of distinctively epistemic *akrasia*?

Generally, the consensus is that if there is higher order defeating evidence, it is stubborn and irrational to retain one's belief against it, because the defeat indicates that one's belief is flawed or mistaken, just not in the straightforward way that one can have propositional defeat from contrary evidence. The higher order evidence and peer disagreement discourses overlap partly because both have questions of response (as in question 3 above), but also because some authors present peer disagreement as an example of higher order evidence or as signalling higher order defeat.

Typically, the argument is something along the lines of the following: once we recognise that some cognitive malfunction has (or is likely to have) occurred, we should no longer think our belief (or whichever state we initially had) is justified, even if it still seems obviously true or apt. This generally assumes some form of evidentialism – the claim that we should only believe as far as our evidence directs us, and that our doxastic states should conform to our available evidence. Higher order evidence is regarded as being of the relevant kind of evidence, so it is not unique in that regard; however, retaining a doxastic state that has been undermined by evidence that signals some flaw or irrationality is clearly a failure to conform one's attitude to one's evidence.<sup>165</sup> Some discuss a further distinction between apparent and actual higher order defeat, considering what to do when evidence *misleadingly* indicates that a belief or attitude is unjustified. This signals another similarity with peer disagreement because there is at least some

---

<sup>165</sup> See Lasonen-Aarnio 2014; 318.

discussion concerning the comparative import of actual disagreement and merely possible disagreement.

Maria Lasonen-Aarnio argues that we should reject the idea of higher order defeat because endorsing it often means acknowledging that we may be compelled to act in ways that are epistemically irrational and akratic. She presumes a two-tiered theory of epistemic justification and a rule-oriented picture of epistemic rationality. On her account, rational action is achieved simply by applying and following correct epistemic rules; she contends that no possible system of rules can make sense of higher order defeat or explain it away without also admitting epistemic akrasia. In the previous chapter, I discussed a typical example from the literature that involves an agent being administered a drug that potentially hinders one's capacities for reasoning. Lasonen-Aarnio instead considers examples wherein agents are susceptible to some brain-altering condition such as hypoxia (i.e., oxygen deficiency.) Both the drugs case and the hypoxia case create a tension between following an initial rule (i.e., believing with justification and one's evidence that  $p$ ) and acknowledging the implications of learning about the higher order defeating evidence. Since one eventually acquires reason to believe that their previous doxastic state was not adequately attained or formed, they should suspend judgment that  $p$ . It would be irrational to ignore that evidence, but adhering to it requires one to violate the epistemic rule that enabled them to justifiably believe that  $p$ , so one acts irrationally either way. One cannot grant that belief and suspension are both rational because the rules that compel them are issuing incompatible recommendations.<sup>166</sup>

Lasonen-Aarnio considers various solutions, all involving distinguishing a particular kind or set of rules. The first is an overarching über-rule which purports to get around defeat and incompatible recommendations by issuing a single correct doxastic response, which she doubts can be sufficiently expressed in such a way that it offers genuine guidance.<sup>167</sup> The other two solutions she considers are, respectively, positing *hierarchies* of rules and *meta-rules* that govern our ordinary rules; both run the risk of issuing conflicting requirements and are susceptible to higher order defeat. The implication is that no failsafe against higher order defeat can be built into systems of epistemic rules, and no prospect of indefeasible rules can be reconciled with higher order defeat. The conclusion Lasonen-Aarnio reaches is that, while one cannot plausibly deny the possibility of higher order defeat, neither can one embrace it and avoid uniquely epistemically akratic action, i.e., behaving in ways one knows to be irrational:

---

<sup>166</sup> Ibid; 329.

<sup>167</sup> Ibid; 332-3.

It may come as a surprise that in some cases a state can be perfectly epistemically rational even if one has what would seem like strong evidence for thinking that it is not. In particular, in so far as there is such a thing as a correct inductive policy or epistemic system, it can be rational to follow the recommendations of that policy or system even if one possesses evidence that in doing so one has committed a rational error.<sup>168</sup>

Recommending that one believe that a rule is flawed is not tantamount to recommending that one stop following the rule. That one should believe that one shouldn't  $\varphi$  doesn't entail that one shouldn't  $\varphi$ . Thinking about defeat easily leads to over-inflating the notion of epistemic rationality. There is an intuition that subjects ought to revise their doxastic states in the presence of evidence that those states are flawed, and that they are criticisable for not doing so. I think these intuitions are correct...

But I very much doubt whether there is a non-paradoxical notion of epistemic rationality that marches step in step with such criticisability. There are epistemic oughts that a subject can violate without thereby being epistemically irrational, or failing to meet the criteria for justification.<sup>169</sup>

I am inclined to agree with this last claim, I am not convinced that epistemic akrasia is altogether unavoidable, but neither do I think that all failures to comply with epistemic obligations constitute epistemically irrational behaviour. Inevitably, some norms, rules, or principles will issue requirements we regard as overdemanding, or which incur a heavy non-epistemic cost, but likewise others will come relatively close to avoiding akratic outcomes. More in the higher order evidence discourse tend to disagree with Lasonen-Aarnio and maintain that, for the sake of rationality, agents *ought* to change their beliefs, states, or attitudes when conflicts between epistemic norms or principles arise. However, I believe tensions between adhering to epistemic rules and avoiding epistemic irrationality are common enough that this should not be assumed as a default.

Darren Bradley disputes the inevitability of a tension or paradox of rationality. He contends that an über-rule need not guide us directly, and that mere parts of it could still achieve their function. There are typically fewer considerations to make sense of with the likes of ethics and aesthetics than there are with epistemology, but we still seem to make sense of moral rules despite there being fewer factors to consider. Likewise, there are many non-epistemic rules that we follow, intentionally or otherwise, that we cannot articulate or grasp in any great depth, such

---

<sup>168</sup> Ibid; 342.

<sup>169</sup> Ibid; 343.

as linguistic rules. We might even get things wrong by trying to formulate such rules, so we need not worry that a complicated epistemic rule cannot be formulated because it is made up of many components. We might be better off positing some kind of good epistemic sense or intuition which enables us to be guided by an über-rule.<sup>170</sup>

While the relation between higher order evidence and akrasia is often discussed, a great deal more attention is directed towards the question of responding to higher order evidential defeat.<sup>171</sup> Subsequently, some use the same distinction between remaining steadfast or conciliating our first order attitudes in response to higher order defeating evidence. Sophie Horowitz discusses a third option which acknowledges the importance of both first- and higher order evidence as well as the fact that each directs different attitudes. On this “Level-Splitting” view, one is permitted to maintain their attitude while recognising its poor evidential support, that is, beliefs along the lines of “p, but my evidence doesn’t support p” are permissible. This view is well-positioned to acknowledge and embrace the import of both propositional and higher order evidence, but it also embraces epistemic akrasia as a norm or even as a rational requirement.<sup>172</sup>

Horowitz considers the example of a police detective, Sam, investigating a theft and drawing a poorly supported judgment that one of various suspects, Lucy, is responsible. His conclusion is poorly received by his colleague, who observes that Sam has a track record of forming similar judgments after lengthy considerations and little rest:

First, suppose that Sam, having reacted as Level-Splitting recommends, reflects on his situation. He is highly confident that Lucy is the thief. And he has some first order evidence, E, which he is highly confident does *not* support the proposition that Lucy is the thief... If Sam takes both “Lucy is the thief” and “my evidence doesn’t support Lucy” as premises, it seems that he can engage in some patently bad reasoning.<sup>173</sup>

Why should we regard this as reasonable? The main motive for Level-Splitting is the differentiation between first order evidence relating to our first order attitudes, and higher order evidence only (directly) being about our higher order attitudes. While there is no actual contradiction from having recommendations that diverge, it would be strange or absurd to

---

<sup>170</sup> See Bradley, D. 2019; 13-15.

<sup>171</sup> As we will see, this is where the parallels with peer disagreement arise. Unlike the uniqueness debate largely influencing or informing positions in the peer disagreement discourse, the higher order evidence debate is indirect, the two run more in parallel.

<sup>172</sup> See Horowitz, S. 2014; 720.

<sup>173</sup> Ibid; 725.



believe or assert that “Lucy is the thief, but my evidence does not support that.” Those who endorse Level-Splitting will not be persuaded of the argument’s intuitive oddness, but Horowitz contends that endorsing it has further-reaching counterintuitive consequences.<sup>174</sup>

The first of these is that it leads to fundamentally bad reasoning, since we usually believe what we believe because we think our evidence supports it, whereas it is just the opposite in the case above. If Sam instead received the higher order evidence before assessing any evidence or making a judgment but still believed as Level-Splitting compels him, he would likely impermissibly and irrationally believe that his evidence is misleading. This is not to say that it is always irrational to believe one’s evidence is misleading, but unless one gains newer, more conclusive evidence that proves this, the prospects of justifying this claim are slim, especially if one knows ahead of time that they will be in a bad position to rationalise evidence. Sam’s evidence, E, must *somehow* bear on the matter, P, of who is the thief, so it still warrants confidence in the guilt of some subject and doubt in the guilt of the others. He cannot conclude that his first order evidence is misleading, nor that his total evidence is inconclusive, and he has higher order evidence that undermines his reasoning and judgment. There are no grounds to suggest that Level-Splitting delivers the right outcome, which is illustrated by supposing Sam were to place a bet of Lucy’s guilt. Horowitz imagines that Sam would offer 1:9 odds on the evidence supporting Lucy’s guilt, but offers 9:1 odds on Lucy actually being guilty: “It seems patently irrational to treat a bet about P and a bet about whether one’s evidence supports P as completely separate. But if Level-Splitting is right, this is just how a rational person will behave in Sam’s situation.”<sup>175</sup>

Moreover, there are substantial and highly worrying epistemic and practical consequences of granting Level-Splitting: Sam’s inconsistent and poorly justified beliefs would lead to the recommendation that the police apprehend and interrogate Lucy, but Sam would be unable to explain or justify any of this beyond what amounts to a punt or a guess. I raised a worry earlier about extreme permissiveness leaving open the possibility of making high-stakes decisions and judgments with an attitude of indifference, and the same worries are present here. We would be horrified if a criminal judgment or conviction was delivered with no evidence or good reason, especially if it turned out to be wrongful and mistaken. As such, Level-Splitting carries a significant burden of explanation or justification for maintaining or permitting epistemic akrasia which it largely cannot meet.

Despite this, some defences have been offered. Timothy Williamson has claimed that a rational agent can come to know various claims and competently believe their conjunction

---

<sup>174</sup> Ibid; 725.

<sup>175</sup> Ibid; 728.

despite being fallible and prone to deductive and inferential errors. One can confidently believe the conjunction, C, while being highly confident that they do not *know* C. As Horowitz notes, however, this has the implication that one has made various errors that have somehow cancelled each other out and resulted in a lucky belief that falls short of knowledge.<sup>176</sup> Brian Weatherson has argued that akratic beliefs and actions are permissible when one has misleading evidence.<sup>177</sup> He draws an analogy with moral action and akrasia, imagining a student facing a murderer at the door and deliberating over whether to lie about her roommate's whereabouts. Because her tutor's pro-Kantian arguments are so compelling, Weatherson argues that she should lie to the murderer, but she should nonetheless *believe* that she is doing the wrong thing since that's what her evidence supports.<sup>178</sup> In other words, sometimes what we should do or believe can come apart from what we should believe we should do or believe. As we have seen, Horowitz contends that Level-Splitting is not able to demonstrate that evidence is always misleading, and when it is not the akrasia cannot be vindicated. It is certainly possible to think that one's evidence or some subset of it is misleading. However, if I have some reason to believe that my first order evidence was misleading or fabricated, or that I was somehow faulty or hindered in my reasoning, then I should straightforwardly be open to re-evaluating it. Similarly, if it then turns out that *that* evidence was misleading, miscommunicated, or misinterpreted, I can again reconsider my higher order belief.

I grant that there will be cases where we encounter internal asymmetries and dilemmas where we feel compelled to believe one way, but we have something else (a second party or our own memory or conscience) pushing us in a different direction. I agree with Lasonen-Aarnio that such dilemmas are clearly possible, but I also agree with Horowitz that we should not therefore deny the possibility of higher order defeat. In that respect, I am rather amenable to conciliation. I am less convinced that peer disagreement functions in a way that is equivalent to higher order evidence, or that facts about peer disagreement are *in fact* items of higher order evidence. Rather, I believe the fact that there is an acknowledged disagreement between two epistemic peers is at most first order evidence about the distribution of opinions regarding the subject of that disagreement. I will elaborate in the following section.

### 5.5: Peer Disagreement and/as Higher Order Evidence

---

<sup>176</sup> Ibid; 721. I am sceptical that there is anything inconsistent with confidently believing that P while maintaining that one does not *know* P, this is quite a common occurrence.

<sup>177</sup> Weatherson, B. (ms); 15.

<sup>178</sup> Ibid; 12.

Many regard peer disagreement as a form of higher order evidence, if not something with an equivalent function. Both share a question of response but more specifically, there is a question of whether and how much the fact that my epistemic peer and I disagree on a matter defeats our justification for our respective belief. Opinions are divided on whether the fact of disagreement *is* higher order evidence or is merely additional propositional evidence. The fact that two speakers are both generally reliable and well-informed agents who consider their evidence before making any judgments implies that their belief or attitude is true or fitting and that it is justified. When coupled with the fact that all evidence in a peer disagreement is shared or common, therein lies the puzzle. However, the fact that anyone is a well-informed, considerate believer merely indicates the likelihood that their beliefs are justified, it does not confirm that the relevant beliefs are justified or ensure that their beliefs are true. Thomas Kelly makes just this observation:

It does not follow from the fact that a generally reasonable individual believes H on the basis of E that it is reasonable to believe H on the basis of E. In a case in which E does not adequately support H but a generally reasonable individual mistakenly believes H on the basis of E, the fact that the individual believes as she does constitutes *misleading* evidence about the character of the evidence E. But misleading evidence is evidence nonetheless.<sup>179</sup>

From your perspective, the fact that you believe as you do is the *result* of your assessment of the probative force of the first order evidence: it is not one more piece of evidence to be placed alongside the rest... Similarly, when you first came to believe that H on the basis of your initial consideration of the first order evidence E, you did not then proceed to treat the fact that *I believe that H is true* as a reason to *increase* your confidence that H is true.<sup>180</sup>

Our respective assertions may count as testimonial evidence, they might be supplemented with further evidence that one of us has made an error or overlooked something, as in my Foul Shot example. Likewise, disagreement between non-peers will likely suggest an asymmetry between the speakers, perhaps one speaker has evidence that the other lacks. The fact that a speaker lacks certain evidence regarding a matter need not prevent them from appreciating its import, but I contend that each speaker's bare assertion in a peer disagreement does not and cannot count for

---

<sup>179</sup> Kelly, T. 2005; 186.

<sup>180</sup> Ibid; 187.

evidence about the matter itself. If epistemic peers disagree about some matter, claim, or hypothesis P, the fact that they disagree is first order evidence for a related matter (e.g., P\*) which concerns the distribution of opinions about the matter P, rather than bearing on P itself.

In his earlier writings, at least, Thomas Kelly argues that Equal Weight View and ascribing equal significance to each view in a peer disagreement is unavoidable, at least not without “some positive reason to think that one of us is more likely to do a better job with respect to assessing the relevant evidence than the other.”<sup>181</sup> I do not think acknowledging the significance of a speaker’s attitude, even if they are our epistemic peer, forces any requirement or obligation to compromise or suspend judgment as Equal Weight maintains. I argue that disagreement in general provides strong reason to reconsider and re-evaluate collectively and cooperatively, but not necessarily to revise attitudes altogether. Consider the following example from David Christensen:

Suppose that two people have shared first order evidence E which bears on proposition P. Alice forms the belief that P on the basis of E *before Ben comes to any judgment* on the matter, and she tells Ben her opinion. Should Ben now take this information as additional evidence supporting P, over and above that provided by E?<sup>182</sup>

Alice cannot appeal to her own judgment as further evidence supporting P, since that would just be circular reasoning and bootstrapping. The thought here, ostensibly, is that there is also something odd about Ben relying on Alice’s judgment. Christensen clarifies that he does not think higher order evidence should be disregarded or dismissed as ‘not really evidence’. Rather, he takes it to motivate an Independence thesis, according to which Ben may only believe something other than P if he possesses some evidence that is not already encompassed within E, or which does not simply reiterate that Emily believes that P. Again, Alice ought not rely on the fact that she believes that P on the basis of E to increase or strengthen her conviction regarding P, doing so would require further evidence that she does not already regard as justifying her belief that P. Moreover, Christensen also thinks this shows that higher order evidence has agent-relative significance, that the bearing some information has on one agent’s justification need not generalise to that of others:

Alice’s having formed the belief that P on the basis of E is evidence for Ben that E supports P, in the following way: If Ben is otherwise inclined to believe not-P on the

---

<sup>181</sup> Ibid; 177.

<sup>182</sup> Christensen, D. 2010; 188-9. My emphasis.

basis of E, the HOE about Alice's belief should raise worries that he's mistaken about the bearing of E, and he should be less confident in not-P (and more confident in P). Alice's thinking thus serves as for Ben as a check on his own. But it obviously cannot do that for Alice. So for Alice, this item of HOE should not affect her confidence in P.<sup>183</sup>

### 5.6: Why Evidence of Peer Disagreement *is Not* Higher Order Evidence

Conventional accounts of higher order evidence define it as any information pertaining to evidence or rationality, rather than propositional evidence:

Higher order evidence, so understood, comes in different forms, including, but not limited to, evidence concerning:

- (i) What evidence one has available,
- (ii) The strength and valence of one's evidence,
- (iii) The normative import of one's evidence and
- (iv) One's capacity to assess one of more of (i)-(iii).<sup>184</sup>

The fact that Alice believes that P based on E, even if she has heavily reflected on E, does not say anything about any of these points – it does not confirm that P actually follows from E, it tells Ben nothing about the evidence they both have, nor about the strength or utility of that evidence. It could even just be a report of Alice's attitude without any normative judgment about the rationality or justification of Ben's (would-be) belief or his capacity to assess the evidence E, and without any attempt to persuade or criticise involved. Alice's belief and assertion does not fit the criteria for higher order evidence, but it must be evidence of some kind. As I have already maintained, insofar as the fact that Alice believes that P based on E *is* first order evidence, it is evidence about the distribution of opinions about P and evidence that at least one person believes that P based on E.

I contend that Alice's belief that P is not higher order evidence in the sense of it calling into question the justification or reliability of any belief that not-P (based on E). At most it is an item of evidence regarding the distribution of opinion about P with respect to E (i.e., it is evidence that at least one person has a P-related opinion based on E.) Moreover, I do not think that the case of Alice and Ben is much like a peer disagreement either. All we know from the way Christensen sets up the case is that they have shared first order evidence. We do not know any

---

<sup>183</sup> Christensen 2010; 189-190.

<sup>184</sup> Whiting, D. 2020; 789.

details about their epistemic standing, which leaves us to guess whether there are any evidential or intellectual disparities between Alice and Ben, so it is certainly not an *explicit* peer disagreement. I regard it as closer to a typical case from the epistemology of testimony, since Ben has not explicitly formed a belief about P yet.<sup>185</sup> However, it is not clear that Ben actually cares about Alice's opinion or is seeking her advice so it might not even qualify as this.<sup>186</sup> Ben could be ignorant or uncertain about what to believe, and is waiting for Alice to make a judgment before deciding himself, which *would* make it more like a testimony case. But it could just be like an ordinary conversation where Ben has neither a settled opinion nor any investment in the matter.

Even then, if it were a peer disagreement, why would or should we (or Ben) take the fact that Alice believes that P on the basis of E as counting against any possible belief that not-P? I am unconvinced that the Alice/Ben example is higher order evidence *or* peer disagreement, but I am also not convinced that peer disagreements genuinely count as higher order evidence in any substantive way. Higher order evidence is supposed to confirm or disconfirm any supposed justification for believing, so Alice's belief that P would not be enough by itself. If Alice knew that P, or if she justifiably believes that the available evidence is obscure and merely appears to support P, this will suffice. However, she does not have knowledge and instead merely believes that P with evidence that is prospectively inconclusive. By its very nature, higher order evidence cannot directly bear on first order propositions – the conditions above state that higher order evidence is (exclusively) evidence or facts about evidence itself.

Alice and Ben are not disagreeing, but the case could easily be reimagined in such a way that Ben concludes not-P and engages in a dialogue with Alice trying to debate the merits of each position and how each has interpreted the evidence E. Suppose that Alice and Ben are in a peer disagreement, that Alice believes P and Ben believes that not-P. Ben's disbelief tells us and them that their shared evidence is *potentially* inconclusive, misleading, or easily misinterpreted, and nothing else. Alice's belief tells us and them nothing about the way Ben has reached his judgment, such as whether he is intoxicated or impeded in his judgment. Their first order propositional beliefs about P cannot be higher order evidence concerning the justification for P.

The same goes for any other purported peer disagreement: disagreement between the diners in the Restaurant Check case does not *confirm* that either of them has miscalculated their share of the bill, it merely raises the possibility. One may wonder why peer disagreement is

---

<sup>185</sup> Crucial here is "If Ben is *otherwise* inclined to believe not-P..." It is not clear whether Ben does in fact believe not-P on the basis of E.

<sup>186</sup> Moreover, one need not be seeking advice or testimonial evidence specifically because there is a perceived or actual disparity, but many examples in the epistemology of testimony are set up in this way.

perceived as higher order evidence, other than the two being regarded as having similar bearing. One defence relies on the claim that any evidence concerning evidence for a proposition, is still evidence for that proposition. One's epistemic peer made a judgment about some topic, theory, or proposition, all based on evidence and their own good reasoning, and one should take their peer's judgment as a reliable indicator of that claim not merely being justified, but true as well.

As Jon Matheson says:

Since a qualified epistemic peer regarding *p* is likely to correctly evaluate the evidence pertaining to *p*, that such an individual has adopted the doxastic attitude *belief* toward *p* on the basis *E* is reason to believe that *E* in fact supports *p*. In this way, evidence about how a qualified epistemic peer has evaluated *E* is evidence about *E* itself.<sup>187</sup>

Along similar lines, one might argue i) *if* it is unreasonable or improper to cite one's own belief as further supporting evidence for that belief, and ii) there is no reason *not* to consider it some kind of evidence, then one *must* consider it evidence that the relevant belief is justified. In other words, the thought is that if my belief that *p* cannot be used as evidence for *p*, it is surely evidence (for others) that belief in *p* is justified. My belief is not first order evidence for anyone, nor is it higher order evidence for myself, but it can be higher order evidence for you and others.<sup>188</sup> This the point David Christensen makes about higher order evidence having agent-relative significance.

One final point of comparison can be made between the peer disagreement and higher order evidence debates. If we grant that the attitudes, opinions, or assertions from speakers have epistemic weight or significance, then we might ask whether this is a fixed, limited significance or if it changes as we identify more speakers and opinions.<sup>189</sup> After all, many will voice their opinions whether it has been solicited or not, and there can be a substantial consensus belief shared by large collectives of speakers who are entirely unaware of it. Do such examples of consensus matter more than just one agent having the same opinion in isolation, and do they count as higher order evidence? I have argued that an individual speaker's opinion about a matter cannot be higher order evidence, so it follows that I would deny that collective consensus is higher order evidence. Having said that, this need not diminish its import as a reason to reconsider our belief when we find it differs from the consensus. If we are in the minority

---

<sup>187</sup> Matheson 2015; 36.

<sup>188</sup> Not that many would typically cite their own belief that *p* as additional evidence for *p*...

<sup>189</sup> Much of the discussion on this topic is framed around whether agents (typically non-experts) ought to defer to groups (of experts) for various judgments and opinions. See, for example, Easwaran et al. 2016; Fricker, M. 2012; Carter, J.A. 2014; John, S. 2011.

especially, but also if opinions are roughly evenly divided, we ought, I argue, to double check our reasoning and perhaps our dissenters' too. This question occurs in the disagreement literature because for many, the number of speakers opposing us or each other is significant. When we find that we disagree with a majority or when there is roughly equal division, we ought at least to reduce our confidence or conviction in our belief or suspend judgment altogether. However, the matter of group testimony and group agreement or disagreement encompasses various individual debates, even if we don't consider epistemic peerage. More specifically, there is the question of the significance of hypothetical or imagined testimony vs. actual testimony, individual vs. group testimony, the nature of group beliefs or attitudes, and then the question itself of whether and why group agreement or disagreement is epistemically significant.<sup>190</sup>

Hilary Kornblith argues that when opinions on a matter are roughly evenly distributed and there is no substantial majority consensus, it is reasonable and justified to withhold opinion on the relevant matter. The distribution of opinions is not the sole determining factor, but it carries 'a tremendous amount of epistemic weight', meaning contrary belief against a majority requires significant independent evidence which shows why the majority are mistaken.<sup>191</sup> The upshot of Kornblith's argument is a widespread suspension of judgment that ultimately leads to scepticism which I find less than persuasive, but I will say more on this later. Conversely, Adam Elga has argued that outside opinions should be factored into consideration only *iff* they are independent (and, presumably, different) from the opinions one has already weighed up. Initial disagreements will matter greatly but the impact is lessened once disagreement becomes known or to be expected, or as more people continue to reiterate the same opinion.

I contend that group beliefs and assertions are no more higher order evidence than those of individuals, but both certainly are evidence of some sort. As I see it, they are clearly relevant to the question of how opinions stand regarding the relevant matter at hand or proposition; if there is a widespread consensus then it tells us just as much as if opinions are equally divided or if no one has reached any strong judgment at all. I contend that the actual number of people who share a belief does not matter for all that much, rather it is the opinion itself and the reasons for believing it that matters more. Thomas Kelly has argued against merely possible disagreement being insignificant by noting that all the advocates for an opinion could be censored or executed by the command of a tyrannical dissenter. I will also say more on the question of whether actual disagreement matters more than merely possible disagreement later. For now, I contend that having no advocates for an opinion is neither here nor there if all the available evidence supports it. Likewise, we could ascertain that 99% of all people believe that P

---

<sup>190</sup> For further discussion, see

<sup>191</sup> See Kornblith, H. 2010.



with little or poor supporting evidence for P, the poor evidential support surely counts less than the majority belief, if it does not also undermine it. If the insufficient evidential support can be demonstrated, *that* is clear higher order evidence that the majority has made a notable mistake.

Regardless, higher order evidential defeat occurs when we have or gain some reason to believe our reasoning is hindered or faulty, or that we have overlooked some bias. Peer disagreement does not ensure that as a matter of fact; at best, it signals the possibility of this occurring, but this possibility could easily run both ways or the matter of who is mistaken or biased could simply be left open and unclear, which is how most peer disagreements are presented in the literature. The fact that two peers disagree on a matter is just evidence about opinion distributions. Some authors in the peer disagreement discourse would disagree with my judgment because they believe that evidence of peer disagreement genuinely *is* higher order defeating evidence. For example, Richard Feldman:

... if the evidence about the other person would be epistemically significant in cases in which one did not have a view of one's own, is it hard to see why it has no impact when one does. You have to think that somehow having your own view renders it completely insignificant. That seems implausible.<sup>192</sup>

The idea is not that one has to weigh the original evidence supporting the proposition against the higher order evidence that goes against the proposition. Rather, the higher order evidence serves as a defeater of the original evidence, in the sense that the conjunction of the original evidence and the higher order evidence fails to support the proposition.<sup>193</sup>

Or Bryan Frances and Jon Matheson:

The reason higher order evidence is important for the topic of disagreement has to do with testimony. If you come to believe something on the basis of some body of evidence, and then you learn that a great many intelligent people believe the exact opposite, that testimonial fact – that they are so intelligent and yet disagree with you – is evidence E<sub>1</sub> that you might have missed something in coming to your belief. What the “missing” amounts to isn't clear.<sup>194</sup>

---

<sup>192</sup> Feldman, R. 2009; 306.

<sup>193</sup> Ibid; 308.

<sup>194</sup> Frances, B. 2014; 42.

Applied to cases of peer disagreement, the first order evidence is the evidence directly pertaining to the disputed proposition, and each peer opinion about the disputed proposition is the higher order evidence (it is evidence that the first order evidence supports the respective attitudes).<sup>195</sup>

I contend that this results from a mistaken conflation of peer disagreement and higher order defeat. If Alice and Ben were peers and were therefore equally intellectually virtuous, well-informed, historically accurate, and so on, then Alice's belief is a strong indicator of the *likelihood* that P follows from E, or at least that it is well supported by E. It does not give any firm confirmation that either one of them is fallible or flawed. The fact that they share evidence and have roughly or exactly similar intellectual capacities and dispositions is enough that it should give both speakers a knee-jerk reaction and signal that they ought to reconsider their reasoning. However, I contend that this is a point about disagreement, rather than something that can be attributed to any higher order defeat.

Alternatively, perhaps the mistake lies in over-inflating the notion of higher order evidence to accommodate it as a matter of convenience. That is, perhaps authors regard peer disagreement as akin to, or an example of higher order evidence because it suits the purposes and fits with their preferred response to peer disagreement. The Total Evidence View relies on the equivocation of facts about peer disagreement as psychological evidence about the beliefs of those involved in the disagreement. As we will see, it recommends that disagreeing peers should reconsider their beliefs by factoring in this evidence alongside the shared first order evidence they each take to favour their respective, incompatible, views. It need not rely on any claim that peer disagreement is higher order defeating evidence, but one may make such a claim to endorse it. Epistemic peers are fallible thinkers and speakers, their beliefs and assertions are the product of their reasoning and evidence and an expression of their expected outcome. I contend that peer disagreement functions in two ways and no more: it signals the possibility of an error without specifying or discriminating between the speakers involved, and it provides first order evidence concerning the distribution of opinions about the disputed matter. In other words, the fact that two peers disagree about P is evidence *merely* for the fact that there is at least some disagreement or dissent about P. I remain unconvinced that higher order evidence has the

---

<sup>195</sup> Frances, B., and Matheson, J. 2019; §5.2.3.

bearing or significance it has been presented as having in the peer disagreement discourse, and I regard it a mistake to conflate the two phenomena.<sup>196</sup>

### 5.7: Summary Thoughts

I have revisited the debates concerning the uniqueness of evidence and the import of higher order evidence and higher order defeat, examining how the two interact and overlap with the peer disagreement discourse in greater depth, as well as the extent to which the peer disagreement discourse depends on former two debates. Both discussions have notable discord and a lack of consensus; I am unconvinced by the extremes of either position in the uniqueness debate, and the hostility directed towards epistemic akrasia and higher order defeat in the higher order evidence discourse. Moreover, there is too great a reliance on stricter forms of uniqueness and the assumption that peer disagreement is higher order evidence in the peer disagreement discourse, which often arises to motivate the specific views of response in that debate. I have argued for the following claims:

- i. Extreme or strict forms of both Uniqueness and Permissiveness are unconvincing. Stringent readings of uniqueness place excessive constraints or demands on rational action and require too much of agents.
- ii. Likewise, the extreme versions of permissiveness make rationality too easily obtained; allowing for any attitude to rationally follow from an item or body of evidence trivialises rationality and makes it a free for all.
- iii. A strict version of uniqueness must be assumed in order to motivate interest in the Equal Weight View, and in the wider peer disagreement discourse overall. If evidence is permissive, or if it at least sometimes allows for multiple rational attitudes, peer disagreement need not be any great concern.
- iv. Higher order evidential defeat (or confirmation) is a highly plausible concept, and it would be strange to try to deny it outright. However, the problems of epistemically akratic action that follow from granting such defeat need not worry us so much that we feel it necessary to disengage from all inquiries and epistemic actions entirely.
- v. Much of the interest in peer disagreement can be attributed to an equivocation of peer disagreement and higher order evidential defeat. However, it is a mistake to try to conflate the two for the following two reasons:

---

<sup>196</sup> As we will see, I am sceptical that peer disagreement even matters all that much, but that is more for reasons to do with peerage than scepticism about the epistemic import of disagreement. I am very firmly convinced that disagreement *is* epistemically significant, but also that peerage adds little, if anything, to this.

- vi. Rather than defeating the shared first order evidence from which peers draw different conclusions from, peer disagreement merely signals the possibility of an error without specifying who it applies to.
- vii. Peer disagreement also functions as evidence regarding the distribution of opinion over the disputed matter, but not as evidence for or against the matter itself.

Overall, I think peer disagreement infringes on and crowds out two debates that are independently substantially interesting and deserve independent consideration. I now address peer disagreement exclusively by raising several concerns I believe are both substantial and stimulating.

## Chapter 6: Peer Disagreement: Uncompelling, Inconclusive, and Insignificant

Over the previous two chapters, I introduced and critiqued the debates concerning the prospect of evidential uniqueness, the bearing of higher order evidence and higher order defeat, and disagreement between epistemic peers. I have argued that the peer disagreement discourse largely subsumes the other debates and requires a particular position in each debate in order to motivate its philosophical intrigue. Unless i) evidence admits of a strictly unique reading (i.e., one exclusively rational attitude) and/or ii) peer disagreement *is* a form of higher order evidential defeating evidence, there is very little reason to be struck or concerned by peer disagreement. I am unconvinced that evidence is either strictly unique or maximally permissive, or that we can ever completely rule out or explain away the possibility of higher order evidential defeat. Where there is any distinct possibility of higher order defeat, it does *not* come from establishing epistemic peer disagreement. I have made the following claims, so far:

1. Stricter readings of uniqueness are overdemanding, unattainable and unrealistic, whereas the extreme readings of permissiveness reduce rationality to a free for all. The optimum is a liberal reading of uniqueness that strikes a balance between the two extremes.
2. Higher order evidential defeat should not be disregarded as a possibility, but it need not always be a bad thing that we are compelled into epistemically akratic actions. Peer disagreement doesn't reflect higher order defeat, it merely mimics its effects.
3. The discourse concerning peer disagreement relies on, and presumes, too much with regards to uniqueness and higher order evidence. Much of the arguments are rooted in some form of evidential uniqueness and a false equivocation of peer disagreement and higher order defeat.

Over the remaining three chapters, my aim is to address the peer disagreement debates more directly by making the following claims:

- i) The various purported replies to peer disagreement, which are offered as rational or rationally required responses, do not stand up to criticism.<sup>197</sup> Some lead to universal dogmatism or suspension, others are too overdemanding to hold true.  
Call this the claim that the discourse is *uncompelling*.
- ii) The peer disagreement discourse doesn't fulfil its aim of explaining why disagreement is epistemically significant. Moreover, none of the replies to peer disagreement fulfil an unspecified but plausibly important crucial aim of developing or advancing inquiries.

---

<sup>197</sup> I will argue that they also do not amount to much when applied to instances of *peer agreement*. I develop this criticism in the following chapter.

- Call this the claim that the peer disagreement discourse is *inconclusive*.
- iii) The peer disagreement discourse both fails to compellingly specify why disagreement (*simpliciter*) is epistemically significant, but it overlooks a simple but weighty account of that significance. Call this the claim that (the phenomenon of) peer disagreement is *insignificant*.
  - iv) Where peerage is a factual state, at best this is a product of circumstances being epistemically fortuitous. Speakers rarely meet the conditions for peerage simultaneously, but they are far more likely to regard or perceive each other as peers or equals.  
Call this the claim that peerage is *unattainable*.
  - v) Peerage is an unrealistic phenomenon with unattainable conditions and a lack of clarity over the extent to which they must be met to qualify as peers. The discourse relies on distortion and misrepresentation of real life and experiences to justify its own continuity.  
Call this the claim that the peer disagreement discourse is *idealised*.
  - vi) The discourse focuses almost exclusively on disagreement between epistemic peers, overlooking peer agreement altogether and limiting consideration of non-peers to inter- and intra-group disagreement and disagreements with or between experts.  
Call this the claim that the discourse is *inattentive*.

I will address each claim in turn, focusing on claims i-iii in this chapter, iv and v in Chapter 7, and vi across Chapters 7 and 8. I will begin by summarising the four views that I regard the main contenders for how we ought to respond to peer disagreement and argue that they are all unsatisfying and leave much to be desired. The claim that any fitting response to peer disagreement ought to advance or progress the relevant inquiry has not, to my knowledge, been offered so far. In arguing for it, I will draw on considerations about what, if anything, we are aiming for when we inquire, emphasizing that disagreement (and agreement) are essential parts of inquiries.

The insignificance claim follows on from the concerns I raise with regards to the replies to peer disagreement, but I am not the only one to have doubts about the discourse's success in identifying why disagreement matters epistemically. I draw on an account of the epistemic significance of disagreement that I regard as compelling, which is entirely free of dependence on peerage and examples of disagreement in narrowly hypothetical or unrealistic circumstances. Others within the literature have argued that there is a high degree of idealisation in peer disagreement, with most acknowledging it and then setting the point aside. It is only recently that the problems with idealisation have become more widely discussed and the discourse criticised.

Most arguments accept that the conditions for peerage are individually difficult and collectively impossible to meet, but few among them explicitly state that epistemic peerage doesn't accurately reflect real human capacities and experiences. Most authors on peer disagreement presume that we only learn why disagreement is epistemically significant by focusing on unrealistic hypothetical cases involving epistemic peers, that we can generalise our response from them to ordinary cases. This distorts reality, not only excessively but, considering my earlier claims, unnecessarily and for little to no benefit.

The claim that peerage is a matter of perception directly contrasts one of the prevailing assumptions within the literature, namely, that peerage is a (factual) state determined by speakers meeting conditions of equality in various respects. However, there is an important distinction between speakers who are peers as a matter of fact, and speakers who merely regard each other as peers. If peerage is a factual state, it's not clear if the conditions for it are necessary or sufficient, and this not often been explicitly discussed. Finally, the claim about the inattentiveness of the discourse is also novel and leads directly to my overall concern that the epistemology of disagreement is hollow or lacking in substance. For all the focus on peer disagreement bearing (I argue) little to no success, progression, or closure, *peer agreement* is never even raised as a mere possibility, let alone given its due consideration. Because peer agreement is given no attention, it means that no-one expresses any concern about it being disregarded, which means that this concern is arguably the most overlooked of the issues I raise. While some attention is given to agreement and disagreement between speakers who aren't epistemic peers, it is severely limited to cases that involve either groups, experts, or both. Even then, it will become clear that terms such as "epistemic non-peer", "epistemic superior" or "epistemic inferior" are vague and imprecise, covering many different possible relations between speakers and grouping them together.

The concerns I will argue for, point to an overall worry that the epistemology of disagreement, as it currently stands, has an excessive and unpromising emphasis on agent epistemic rationality, and that it lacks substance and depth by overlooking crucial matters and, in the case of overlooking peer agreement, committing a fundamental misstep. I contend that any epistemology "of disagreement" should accommodate and examine a variety of cases, not simply those that involve (purported) epistemic peers. More importantly, it should acknowledge that agreement is certainly pervasive enough to warrant consideration, and arguably at least as significant as disagreement. After all, how can there be substantial discussion of why disagreement is epistemically significant, but no consideration or comparison paid whatsoever to agreement? This question will feature frequently through the following three chapters, let us



begin with the crux of the discourse: the debate concerning how to reply or respond in instances of peer disagreement.

### 6.1: Peer Disagreement is Uncompelling (In Four Parts)

In my initial characterization of the peer disagreement discourse, I contended that it has two chief aims: the first aim is to identify and demonstrate why disagreement is epistemically significant, and the second aim is establishing what responses are rationally required or rationally optimal for speakers that acknowledge they are in a peer disagreement. However, the discourse has a bit more going on than these two aims alone. We've seen that peer disagreement relies on at least partial attempts to address the debates regarding uniqueness and higher order defeat, but the discourse does not (directly) address the matter of the nature of disagreement, as I did in the first part of this thesis. It emphasizes hypothetical cases of disagreement in belief between epistemic peers, and whether it is intentional or not the discourse typically excludes disagreement that does not involve doxastic or credal states:

My aim... is to explore the *epistemic significance of disagreement*. A central concern is whether the practice of retaining beliefs that are rejected by individuals over whom one claims no epistemic advantage is a defensible one. It is, of course, far from clear that the relevant practice *is* defensible.<sup>198</sup>

How should I react when I discover that my friend and I have very different beliefs on some topic? Thinking about belief in a quantitative or graded way, the question concerns cases in which my friend and I have very different degrees of confidence in some proposition P. Should my discovery of her differing degree of belief in P lead me to revise my own confidence in P?<sup>199</sup>

This book is about disagreement, but what is disagreement? Two individuals *genuinely disagree* about a proposition just in case they have adopted incompatible doxastic attitudes toward that proposition.<sup>200</sup>

The implication is that peer disagreement (or any epistemically significant instances of disagreement) only involves speakers with no discernible differences and common evidence, who

---

<sup>198</sup> Kelly, T. 2005; 168.

<sup>199</sup> Christensen, D. 2007; 188.

<sup>200</sup> Matheson, J. 2015; 6.

happen to differ in their belief or degrees of confidence regarding some proposition or hypothesis. While I am tempted to say that clarifying what disagreement *is*, is a preliminary to the question of why disagreement matters and to call it a misstep that the nature of disagreement is not discussed in the epistemology of disagreement discourse, the two are very much separate matters. Nonetheless, I do think that any substantial philosophical analysis of disagreement should give at least *some* consideration to both topics. Even if, as I have argued, specifying what disagreement is in narrow or simple terms is difficult, ascertaining what is epistemically significant about disagreement will require some clarity as to what disagreement *is*, either generally or specifically.

While the peer disagreement debate is typically presented as a debate between conciliating and remaining steadfast, I believe the Equal Weight View is perhaps the focal point of many of the responses in the debate and thus the main target for dissenters. I will argue that none of the typically cited positions in the discourse, which are either derived from Equal Weight or offered in response to its shortcomings, are especially compelling. I will argue that some of them risk collapsing into Equal Weight or place overdemanding expectations or requirements on agents, and that they all fail to achieve their aim of demonstrating what is epistemically significant about disagreement. I will argue that they all fail to adequately advance or progress debates and inquiries which, to my knowledge, is a criticism that has not been offered in the discourse so far. Disagreements occur naturally in many inquiries, so if peer disagreement *is* a live possibility, the lack of progress should be acknowledged, and subsequently any response to peer disagreement worth its salt must avoid stopping a debate or inquiry in its tracks and be able to advance or progress discussions. All the main contender replies to peer disagreement lead to stagnation, either from indefinite belief suspension or dogmatism.

## 6.2: Equal Weight

Let's begin by revisiting the Equal Weight View, the claim that epistemic peers who disagree each ought to consider the other's view at least as likely and as important as their own, and either revise their attitude or degree of confidence such that they mirror each other or suspend belief.

But absent some special reason for treating her evidence or reasoning differently from mine, it seems clear that in evaluating explanations for our disagreement, I should regard our opinions as equally likely to be accurate, and thus I should alter my opinion towards hers.<sup>201</sup>

---

<sup>201</sup> Christensen 2007; 212.

I originally discussed this view in the context of the uniqueness of evidence, noting that a strong reading of the Uniqueness Thesis underpins Equal Weight. In addition, I discussed Thomas Kelly's criticism that Equal Weight doesn't accommodate the fact that there is disagreement between two peers as relevant (higher order) evidence. Equal Weight's treatment of disagreeing peers is compelling, it acknowledges the fact that they are in epistemically symmetric positions and does not discriminate between rational and irrational speakers. Unfortunately, the view does little more than compel or demand that each peer abandons their respective view for an indefinite and unspecified length of time. Peers are left to simply withdraw from engaging in the disagreement (and thus, the inquiry) until such time as additional evidence is acquired that breaks the symmetry or one peer concedes to the other, which may well be indefinite.

The lack of discrimination also means that both speakers are compelled to change their view and suspend belief, when one of them could have initially had a justified, rational, or accurate belief. Rational agents are compelled to suddenly act in a way that is rationally sub-par, essentially because of fairness or parity. I contend that any advocate of Equal Weight who does come to disagree with an epistemic peer may as well shrug their shoulders and say "Huh... so we disagree, how about that." We can reduce our degree of confidence or conviction, perhaps even temporarily suspend belief on the basis of peer disagreement (perhaps even in non-peer disagreements as well), but we surely have to move on eventually. On the one hand, it is only sensible for a speaker to compromise with their peer if they were mistaken from the outset and would have believed irrationally regardless, or for both to suspend judgment *iff* they would both have drawn a wrong conclusion. On the other hand, it is entirely plausible that any rational agent should strive to help irrational agents improve and believe correctly.<sup>202</sup> Unfortunately, Equal Weight is incapable of delivering this because the rational peer would need foreknowledge concerning the relation between the evidence and the disputed claim. Were one speaker to possess such knowledge, the prospect of disagreement becomes implausible, and peerage falls out of consideration because knowledge provides a clear epistemic advantage.

Thomas Kelly claims that Equal Weight cannot make sense of the higher order evidence that arises out of facts concerning peer disagreement and he offers various arguments that purport to demonstrate its odd and undesirable consequences. First, there is an appeal to intrapersonal conflict – the recognition of an internal inconsistency between a pair of one's attitudes. We can imagine a view that rationally requires peers to abandon both their respective beliefs when they discover that they are inconsistent in this way, and such a view is the closest

---

<sup>202</sup> I will say more about this in Chapter 8, particularly with regards to agreement and disagreement between '*non-peers*.'

analogue to Equal Weight for resolving internal inconsistencies. Kelly contends that we have little reason to motivate such a view about intrapersonal conflicts, since there might be equal or stronger reason to only drop one belief and not the other. If one's total evidence supports one view over the other, this surely overrides the reason to abandon both beliefs that comes out of the fact that it is inconsistent with another of our beliefs. Kelly thinks that the same can be said of following the Equal Weight View and revising our beliefs in peer disagreements, since the requirements are essentially the same:

A normative view about how it is reasonable to resolve inconsistencies among one's beliefs that completely abstracts away from facts about which beliefs are better supported by one's evidence, and that would have one treat one's prior beliefs on a par, regardless of how well or ill supported they are by one's total evidence, would not be an attractive one. But the features that make such a view unattractive are shared by the equal weight view.<sup>203</sup>

The second of Kelly's concerns is that Equal Weight licenses a form of *bootstrapping*: believing that one's beliefs or attitudes are reliable without any justification or knowledge. He regards it as dubious to think that one can merely rely on one's own belief as reliable or as justifying a high degree of confidence, when one's evidence could just as easily be misleading. My point about Equal Weight compelling irrationality is derived from this claim. For Kelly, even if epistemic rationality requires that disagreeing peers split the difference between their attitude, doing so will not ensure or guarantee that their subsequent attitudes are rational, since they could both have evaluated wrongly in the first place.<sup>204</sup> On the Equal Weight view, rational belief requires ideal or a God's eye rationality, or foreknowledge of what rationality requires. I contend that similar issues arise for at least one competitor to Equal Weight, but I will say more on this later.

Kelly's third argument against Equal Weight is based on his observation that in at least some cases there will be a scarcity of propositional evidence and one may be limited only to "psychological evidence" concerning the distribution of opinions. That is, there will be cases in which our only evidence with which to form a belief about some matter,  $p$ , is evidence to the effect that "Speaker S believes that  $p$ ", whereas in other cases we will have no evidence of this sort whatsoever. An Equal Weight proponent will have to concede that, in the latter case, one's judgment can only be determined by one's propositional evidence, and that the psychological evidence would otherwise swamp the former:

---

<sup>203</sup> Kelly, T. 2010; 192-3.

<sup>204</sup> Ibid; 193-4.

In general, for any time  $t_n$ , a proponent of the equal weight view will hold that what one is rationally required to believe is entirely fixed by the opinions of the  $n$  peers. Why then should things be any different at time  $t_1$ , when the number of peers is 1? It seems as though the only principled, not ad hoc stand for the proponent of the equal weight view to take is to hold that the psychological evidence swamps the nonpsychological evidence even when the psychological evidence is exhausted by what you yourself believe.<sup>205</sup>

The Equal Weight View requires only that disagreeing peers should give equal weight to each other's views, even if one of them has misread or miscalculated things, arrived at the wrong conclusion, or otherwise exhibited some error in performance. If the number of peers and peer-opinions increases, each individual peer's actual performance and reasoning becomes benign.<sup>206</sup>

Kelly's fourth objection is based on the observation that we will generally accept non-agential evidence (such as a litmus paper) as consistently reliable and often without question, and that we regard a record of good judgment as indication of an agent's reliability. If we suppose that our non-psychological evidence for a proposition is substantial but our psychological evidence is limited to two items – the fact that I believe it and the fact that you do not, respectively – Equal Weight effectively sidelines the psychological evidence. Kelly claims that it is far from clear why the non-psychological evidence should be set aside and ultimately given no weight at all in favour of only the psychological evidence.<sup>207</sup> Kelly takes these reasons to clearly indicate the inadequacy of Equal Weight, and proposes the Total Evidence View as an alternative, which I will discuss shortly.

One final and important worry for the Equal Weight View can be attributed to Adam Elga, who claims that it not only compels disagreeing peers to reduce their confidence and suspend belief, but “to think, of each stance, that it is very unlikely to be right.”<sup>208</sup> The implication is a kind of *spinelessness*. Endorsing Equal Weight is a commitment to pre-emptively avoiding or suspending belief for any matter in which we think there is a high possibility of peer disagreement or in which we genuinely experience peer disagreement. Elga thinks that we can get around the worry of spinelessness by distinguishing clean, simpler disagreements that more naturally lend themselves to Equal Weight from messier disputes on “hard issues.” Topics of the latter kind are those for which endorsing Equal Weight leads to absurd conclusions and

---

<sup>205</sup> Ibid; 195.

<sup>206</sup> Ibid; 196.

<sup>207</sup> Ibid; 197.

<sup>208</sup> Elga, A. 2007; 484.

misguided judgment suspension. As will become clear, I am unconvinced that genuine (factual) peer disagreements are a live possibility, but I agree that taking Equal Weight to its logical conclusion leads to stagnant inquiries and even an intellectual defeatism. I agree with Elga's assessment that some matters lend themselves more easily to identifying peers and others less. However, Elga's conclusion that we should suspend judgment in many disagreement cases because they are over subjects that cannot be narrowly refined is one I regard as more about getting around spinelessness than genuinely trying to vindicate Equal Weight. Even if we think the chances of identifying epistemic peers about some matters are low, we should not consider *indefinite* belief suspension to be our default response. It only stunts inquiries and debates in the long term, and forces otherwise rational or justified speakers into epistemic akrasia in the short term.<sup>209</sup>

### 6.3: Extra Weight

While Thomas Kelly proposes the Total Evidence view in response to Equal Weight, I will first consider the Extra Weight view that Adam Elga offers. For Elga, this is the view that “you should be pulled a *bit* in the peer's direction – but not halfway. You should still think it is more likely that you are right than that the peer is.”<sup>210</sup> Similarly, for Thomas Kelly, it is the view “according to which each party to the dispute is permitted to give some special, presumptive weight to his or her own judgment.”<sup>211</sup> The thought is that disagreeing peers should each give the other's view due consideration and value it as epistemically weighty, while also granting it rationally permissible for each to privilege themselves and value their own opinion view as more likely right, significant, justified, or rational than their opponent's.<sup>212</sup> Elga thinks the appeal of this view lies in the fact that it avoids the problem of spinelessness, forced akrasia, indefinite belief suspension, and a lack of self-trust that follows from Equal Weight. However, it instead faces charges of dogmatism and the unwarranted privileging of one's own beliefs and attitudes. Dogmatism and privileging are only surface level issues, however. Equal and Extra Weight are only about how disagreeing peers rationally ought to respond to each other. Neither view, I contend, reveals whether either peer's attitude is genuinely justified or rational, nor do they have

---

<sup>209</sup> I also think Equal Weight faces worries when considering group disagreements or disagreements about matters that admit of more than a straightforward affirmation, denial, or judgment suspension.

<sup>210</sup> Elga 2007; 485.

<sup>211</sup> Kelly 2010; 161.

<sup>212</sup> There is an important difference concerning “more likely right” here; peerage rests on either a mutual perception between speakers (i.e., that they both think they are equally likely to be correct and rational, based on related previous exchanges) or a fact of the matter (i.e., these speakers *are* peers.)

But the use of “more likely right” that is relevant for Extra Weight refers to thinking that one's own attitude is, as a matter of fact, correct on this topic. Thus, the speakers remain peers and there is no need for them to question their peerage.

the resources that enable peers to work out for themselves anything regarding the accuracy, fittingness, or justification of their opinions.

Ordinarily one might want an analysis of disagreement to yield a solution to that disagreement, where possible; for example, in simple arithmetic cases such as the Restaurant Check case, we do not want to ascertain just whose judgment the evidence best supports, but whose judgment (if any) is true and correct. Neither Equal nor Extra Weight deliver any such result.<sup>213</sup> In other words, both Equal and Extra Weight are hindrances, preventing most relevant notions of progress within an inquiry or disagreement: Equal Weight forces suspension on both sides whereas Extra Weight permits mutual dogmatism. Two peers could endorse Extra Weight, find out that they disagree about P, but agree that neither of them is required to change their mind, or even to try to explain themselves or reconsider, even if it turns out that one or both of them has genuinely misjudged things. In this respect, it shares a motive with Equal Weight because both views preserve the parity between epistemic peers and compel them to respond in the same way. The privileging that Extra Weight permits faces challenges of the sort raised against self-trust in the epistemology of testimony, pertaining to whether we are *a priori* entitled to trust our own perceptions and attitudes over others. Kelly notes that a speaker could make consistently poor judgments and cite Extra Weight as a reason to not concede when they encounter a disagreeing peer. One might use the same line of reasoning to think they are epistemically superior (which would clearly be absurd) but the justification for Extra Weight just is that some speakers and even some peers judge poorly!<sup>214</sup> Extra Weight gets into further complications when we consider its application to agreement and disagreement with non-peers, but I will say more on this later.

Let me make one final point. Extra Weight licenses privileging one's own view over one's opponent in a peer disagreement, but it is not clear from the Extra Weight view whether doing so is *evidentially* justified. Perhaps we are *a priori* epistemically entitled to self-trust, but the fact that epistemic peers disagree over some matter with shared evidence does not ensure a clear answer. We can grant that there may be some absurd cases where one view is obviously wrong. However, without any way to determine that one peer's attitude is more (or most) justified, fitting, or true, we would only be able to dismiss views arbitrarily, which is clearly unreasonable and counter intuitive. It is for this reason that some proponents of Equal Weight explicate a requirement for *dispute-independent evidence* to break the symmetry to justify privileging. The thought is that we may only favour our own view after ascertaining that we disagree with an

---

<sup>213</sup> However, again, this is not within the remit or aims of the epistemology of disagreement, and certainly not an aim of the peer disagreement discourse, since it is an empirical matter.

<sup>214</sup> See Elga 2007; 487. See also Kelly 2010; 211.

epistemic peer *iff* we have some evidential defeater that either has not already been used to justify our belief or does not relate to the fact that we disagree with our epistemic peer.<sup>215</sup>

I am unconvinced that peer disagreement generates or constitutes higher order defeating evidence. The mere fact that an epistemic peer disagrees with us does not by itself challenge the reliability of our judgment-forming, but I do grant that genuine peer disagreement may give participating peers reason to doubt what to believe or whether to maintain their attitude. However, this position will be contentious, perhaps even problematic since many regard doubts and doubt-inducing objects as clear epistemic defeaters. Jennifer Lackey, in particular, contends that beliefs and attitudes can be defeated by two kinds of belief or doubt. There are *psychological defeaters*, which are a speaker's own doubts or beliefs about the reliability of their belief-forming or belief-acquiring process with regards to another given belief. Secondly, there are *normative defeaters* – doubts which one *ought* to have that indicates that one's belief is either false or has been formed unreliably.<sup>216</sup> The consensus view would seem to be that peer disagreement functions initially as a defeater of the second kind but instils in each peer doubt of the first kind, thus undermining their judgment and generating reason for each peer to reconsider their own attitudes. I grant that the acknowledgment of peer disagreement would give each peer reason to reconsider in this way, but I am simply not convinced that *that* reason is equivalent or synonymous with higher order defeat, or that it is sufficiently strong enough to challenge the reliability of one's judgment or reasoning. More generally, it simply does not strike me that the fact that two epistemic peers disagree on a matter or judgment based on common evidence is anything more than a fact or evidence pertaining to the distribution of opinions on that matter.

When we have this kind of uncertainty, especially in cases like peer disagreements where evidence is shared and seemingly inconclusive, neither is it enough to rely on self-trust to justify steadfastness.<sup>217</sup> Both Equal and Extra Weight are inadequate and unconvincing as replies to peer disagreement and they fail to deliver on the aims of the discourse. Equal dogmatism and equal indefinite suspension are equally poor outcomes and neither shows why disagreement (in general) is epistemically significant or indicates which attitude (if any) is more or most rational, justified, or fitting. Neither view is apt to resolve or advance inquiries and debates, both lead to an undesirable stalemate, with Equal Weight requiring at least a partial downgrading of rationality and Extra Weight permitting belief-privileging that's aptly been described as chauvinistic.<sup>218</sup> Therefore, we need a better alternative.

---

<sup>215</sup> See Christensen 2007; 212.

<sup>216</sup> See Lackey 2006; 438.

<sup>217</sup> Even though I am sceptical about the possibility of genuine peer disagreements, there are plenty of disagreements between speakers who have common evidence.

<sup>218</sup> E.g., Matheson: "it seems to be an act of hubris to cite *one's own belief that p* as evidence for *p*." 2009; 272.



#### 6.4: Right Reasons

Of the other responses to peer disagreement, there are two more that I will consider, one being the position Adam Elga names the Right Reasons view and attributes to Thomas Kelly. In summary, this is the position that each peer's rationality, given their disagreement, is determined by whether they had *in fact* correctly reasoned or interpreted the evidence and believed or judged accordingly. I will explain how it goes by borrowing an example from Michael Titelbaum involving two peers, Greg and Ben, who share evidence related to a hypothesis *h* but disagree over *h*:

At *t*<sub>1</sub> Greg and Ben have adopted their own attitudes towards *h* but each is ignorant of the other's attitude. At *t*<sub>2</sub> Greg and Ben discover their disagreement about *h*. They then have identical total evidence *E'*, which consists of *E* conjoined with the facts that Greg believes *h* on the basis of *E* and Ben believes  $\sim h$  on the basis of *E*. The question is what attitude Greg should adopt towards *h* at *t*<sub>2</sub>.<sup>219</sup>

Importantly, *E* is understood simply as Greg and Ben's shared body of propositional evidence for *h*. Right Reasons acknowledges that additional evidence comes from the fact that they acknowledge their disagreement – the fact that Greg believes and asserts *h* and the fact that Ben believes and asserts not-*h*. Advocates for the view maintain that Greg and Ben should each reconsider their views, factoring in these additional pieces of evidence. However, they also typically claim that the two pieces of psychological evidence neutralise or cancel each other out, leaving both speakers only with their original shared evidence *E*. What is then required of each speaker depends on which speaker (if either) is correct, and whether this is known or made explicit.

As Kelly and Elga both argue, Right Reasons says that you are rationally permitted to stick with your initial judgment if and only *iff* you initially evaluated the evidence correctly:

If you and I have access to the same body of evidence but draw different conclusions, which one of us is being more reasonable (if either) will typically depend on which of the different conclusions (if either) is in fact better supported by that body of evidence. No doubt, especially in the kinds of cases at issue, it will often be a non-trivial,

---

<sup>219</sup> Titelbaum, M. 2015; 283.

substantive intellectual task to determine what the totality of the relevant evidence supports.<sup>220</sup>

For example, suppose that the shared batch of evidence *in fact* strongly supports the disputed claim. You correctly apprehend this, but your peer misjudges the force of the evidence and as a result disagrees. In this case the right-reasons view says it can be reasonable for you to stick to your (correct) evaluation.<sup>221</sup>

On another understanding that is concerned with evidential relations more generally, Right Reasons is the claim that any judgment that results from “actually correct reasoning” should never be subsequently modified because of higher order evidence or “calibrating features.”<sup>222</sup> Speakers who misjudged, misinterpreted the evidence, or were otherwise wrong before ascertaining their peer disagreement, should now adopt the right view. As Titelbaum sets up the example with Greg and Ben, their shared propositional evidence strongly supports Greg’s belief and not Ben’s. Right Reasons therefore says that because it is Greg who has reasoned correctly, it is Greg who is permitted to retain his belief, whereas Ben was mistaken beforehand and must subsequently change his view. The result is similar to the Extra Weight view, but it differs in an important respect because Greg’s steadfastness is not permitted because of any entitlement to self-trust, but because the evidence genuinely favours his belief. Likewise, Ben’s requirement to change his belief is not because of any conciliatory pressure from Greg’s disagreement, but because he had genuinely misjudged things. So far, so good. However, this interpretation faces issues when we consider that not all peer disagreements are set up such that we know exactly which peer (if either) is correct. Indeed, most cases are set up in such a way that the correct judgment is a complete mystery. After all, how can we say that one disagreeing peer must change their view, but not the other, when all that is available is either neutral psychological evidence or non-psychological evidence read in conflicting ways? This uncertainty brings about the same results as Extra Weight, and therefore the same problem – dogmatism for everyone becomes the default. Just as was the case for Extra Weight, this is neither compelling nor productive for inquiries.

---

<sup>220</sup> Kelly 2005; 180.

<sup>221</sup> Elga 2007; 485. Taking the lead from Kelly, Elga distinguishes between extreme versions which rationally permit a rational agent to be completely unmoved, and modest versions that permit them to be moved less than halfway towards their peer’s judgment. 2007; footnote 20.

<sup>222</sup> See Steel, R. 2019; 432. Calibrating features are factors which indirectly bear on our beliefs and justification and indicate inability to reliably assess evidence, such as biases or, indeed, testimony. In other words, higher order evidence.

In his initial characterisation of the view, Kelly argues that it is permissible for a rational agent to maintain their attitude in virtue of having reasoned correctly. In his later analysis, however, he modifies his position and contends that one should move at least partway in the direction of one's peer's view, "even if one responds to the original evidence in an impeccable manner and one's peer does not".<sup>223</sup> Ultimately, Kelly comes to the same conclusion as myself concerning the inescapable uncertainty about who, if anyone, has reasoned correctly. He discusses an example of a professional mathematician who believes they have discerned a proof for a complicated mathematical conjecture, but who later consults various colleagues who all independently declare the proof to be unsound. For Kelly, this is sufficient to indicate that the mathematician should be less confident in the conjecture being true, and that this downgrading of confidence is rationally required from the fact that all the peers drew the same contrary judgment. Kelly contends that since there is a requirement to downgrade one's confidence from the fact that all of one's peers disagree, the fact that a single peer disagrees must also generate at least a partial requirement to downgrade. If giving zero weight to the opinions of others were rationally permissible in that sort of situation, then the mathematician is left in the same state of hesitant uncertainty they were in before consulting their first peer. One rarely occupies the "God's-eye points of view" about what the evidence indicates, so one may have reason to doubt their judgment, which (purportedly) is just the sort of reason provided by peer disagreement. Disregarding a peer's dissent presumes that one's peer has misjudged.<sup>224</sup>

This is why I reject Right Reasons. Not only is it implausible that we can know, from a God's eye perspective, that our evidence conclusively and indefeasibly supports our judgment, but we will often think it is likely that we have, even despite there being disagreement on such matters. Disagreement should throw us into a position of caution or uncertainty, but it doesn't ever tell us what attitude is justified, fitting, or correct. Much like Equal Weight, it presumes a strict version of Uniqueness that grants only one rationally permitted attitude from the evidence, but the idea that only the person who *has* responded correctly can retain their attitude is implausibly overdemanding. What grounds does any speaker have to think they have reasoned correctly other than the fact that they have interpreted the evidence and subsequently arrived at their judgment? It's not enough for an agent to merely suspect or believe that the evidence justifies their attitude in order to retain it when they disagree with a peer, they must know that they are in the right, otherwise why bother endorsing Right Reasons at all? Even then, if one speaker knows that the evidence justifies their belief or attitude, this is clearly an advantage over

---

<sup>223</sup> Kelly 2010; 198.

<sup>224</sup> Ibid; 199-200.

their opponent, so the disagreement between them can't be classed as a peer disagreement.<sup>225</sup> Moreover, if one has this knowledge, surely the prospect of a disagreement would be mitigated because their knowledge would enable them to pre-empt any dissent and correct it.<sup>226</sup>

Right Reasons delivers what is needed by way of a response to peer disagreement, but it sets unrealistic expectations on would-be peers and doesn't explain why disagreement matters, epistemically, unless one or both peers are *knowingly* mistaken. Might there be a way of vindicating the view? Michael Titelbaum addresses the matter of rational foreknowledge being required:

One can't object to RR on the grounds that it will lead Greg to a conclusion forbidden by his initial evidence; by stipulation the view only applies when he's read that evidence right. But one might ask: How can Greg know that he's the one to whom the view applies – how can he know *he's* the one who got it right?<sup>227</sup>

One might think that Right Reasons is an insufficient position in the peer disagreement discourse because it cannot tell us what view is true in any given disagreement, but as I noted with Equal and Extra Weight, this is a separate matter altogether. Titelbaum notes that Right Reasons is a conditional view, its job is only to describe what one is rationally required or permitted to do when encountering disagreement, having drawn (or not drawn) the rationally required judgment. Right Reasons is not beholden to tell us what beliefs are true and false, neither is it the fault of the view if an agent mistakenly thinks it justifies some belief which leads to further mistaken beliefs. For Titelbaum, the challenge regarding the need for God's eye omniscience comes from anti-Right Reasons advocates presuming that the relation between one's attitude and situation must be either obvious or at least discernible.<sup>228</sup> This is all well and good when circumstances are like the situation he describes, in which the relevant relation between *h* and *E* is discernible, accessible, perhaps even obvious. But it is often far from clear how the evidence points and insofar as there can be peer disagreements, few of them have solutions that are clearly indicated as following from the evidence, let alone known.

---

<sup>225</sup> Of course, whether knowledge that the evidence favours or justifies one's attitude, *p*, is tied to knowledge that *p*, is a separate question, but both surely constitute an epistemic advantage that is inconceivable in peer disagreement cases.

<sup>226</sup> This last point presumes that possession of the rational foreknowledge Right Reasons requires *also* entails that knowledge ensures the sufficient explanatory skills or understanding to dissolve a disagreement. Whether knowledge is a requirement for understanding and/or explanation is a different matter entirely, it should be sufficient that I've raised it here and can move on.

<sup>227</sup> Titelbaum 2015; 286.

<sup>228</sup> Ibid; 287.

Right Reasons is not responsible for giving us the answers about what to believe objectively. Fine. It still presumes a lot to believe that all evidence for every dispute is fully transparent and easy to read, but the opposite view would be to think no evidence is ever clear and we can never tell who is correct in a dispute, and this is just as implausible. In many real disputes we have a mix of clear and unclear answers, with some more obviously pointing to one side or neither side being correct. Either Right Reasons collapses into Extra Weight and leads us to justified dogmatism (resulting from the higher order uncertainty about how evidence points) *or* it is overdemanding and requires implausible foreknowledge for agents to be rational. Regardless, it does not reflect our typical experiences of at least some errors in judgment being clearly identifiable.

### 6.5: Total Evidence

Now I'll consider the Total Evidence View, which is suggested by Thomas Kelly as his solution to the shortcomings of the Equal Weight and Right Reasons views, specifically the spinelessness and lack of self-trust implied by the former and the latter's need for omniscience about rationality. Once again, this view is chiefly about how agents ought to respond when they encounter an epistemic peer about some matter who disagrees with them about that matter. Kelly claims that rationality is not determined by the prior correct assessment of the evidence, but instead it is the sum of *all* the evidence (including the higher order evidence gained from the fact that there is peer disagreement) that is decisive:

In some cases, the first order evidence might be extremely substantial compared to the higher order evidence; in such cases, the former tends to swamp the latter. In other cases, the first order evidence might be quite insubstantial compared to the higher order evidence... In still other cases, the two kinds of evidence might play a more or less equal role in fixing facts about what it is reasonable to believe. So the question of which counts for *more* – peer opinion, or the evidence on which the peers base their opinion? – is not, I think, a good question when it is posed at such a high level of abstraction.<sup>229</sup>

In a similar setup to Right Reasons, two peers disagree over the matter H, drawing their judgments from the first order evidence E (which, in-fact, supports H). Once they acknowledge that they disagree over H, they each gain two additional pieces of evidence: the fact that one believes H is true, and the fact that the other believes H is false. These two pieces of evidence

---

<sup>229</sup> Kelly 2010; 202.

are taken to be “more or less equally strong pieces of evidence that point in opposite directions”<sup>230</sup> and the newer, broader body of evidence, E\*, consists of the original propositional evidence and these new pieces.

One might think that the belief-facts neutralising each other means that E\* makes it just as reasonable to believe H as the original evidence E. However, Kelly contends that E and E\* do not have the same degree of support for H, that the counterbalanced (psychological) evidence makes a greater proportion of E\* support agnosticism than E alone. It may not be enough to defeat E and make it all things considered rational to suspend judgment, but it does push the disagreeing peers at least a little further to neutrality. Ultimately, in this instance (and presumably many others like it) E\* favours H, but not to as great a degree as the original E. If we have reasoned correctly, then great, we come out of the disagreement only needing to be slightly less confident in our belief, otherwise nothing else has changed all that much. But the Total Evidence View faces similar problems to Right Reasons in that we still require some degree of foreknowledge or God’s eye omniscience to discern what we are permitted to believe. Even if we have a strong degree of confidence in our evaluation and reasoning, the fact that someone we regard as an equal or who is our epistemic peer disagrees with us still gives us reason to stop and reconsider. As such, we cannot simply be dogmatic and dismiss any dissenting view outright.

Total Evidence was motivated by Kelly’s concern that endorsing Equal Weight would allow psychological facts to swamp the original propositional evidence that peers possess, letting their subsequent beliefs be dictated by facts about opinion distributions. It succeeds in preventing psychological evidence from swamping propositional evidence, but the implications are comparable to Right Reasons and Extra Weight. Kelly maintains that facts about disagreement only make a small difference between the degrees of support for a hypothesis afforded between two peer’s pre-disagreement evidence and their post-disagreement evidence. A full body of first and higher order evidence will point towards *an* answer but not to the same extent as the first order evidence by itself. For this reason, one might be tempted to think that the upshot of endorsing Total Evidence depends on one’s objective or subjective accuracy, as is the case according to Right Reasons. If this were so, then Total Evidence would permit rational agents to maintain their belief, albeit with a slightly reduced degree of confidence and a slight inclination towards scepticism. However, irrational agents would also be permitted to retain their beliefs with a marginally weaker degree of confidence. In either scenario, it comes to the same result as Extra Weight: dogmatism and marginal shifts in rationality. Now suppose that Total Evidence does not factor in the caveat from Right Reasons (which, of course, it doesn’t).

---

<sup>230</sup> Ibid; 202.

Regardless of whether either disagreeing peer evaluated the evidence rationally beforehand or not, they can retain their belief but must still be slightly less confident.

While Right Reasons has the caveat that peers' post-disagreement rationality depends entirely on pre-disagreement rationality, Total Evidence appears to disregard the matter altogether; it says they can carry on believing as they were as long as they reduce their confidence proportionately. The psychological evidence does not swamp their propositional evidence, but the effect is such that they must acknowledge it and adjust their confidence after establishing disagreement. Agents who are in a peer disagreement do not know how their evidence points, although they likely each have a strong suspicion that they have evaluated the evidence correctly.<sup>231</sup> If Total Evidence is true, it does not alleviate this uncertainty or answer the question of what to believe about (any given)  $P$  or  $H$ , from  $E$  and  $E^*$ . Instead, dogmatism prevails. Whose (if anyone's) judgment is correct is correct remains uncertain and obscure, and each speaker's rationality is only marginally changed. The results of Total Evidence and the results of Extra Weight are no different.

I contend that Total Evidence has it right that we should be sensitive to our fallibility and bear in mind evidence about the distribution of opinions, but it makes the wrong recommendations regarding how to do this. It does not tell us why disagreement is epistemically significant, and proponents are committed to doing no more than stating the obvious: speakers are peers who disagree over some matter, and once they acknowledge that they disagree, they should be sensitive to the fact that they are peers who disagree! The disagreement ends with mutual stubbornness and dogmatism, neither it nor any wider inquiry progresses and neither speaker is any better off except that they both now know their view faces at least some dissent, which they may have suspected beforehand anyway. To be clear, I agree that insofar as there are requirements of rationality, those requirements should include sensitivity to our fallibility as epistemic agents. We are not omniscient beings who know always and everywhere what evidence points to, and there might be evidence that is relevant about which we are unaware. Nor are our circumstances ideal or perfect, in fact a great deal of the time they are not even optimal, and we should acknowledge this when we deliberate and form or revise our beliefs. Evidence or facts of the kind " $S$  believes that  $p$ " are not, I argue, propositional evidence in favour of  $p$ , nor are they

---

<sup>231</sup> There is a possible, but rather absurd situation in which two peers disagree about  $P$  but neither of them believes they are correct in their own belief, or a less absurd case where they each merely have low credences. In that case Total Evidence permits them both to retain their (akratic or weak) beliefs, bringing it more in line with Equal Weight than Extra Weight.

higher order defeating evidence against not- $p$  and they should not be regarded as such. They are evidence concerning the distribution of opinions about  $p$  or facts about  $S$ 's beliefs.<sup>232</sup>

Disagreeing peers (and non-peers) with inconclusive first order evidence and neutral higher order evidence or psychological evidence should therefore suspend judgment, but not in such a way that is indefinite or that abandons the inquiry. Rather, the optimal solution for all parties is to *collaborate* and engage in collective inquiry, taking the time to explain their reasoning and clarify any ambiguities or differences in terminological interpretation. Sometimes our non-psychological evidence will swamp the psychological evidence, as is the case when we are isolated but have access to plenty of well-documented, reputable data regarding a matter. Other times, we will have little propositional evidence but find that plenty of people we trust have given the matter much consideration, possess evidence that we lack, and have a firm opinion on the matter, meaning the psychological evidence dominates.<sup>233</sup> If we are unsure about what to believe, whether it is because of disagreement or not, rather than simply letting psychological evidence swamp by conceding to a majority or deferring to authority beliefs, we should instead seek their support and clarity. One final point is that although most examples of peer disagreement in the literature are described such that it is often indeterminate who is correct or justified and who is not, most real-life disagreements have a clearer answer. Many cases where it is more obvious will likely involve absurd conclusions or obviously faulty reasoning, such as the adapted Extreme Restaurant Check case that David Christensen considers, in which one of the guests calculates their share of the total bill as \$43 and the other concludes it is \$450! This case obviously demonstrates an error in judgment because \$450 is greater than the total bill, so in light of this the diner who calculates a \$43 share is clearly justified in maintaining their view without any additional consideration.<sup>234</sup> However, the motive for Total Evidence is still far from clear in cases like this, for the obvious reason that it licenses absurd conclusions. I'm much more amenable to something like Total Evidence, although I maintain that collective inquiry would likely make the error more salient.

## 6.6: Peer Disagreement Does Not Advance Inquiry

So far, I have examined the primary views within the peer disagreement discourse, which all purport to explain how we rationally ought or should respond to peer disagreement. I have

---

<sup>232</sup> Recall, from Kelly: "it does not follow from the fact that a generally reasonable individual believes  $H$  on the basis of  $E$  that it is reasonable to believe  $H$  on the basis of  $E$ ." 2005; 186.

<sup>233</sup> I will say more on this in Chapter 8 when I consider interactions between speakers who are epistemically asymmetric, between epistemic superiors and inferiors of various kinds.

<sup>234</sup> See Christensen, D. 2007; 199. The steadfastness is warranted, I suspect because of the arithmetic impropriety rather than considerations to do with peerage and disagreement...



argued that they are all *uncompelling*, they each face particular flaws or invite undesirable and problematic implications, as follows:

- 1) The Equal Weight View forces indefinite suspension of judgment and potentially compels otherwise rational agents to be epistemic akratic.
- 2) The Extra Weight View permits a privileging of beliefs which is not always clearly or discernibly evidentially justified and could be relied upon *ad infinitum* by irrational agents.
- 3) The Right Reasons View requires an unrealistic God's eye omniscience with regards to rationality, precluding any prospect of disagreement and rendering one an epistemic superior.
- 4) The Total Evidence View ultimately collapses into Extra Weight. It still requires omniscient foreknowledge of evidential rationality in non-absurd instances of peer disagreement.

While explaining how disagreeing peers rationally ought to respond to their disagreement is one aim of the peer disagreement discourse, I've also alluded to each of these views failing at the other, primary aim. I will now turn to this charge directly. I contend that the discourse is *inconclusive* – it fails to deliver on the chief aim of clarifying what makes disagreement *epistemically significant* or weighty, as well as a further aim of *advancing inquiries*. Each of these views have implications or requirements that stop disagreement, either resulting in mutual dogmatism or all parties in a state of suspended belief. This is not widely or explicitly discussed in the literature, nor is progress in a domain posited as an aim in or of disagreement (insofar as the phenomenon of disagreeing has aims). However, disagreement and wider dialogue are typical features of inquiry, so it would be strange to maintain that peer disagreements are a unique, isolated sort that fully make up an inquiry. That is, our engaging in just about any inquiry does not simply begin and end with disagreement. Consider the following hypothetical:

Imagine two alternative worlds to our own, where no agent possesses knowledge or certainty about the nature or content of truth in any domain. In one such world, there is no consensus or agreement. Everyone possesses differing sets of first- and second-order beliefs, and no two agents share the same combination of ideas or beliefs. In the second world, just the opposite is the case, there is total uniformity of belief about every possible contentious matter and there is no doubt or dissent. Still, no one knows for certain that their beliefs are in fact true or correct.

While these are exaggerated cases, I believe they show, firstly, that disagreement is not as problematic as one might believe and, secondly, that agreement and disagreement are both essential and unavoidable parts of inquiries. There must be some degree of both agreement and disagreement in order for us as epistemic agents to function, and for inquiries to advance and for communities and societies to progress. Without agreement we have nothing to reinforce our confidence as inquirers or judges, whereas without disagreement we risk neglecting our own fallibility and falling into anarchy.

My charge that peer disagreement is inconclusive touches on a few matters, one being the relation between inquiry and disagreement (and, subsequently, the relation between inquiry and peer disagreement). Agents will sometimes disagree over a matter without actively inquiring into it, and while it is possible to have inquiries without disagreement, it would be highly implausible. I regard both points as fairly trivial truths, so I will not explore the relation between inquiry and disagreement (and agreement) at length. The other relevant question is what inquiry aims at, which has been phrased both in terms of what it is that agents and speakers are looking for when they are engaging in inquiry, and what is the natural point at which to cease inquiring – the *inquiry stopper*. Many epistemic goods have been regarded and defended as a natural inquiry stopper, such as knowledge, indefeasible belief, and consensus. Klemens Kappel has argued that inquiry has no formal stopping point because human agents can never be in the epistemic position to rule out any possibility of error. However, he grants that the concept of knowledge fills the natural role of an inquiry stopper, indicating that an inquiry has gone on long enough despite having *pro tanto* reasons to continue it.<sup>235</sup> Similarly, Nathan Ballantyne contends that philosophical and epistemological methodology may clash with the desire to settle questions or inquiries and attain knowledge or beliefs that are justified or undefeated. As long as we have undefeated defeating evidence and inquiries remain open, we may have unsettled doubts about what to believe.<sup>236</sup>

In terms of accounts of the *aim of inquiry*, Christoph Kelp has recently suggested various lighter aims, such as settling “whether” questions, coming to know, to truly believe, or to justifiably believe propositions. He settles on knowledge as the ultimate aim and the signal to close, cease, or end inquiry.<sup>237</sup>

---

<sup>235</sup> See his 2010; 76-8.

<sup>236</sup> Ballantyne, N. 2019; 271-2. It’s worth noting that he then dismisses that conclusion as too simple, noting that while our methods may sometimes lead us to question the value of continued inquiries, continuing could lead us to evidence that defeats our present evidence-defeaters, and that we can reasonably anticipate doxastic openness even after continued inquiry.

<sup>237</sup> See Kelp, C. 2021

Any belief on whether p whatsoever will *ipso facto* be improper if and only if it does not attain the aim of inquiry into whether p. Crucially, it is implausible that there should be a requirement stronger than knowledge such that any belief that doesn't satisfy it will *ipso facto* be improper.<sup>238</sup>

Finally, others still have located the question in broader considerations of how debates and disciplines *advance* and *progress*, with some offering sceptical doubts about certain debates progressing compared to others and citing widespread disagreement as a reason for doubting. I am particularly interested in this question, since some notion of progress will be at least an implicit aim of inquiries in most disciplines or domains. I will argue current responses to peer disagreement fail in this respect because they aren't targeted towards progressing or advancing a dispute in some relevant way when they should be. In some cases, perhaps, once epistemic peers acknowledge that they disagree, they may aim to resolve their disagreement by either establishing that one or all parties are mistaken, by coming to a consensus, or by identifying new difference-making evidence. These are all different ways that a disagreement and a domain can progress, but it has not often been discussed within the literature except to draw sceptical doubts. For example, Hilary Kornblith contends that the likes of the mathematical or scientific communities are far from on par with the philosophical community in terms of historically progressing and resolving issues. He contends that this gives us reason to avoid deferring to community or majority views, and since the philosophical community consists of intelligent and intellectually considerate thinkers who are nonetheless widely divided on several issues. Because philosophers are historically inclined to disagree, there is little reason to have confidence in individual philosophers or in the wider discipline, and we have no choice but to suspend judgment on the majority of philosophical matters.<sup>239</sup>

I am unconvinced by Kornblith's argument, simply because I think he makes an unrealistic comparison between philosophy and the sciences or mathematics, but more importantly I do not think consensus and agreement are the sole measure of progress in a domain. Arguably, we could say that there is progress when one better understands an idea, claim, or theory, and is more competently able to discuss it or participate in inquiries about it.<sup>240</sup> *This* notion of progress is arguably far more common and more easily identified in philosophy

---

<sup>238</sup> Ibid; 18-19.

<sup>239</sup> See Kornblith, H. 2010; 44-5.

<sup>240</sup> This echoes a distinction that some maintain between understanding and knowledge, where knowledge is merely propositional knowledge, whereas understanding involves the grasping of relations between items of information, or the ability to explain those relations to others. See, for example, Jonathan Kvanvig's 2003 (197) and Alison Hills' 2009.

than consensus. For example, for a great deal of the 20<sup>th</sup> century, the concept of *knowledge* was defined and regarded in terms of justified true beliefs. Before that, it was more conventionally or traditionally understood in terms of beliefs with discernible truth-entailing properties.<sup>241</sup> Were it not for the intervention of Edmund Gettier (and Bertrand Russell beforehand), we would not have come to recognise concepts such as epistemic luck, justification, or (obviously) Gettier cases, or at least not in the same manner and timeframe. Their theories concerning knowledge were foundational to a decades-long debate that remains open. Philosophers still have no firm or universal consensus on what knowledge is and is not, but there is at least a common recognition of its deep complexity and of the continual opportunity to theorise, conceptualise, and analyse, which is all indicative of progress without knowledge or consensus. Inquiries can progress without closing, and disagreements keep inquiries from closing. Therefore, regardless of what any individual regards as the aim of philosophy, we need not be entirely sceptical or accept that suspending judgment is all that philosophers are destined for.

The main contender views in the peer disagreement discourse aim to identify what is epistemically significant about disagreement and what is required or rational by way of a reply to peer disagreement. There is no explicit aim of advancing or closing inquiries, and I've already argued that endorsing them simply forces disagreeing peers into either dogmatism or suspension of belief. That serves no desirable purpose other than avoiding further disagreement, and I am hesitant to even regard that as a desirable end, let alone one that can be called 'rational' in any sense. Surely a more likely motive is simply to avoid going around in circles? Even if we think there is little to gain from further inquiry, additional dialogue, or extended disagreement, the alternative certainly provides less reassurance. I think we can confidently say that what motivates most inquiries is (at least in part) an urge or itch to answer a question, resolve a problem, or build on existing considerations that have been settled, which fits with the idea that some progress consists in understanding rather than consensus. I've already noted that Klemens Kappel maintains that inquiry has no natural stopping point, and he is right in a sense because for many inquirers there will be an ever-present urge or a desire to go further.

Knowledge or justified belief may serve as a natural point of closure to many inquiries, and while keeping a disagreement open does not often lead us to knowledge or justified belief, neither do we get them from closing an inquiry prematurely either. *This* is what I contend the replies to peer disagreement would otherwise lead us to. The logical implications of Equal and Extra Weight are, respectively, merely suspending judgment or licensing dogmatism for all. Right Reasons and Total Evidence fare no better because they both require omniscient foreknowledge,

---

<sup>241</sup> See Dutant, J. 2016.

which *would* close an inquiry and possibly prevent any disagreement in the first place but is nonetheless unattainable and unrealistic. Even setting that aside, their bare assertions are, respectively, that whoever had things right in the first place can stick to their original view, or that whoever's view is best supported by the combined evidence can carry on almost as they were. Endorsing either of these views in a disagreement case, whether we are disagreeing with a peer in some domain or inquiry, would fail to contribute anything new or novel to the disagreement or inquiry or clarify any complexities. Disagreements are inevitable in most matters, we cannot separate the two apart, which surely means any response to disagreement must also aid, address, or serve inquiry somehow. The main positions in the peer disagreement are unconvincing and inconclusive; they fail both at their explicit aims *and* the unacknowledged but plausible aim of advancing inquiries. I contend that the epistemology of disagreement needs to reflect on this more.

#### 6.7: Peer Disagreement is *Insignificant*

I now want to explore some other problems with the wider peer disagreement discourse, some which have been discussed in some detail, and others that I consider novel and unacknowledged. So far, I have not specified what precisely I think is epistemically significant about disagreement *simpliciter*. The discourse affords epistemic weight or bearing to some cases of disagreement, but not necessarily to all, which may be apt. After all, why should we think that it matters that an infant says *p* with no consideration or comprehension of what they are saying when we believe that not-*p* on the basis of substantial reflection and understanding of the implications of not-*p*. Signalling the possibility of error is a viable candidate for why disagreement is epistemically significant, indeed it is the most salient option; my intuition is also that disagreement can be a means of attaining or developing several epistemic goods (as indicated in the previous section.) As I've noted already, much of the peer disagreement discourse contends that disagreement's significance only becomes apparent when we abstract or bracket inequalities and differences:

By examining idealized cases of disagreement, we can better isolate the epistemic significance of *disagreement* itself. That is, in idealized disagreements we can neutralize other muddying factors that might cloud the epistemic significance of the evidence about disagreement itself... idealized cases of disagreement can set the 'baseline' for the epistemic significance of disagreement in general. How we should respond to evidence about everyday, real-world disagreement will depend at least in part on how we should

respond to idealized disagreement (even if we never in fact encounter idealized disagreements).<sup>242</sup>

By beginning with ideal cases and paring the circumstances to focus chiefly or purely on the fact that there is disagreement, we (ostensibly) see why that fact matters so much, and what it implies. Consequently, we have our foundation for identifying what response to peer disagreement is more or most rational: Ideal cases have no factors that would otherwise obscure or make a difference to what we ought to do, so we can start with such cases where the compulsion is at its clearest and generalise to typical and ordinary disagreements between non-peers. Or so the thought goes. I am not convinced that excessively focusing on peer disagreements demonstrates what is epistemically significant about disagreement, my intuition is that peer disagreement actually obscures it.

Still, the question remains: what *is* so significant, epistemically, about the fact that one agent disagrees with another, or that one disagrees with some merely possible dissenting attitude and, moreover, is it the case that *all* disagreements are (equally) significant? As we have seen, all the positions in the discourse except for Total Evidence claim that peer disagreement indicates or entails that at least one speaker ought to revise their attitude. N. Gabriel Martin argues that the significance of disagreement is obvious, but not trivial. Much like myself, he maintains that it is missed or overlooked by focusing only on peer disagreement. However, he stresses that it is only *actual* disagreement that is significant, and not merely possible disagreement. I can imagine, for any object, event, hypothesis, idea, or intuition, there likely will be a possible speaker who has a conflicting attitude, view, or conviction to my own, and I can dismiss this as irrelevant because there has been no confirmed dissent or conflict. Actual disagreement is significant because it is a genuine confrontation – we are actually facing an opposing view – and this fact compels us to consider or believe that our dissenter is mistaken. When we encounter an opponent, we have no choice but to acknowledge their dissent and, more importantly, the possibility that they are mistaken. Actual disagreement signals that a speaker has committed an epistemic failure and pre-empts assessment of that speaker.<sup>243</sup>

It does not demonstrate conclusively that the other is absolutely wrong in their belief, thereby rendering any further inquiry superfluous. Instead, it raises the question of the other's wrongness precisely by *indicating* that that is the case.<sup>244</sup>

---

<sup>242</sup> Matheson, J. 2015; 33.

<sup>243</sup> See Martin, N.G. 2019; 294-7.

<sup>244</sup> Ibid; 297. My emphasis.

Disagreement does not signal the end of an inquiry, but the intersubjective possibility of error; it raises the question of how accurate one's interlocutor is. This cannot be bracketed, abstracted, or ignored, unless we are to deny that disagreement is significant altogether. Since peer disagreements are a unique subset of disagreements, they build on its basic significance because they occur between speakers who are equally epistemically positioned and, consequently, are both equally as likely to be correct or mistaken. When we experience peer disagreement, it indicates the possibility that we are mistaken as well as our interlocutor's possible error.

I am sympathetic to this position, especially since it maintains that peer disagreements and non-peer disagreements are epistemically significant for at least somewhat distinct reasons. This stands even if one does not endorse peer disagreements as a possibility or focus only on disagreement cases that are more plausible and realistic. Martin's theory doesn't render peerage and peer disagreement entirely trivial, even if we agree that they are unrealistic, although I am hesitant to say there are enough cases where peerage (factually) holds in virtue of speakers meeting all the conditions. I am far more inclined to say there will be instances where speakers perceive each other as peers and justifiably think of each other as peers, even if they do not know they are. I'll say more on this in the next chapter because peer disagreement's failure to be realistic and its emphasis on idealisation deserves more consideration than a few passing remarks here. Signalling potential error for both or all parties need not occur only in full or explicit peer disagreements, it surely holds for cases of perceived peer disagreements or disagreements with partial or would-be epistemic peers. Arguably, it might well be elicited by some non-peer disagreements as well, such as a disagreement with someone I regard as a well-informed epistemic superior or expert on the matter. I would likely think that, of the two of us, I am more likely to be mistaken and they are more likely to be correct or justified (more on that later.)

Signalling of error need not be the only reason that disagreement is epistemically significant, however. Fernando Broncano-Berrocal and Mona Simion argue that disagreement is a fundamental and inextricable part of inquiries, and that insofar as inquiry aims at something, plausibly so too do disagreements aim at just that same thing.<sup>245</sup> For Broncano-Berrocal and Simion, any norm regarding the response to disagreement that takes knowledge as its central aim presents inquiry as an 'all or nothing' affair. On such a model, a speaker who holds a well justified (but not *known*) belief and disagrees with a speaker who holds an unjustified contrary belief is compelled to withhold their belief because of the disagreement between them. Disagreement is a move in inquiry, and our inquiries aim either directly or indirectly at epistemic

---

<sup>245</sup> See Broncano-Berrocal, F. and Simion, M. 2021; 14654.

goods such as knowledge, understanding, or perhaps mere true belief, which are all encompassed by some semblance of epistemic self-improvement. As such, the response to any disagreement is to adopt the attitude (belief, disbelief, or suspension) with the best set of epistemic properties available, or which best enables the acquisition of epistemic goods, in virtue of disagreement:

That is, what matters is not whether, post-disagreement, my doxastic state enjoys a better set of properties than yours, but rather that, in light of the disagreement, the set of epistemic properties my doxastic state enjoys improves from  $t_1$  (the time after disagreement occurs) to  $t_2$  (the time after conciliation in virtue of the disagreement). If it does, I should conciliate. If it does not, I should hold steadfast.<sup>246</sup>

Our response to disagreement should thus be motivated by the disagreement itself (not for practical or unrelated and arbitrary reasons) and determined by whether our epistemic position is improved by changing our attitude. If we do not stand to improve by conciliating, then we should opt for an approach of damage limitation by holding steadfast and avoiding potential epistemic demotion. I have drawn similar conclusions, just in terms of individual and collective notions of inquiry-related progress, we should respond to disagreement in a way that benefits our own epistemic standing but also stands to advance our inquiry or discipline, where possible. I am confident that this is consistent with Broncano-Berrocal and Simion's approach to disagreement and inquiry.

Ultimately, the peer disagreement discourse is motivated by the thought that bare facts concerning disagreement of particular kinds or in particular circumstances tell us something noteworthy, and that what is so significant can be determined by focusing largely on ideal cases of disagreement. While I believe that most cases of disagreement yield something of note, and furthermore that many agreement cases will also do this, I am not convinced that peer disagreements are the litmus test for this significance. I am sceptical about the pervasiveness of peerage itself, and that they tell us something unique that ordinary disagreements do not. Namely, that there is an equal chance that each speaker is mistaken; this possibility could arise in many other disagreements where the speakers are not peers or equals. Moreover, if disagreement's epistemic significance can be identified by focusing on ideal and rare cases, one might think this gives us reason *not* to focus on *non-peer* disagreements, because they only obscure or hinder it. I contend that conceptualising the epistemology of disagreement in this way is both wrong and misguided. We could (and should) instead ask what reasons we have to think that our

---

<sup>246</sup> Ibid; 14659.



reactions to peer disagreements will translate or reflect our response to non-peer disagreements? We might think there is a stronger pressure to reconsider our judgment when we disagree with an epistemic superior, or a weaker demand when we disagree with an epistemic inferior. The differences and inequalities between dissenting speakers might make it seem as though these cases only differ because the likelihood of error is greater or more apparent. However, there is actually much more going on in cases of disagreement (and agreement) between speakers who are epistemically asymmetric. If at least some kinds of non-peer or not-quite-peer disagreements can bear the same epistemic significance, perhaps it does not matter all that much if we can genuinely be epistemic peers, and thus we can simply set aside the notion except for theoretical considerations. Regardless, focusing only on peer disagreement neglects far more realistic and easily identified cases involving non-peers of various stripes. More importantly, it gives the false impression that peerage is sufficiently common to assume peer agreement and disagreements as defaults, which misleads and misinforms anyone who is unfamiliar with the discourse. Those who are more acquainted with peer disagreement may differ, some certainly acknowledge the use of idealisation and the substantial use of a minor set of cases. I have much more to say on this in the remaining chapters, in which I will continue to offer strong charges against the peer disagreement discourse and literature.

So far, I have argued that the primary positions regarding peer disagreement are *uncompelling* and *inconclusive*. I am unconvinced that any of them offer a compelling account of the rational or rationally required response to peer disagreement, or that they fulfil the aims of the epistemology of disagreement (both those that are explicit and another that I offer and regard as unacknowledged but important.) These point to a failure to explain why peer disagreement is uniquely significant, or why disagreement in general is epistemically weighty or significant. The discourse overlooks a plausible account that can stand independent of the possibility of peer disagreement. The focus is more on the question of response, and on how rational peers are in responding to disagreement. In other words, much has been said and very little has been accomplished. The discourse is drifting in circles rather than moving forward, and I contend that this is primarily due to a misplaced emphasis on peerage. I contend that it would not be unreasonable to say that the “epistemology of disagreement” focuses mostly on unrealistic disagreements between epistemic peers who we are unlikely to encounter or identify, let alone actually disagree with. While I think there is room for peerage as a hypothetical concept or an ideal to aspire to, we should not elevate it to the level it has acquired thus far.

First and foremost, we should recognise and acknowledge that we are fallible, limited beings, and that dissent and disagreement are always a lingering possibility, but also that we do not achieve anything if we acknowledge that and then simply give up. Instead, we should recognise that disagreement is to be expected and look for ways of utilising them and advancing or improving the inquiries in which they are a natural feature. Recall my example of the two worlds, one where there is all agreement and the other where there is only disagreement on all matters. Both are unrealistic but demonstrate that we cannot have one without the other and still expect any semblance of intellectual progress or development, as individuals and in communities and disciplines. We need to look beyond peer disagreements and consider the implications of agreement between peers and both agreement and disagreement between those who are not our peers if we are to identify why either agreement or disagreement are epistemically significant. In the following chapter, I will continue to address the critiques to which I believe peer disagreement is open. I will focus on idealisation – the worry that the concept of peerage has excessive requirements and that the discourse is grounded in assumptions that fundamentally misrepresent and distort reality, and reflect on the need for peerage in general, either as a factual state or as a mere perception.

## Chapter 7: Peer Disagreement is Idealised, Peer Agreement is Overlooked

In the previous chapter, I introduced several important criticisms I regard as overlooked or ignored in the peer disagreement discourse. Firstly, I argued that the main views in the debate, which are all concerned with how epistemic peers ought to respond when they disagree, are unconvincing and fail to meet their aim of showing what is epistemically significant about disagreement (the general phenomenon and specific instances). Secondly, I claimed that the peer disagreement discourse is inconclusive, not only does it fail to achieve the aim of showing why disagreement is epistemically significant, but it also overlooks a further aim of advancing disagreements or wider inquiries. While this is not an explicit aim proposed in the literature, it would be difficult to show how disagreement and inquiry are entirely unrelated or separate; agreement and disagreement are undeniably fundamental and inevitable in most inquiries. The third criticism I offered was that peer disagreement is *insignificant*. The discourse fails to specify or explain what is significant about disagreement, beyond offering paltry explanations of why a peer's opinion or the fact that peers disagree is weighty. More importantly, it overlooks a highly plausible account that is at once simple and compelling without being trivial. Disagreement in general indicates the possibility of a speaker's error in judgment, disagreement with a peer indicates our possible error as well.

Over this chapter and the next, I will consider the other issues with peer disagreement that deserve attention. The first issue is that there is excessive and unnecessary *idealisation* in the peer disagreement discourse, which consists of two different worries. Firstly, most authors assume that speakers genuinely do meet the conditions for peerage and that they encounter others who also meet them often enough. This level of assumption misrepresents real life, misleading and misinforming others all for the sake of an aim which I have already argued it fails to achieve. The second part of the idealisation worry is that authors assume peerage is a matter of fact. The conditions that speakers must meet (or the respects in which they must match each other) are unrealistic and highly difficult to attain. I will argue that if peerage is a factual state, it is the product of good fortune at most, and that we are better off thinking of peerage instead as a matter of perception or regard, wherein a speaker is our peer if and only *iff* we regard them as our peer. Thirdly, I will argue that there is no consideration of agreement between epistemic peers and non-peers, and non-peer disagreement is generally sidelined. This is characteristic of the (so-called) epistemology of disagreement being dominated almost entirely by peer disagreement. However, the criticism concerning idealisation largely rests on the assumption that peerage is a factual matter, so I will address that accordingly and conclude by considering agreement between epistemic peers at the end of this chapter, and agreement and disagreement between non-peers in the next chapter.

### 7.1: Epistemic Peerage is Unattainable

In Chapter 4, I introduced the notion of epistemic peerage by saying that epistemic peers are just those speakers who are equal and identical in their epistemic position or standing, with respect to some matter or domain and body of evidence. That is, they are (or regard each other as) equal in the following respects:

*Evidence* – Evidential equality, relative to the question whether  $p$ , occurs when two agents share mutual evidence or are (otherwise) familiar with the same relevant evidence.<sup>247</sup>

*Cognitive* – Cognitive equality occurs when two agents are equally competent and intelligent and have the same faculties for processing evidence (e.g., perception).<sup>248</sup>

*Dispositional* – Two agents are of equal disposition if they are both equally disposed to respond to evidence in the same way, or in an epistemically appropriate way.<sup>249</sup>

Alternatively, disposition may be conceived in terms of a shared distribution of intellectual virtues and vices, such as honesty and freedom from bias.<sup>250</sup>

*Accuracy* – two agents are equal, relative to the question whether  $p$ , if both are equally likely to be accurate in their respective belief about  $p$ .<sup>251</sup>

Not every author on peerage endorses these conditions as exhaustive or even as relevant. The conditions above were not all endorsed from the outset, many earlier conceptions of peerage limited the notion to one or two conditions. As the discourse on peerage has developed, so too have conceptions of what is required of speakers for them to be epistemic peers. Some authors continue to focus on isolated conditions, whereas others add supplementary conditions. For example, Axel Gelfert contends that:

By extension, we demand of epistemic peers (and, even more so, of experts whom we entrust with policy advice) not only that they be as reliable and well-informed as us, but also that they share, by and large, our commitments as to *what it is important to know*. Epistemic peers should not only get their facts right, but should also agree on which facts it is important to get right.<sup>252</sup>

---

<sup>247</sup> See Kelly, T. 2010; 112. See also Christensen, D. 2009; 757.

It's worth noting that access to evidence need not entail extensive familiarity with it. Two speakers could *just now* have been presented with all their evidence for the first time and still be equals in this sense.

<sup>248</sup> Kelly 2010; 113. Oppy, G. 2010; 185.

<sup>249</sup> See King, N.L. 2012; 251.

<sup>250</sup> See Kelly 2005; 174-5.

<sup>251</sup> See Conee, E. 2009; 313. See also Kelly 2010; 112, and Enoch 2010; 956.

<sup>252</sup> Gelfert 2011; 514.

Other authors have commented on the peerage relation being determined by speakers meeting these conditions, or more generally they have questioned the prevalence of peerage in general. Nathan King focuses on a limited notion of peerage, but notes that peer disagreements themselves have their own additional requirements, and that any genuine, acknowledged instance of peer disagreement requires that speakers believe the following, *with justification*:

- i. that they genuinely disagree with a speaker,
- ii. that both speakers have the same evidence or,
- iii. that they each have immediate access to their evidence
- iv. that they both can articulately and comprehensively communicate their evidence,
- v. that they are both disposed to respond to that evidence rationally,
- vi. that they are each equally cognitively impaired or free from impairment,
- vii. that they each have any relevant background beliefs and skills or competencies, and
- viii. that they each have identical or equally distributed intellectual virtues and vices.<sup>253</sup>

In what follows, I argue that the basic conditions of equality in evidence, cognition, disposition, and accuracy are *unattainable*, and that it is unrealistic to think they are met rarely, let alone often. This leads to further questions over whether peerage is a matter of fact or of perception, whether actual and merely possible disagreements matter equally, and whether we the conditions for peerage are necessary or sufficient. Much of the literature not only acts as though these conditions are attainable, but that it is common for speakers to meet them and encounter peers. I will argue that the issues that arise out of this assumption lend credence to peerage more plausibly being a matter of regard or perception, rather than a fact, and the above conditions for peerage being sufficient, but not necessary.

I am not the only one to contest the assumption that peerage is common, or to argue that peerage is unrealistic or implausible. King considers a very constrained set of conditions for peerage, wherein speakers must have the same relevant evidence, and be equally disposed to respond to that evidence in an epistemically appropriate way. However, he notes that merely satisfying these conditions and disagreeing is not by itself epistemically significant or interesting. The only philosophically interesting cases of purported peer disagreement are those that do not depend on any facts concerning whether speakers are epistemic peers and/or genuinely disagree.

---

<sup>253</sup> King, N.L. 2012; 262-3.

Consequently, they are the cases which are the least likely to be peer disagreements.<sup>254</sup> There are at least three important issues to consider: First, there is a distinction to be made between speakers *thinking* they are peers, and indeed thinking they disagree, and them being so as a matter of fact. Second, is the matter of which conditions must be met to be a peer, and whether we should think these conditions are necessary or sufficient. The third issue is whether merely possible peer disagreements are epistemically significant, or if it is only actual instances that matter and bear on participants' judgments. We need not think of these issues as a hierarchy or listed in order of significance, but they are all important considerations that relate to the matter at hand, and I will discuss each one in turn.

### Fact or Perception?

The first point concerns whether we should maintain a distinction between being peers and regarding each other as peers. If we think that it only matters when speakers *are* peers then a very slim or limited set of disagreements are likely to be peer disagreements, especially as the conditions for peerage widen. If we emphasise the perception of peerage more, then it invites the possibility of errors in perception, since one can think they are in a peer disagreement when they are not, or one can think they are not when they are. King denies that peer disagreement is especially epistemically significant based on two sceptical doubts. He contends that it is not simply that speakers rarely ever meet the conditions for peerage (or that this fact is not enough to dispel the possibility of peer disagreement), but that speakers are rarely able to justifiably believe they are in a peer disagreement. Sharing the same evidence is not an easy condition to meet, no matter how widely or narrowly evidence is construed, and it is hard to acquire good reason to believe that a dissenter has just the same evidence as oneself. More often one will have good reason to deny that a dissenter has the same evidence they have, as well as the same relevant abilities or faculties. For King, the best way to ascertain this is with a dissenter's track record, but he maintains that even these are hard to come by.

Consequently, there are many ways in which we can fail to satisfy the conditions for peerage while genuinely disagreeing, even in cases with topics or subject matters that are not especially complicated. Acknowledged peer disagreements, those in which speakers satisfy the peerage conditions, genuinely disagree, *and* justifiably believe they are in a peer disagreement, are possible only in very rare cases, either with regards to very simple matters or between novices who are equally lacking in evidence and track records of reliability. Such cases are of no interest

---

<sup>254</sup> Neither does it require any familiarity with the concept of epistemic peerage or peer disagreements. A non-philosopher could, in principle, believe they are genuinely disagreeing over a matter with someone they believe has the same evidence and is equally disposed to respond in an epistemically appropriate way.

in the peer disagreement discourse. The sorts of cases which are regarded as interesting (both in the discourse and in general) are unlikely to ever be peer disagreements in reality, at least on King's conception.<sup>255</sup> Simon Kirchin and Steve Pethick also discuss this distinction, noting the difference between there *seeming* to be a disagreement and there genuinely *being* a disagreement, and likewise between speakers *being peers* and *thinking of themselves* as peers and focus on peerage as resting on one speaker's perceiving another as their epistemic peer on the relevant matter.<sup>256</sup> David Enoch extends the distinction further and argues that the focus should *only* be on whether to revise belief given a disagreement between someone we take to be our peer (like Kirchin and Pethick) and someone we *justifiably* take to be our peer.<sup>257</sup>

One might worry that if peerage is a matter of perception rather than a matter of fact, then speakers may be more prone to making mistaken judgments about who is and is not their peer in a domain. However, if the thought is that it is difficult enough for multiple speakers to match each other in each individual respect required for peerage (i.e., the conditions specified above), then we surely rarely ever have reason to think speakers match us in some or all of them. We can grant that judgments can be mistaken, and sometimes someone who we initially perceive as our peer can turn out to be better or worse than our peer and equal. Character and intelligence are often far-removed. Good intellectual character does not ensure competence, and neither does epistemic superiority or authority entail virtuous dispositions. Someone can feign expertise if they are charismatic and convincing enough and those we regard as experts can be self-interested, deceptive, misleading, or otherwise uncaring.<sup>258</sup> One could be as engaged and well-informed as possible about a subject, and it may still not be enough to be able to identify speakers who are reliable, sincere, clued in, and well-intentioned, although one might expect that this ability develops naturally alongside interest, familiarity, and competence regarding a subject. Richard Feldman pre-empts this concern by noting that disclosure of evidence and reasoning is important for ascertaining peerage, with this sort of disclosure being a key indicator. In this respect, Feldman's argument runs contrary to King's, since disclosure would give an indicator of just how much evidence is shared between dissenting speakers. But again, I am sceptical of how many real-life situations would involve speakers sharing all of their evidence openly. It seems that the nature of peer disagreements is that they are either hypothetical or (at best) artificially engineered. For example, juries seem to be the closest possible case to a genuine epistemic

---

<sup>255</sup> See King 2012; 263.

<sup>256</sup> See Kirchin, S. and Pethick, S. (ms)

<sup>257</sup> See Enoch, D. 2010.

<sup>258</sup> See Nathan Ballantyne's 2019 paper on what he calls "epistemic trespassing", where agents make judgments or claims, or profess to be sufficiently or expertly informed about matters which (at best) border on their own areas of expertise and competence.

See also King 2012; 259.



peerage between speakers, and that is because all jurors are screened for biases and competence and then presented with all the same evidence. Regardless, I contend that letting peerage be a matter of perception has fewer and less substantial implications than those that come from maintaining peerage is a factual matter. If epistemic peerage is a matter of fact, then we will inevitably fail to identify any genuine peers.

### Necessary or Sufficient?

For the sake of consideration, suppose that peerage is a matter of fact, and that the conditions are at least met by speakers often enough to warrant consideration. We can now turn to the second question raised by the prospect of peerage being unattainable, which is whether the conditions are necessary or sufficient. That is, whether peerage *requires* equality in the respects listed above, if not others besides them, or if two speakers being peers (merely) ensures or indicates equality in some or all of them. It may be tempting to think that this simply reiterates the question of whether peerage is a matter of fact or a matter of perception, but this would be a slight misconception. After all, peerage might be a matter of fact, in which case it *is* determined by whether speakers meet the conditions for peerage, but those conditions may not be jointly necessary. However, if this is so, and speakers need not meet all of them to qualify as peers, then it becomes *too easy* to explain why there is any given instance of peer disagreement when some condition is not met, namely, *because* not all of the conditions have been met.<sup>259</sup> Recall that the reason peer disagreement is presented as the prototypical epistemically significant disagreement case is because all the difference-making factors are bracketed, so as not to provide an easy explanation for disagreement. This suggests that peerage is either a matter of perception, or an all-or-nothing factive state. Otherwise, all we can say is that asymmetries between speakers in one of those respects is a clear sign that the speakers fall short of peerage, and that their disagreement is not, or not quite, a peer disagreement. Regardless, the peerage conditions must be either jointly necessary or not necessary at all. If they are jointly necessary then, as we have established, they are not impossible to satisfy but they are so difficult to meet simultaneously that they fall just short of impossible. King has suggested that even on a constrained account of peerage, speakers do not typically match each other or meet the conditions, so we have reason to regard it implausible that the conditions for peerage are jointly necessary.

The conditions for peerage are sufficient, then; speakers are not required to meet them, either individually or together, but at least some of them are brought about in virtue of speakers being epistemic peers. This matter has not been given much consideration because much of the

---

<sup>259</sup> This also suggests, paradoxically, that peer disagreement can only arise when there is some asymmetry: explaining why there is peer disagreement becomes more difficult when all the conditions are met, which is counterintuitive.

discourse assumes without argument that epistemic peerage is both possible *and* prevalent. One might still worry that because it is harder to identify equals in some of the conditions than others, that peerage is in fact an all or nothing relation but again, as we have seen, treating peerage as an absolute relation actually limits the likelihood of peer disagreements. We can certainly imagine circumstances in which we disagree with a speaker who falls just short of matching us in some or all these regards, and others where there are far more substantial asymmetries between speakers. I contend that we can point to the development of peerage as a concept and the gradual addition of further conditions as lending credence to regarding the peerage conditions as sufficient, and indeed to treating peerage as a matter of perception rather than one of fact. Ultimately, if the peerage conditions are necessary, then King is right to point out that we rarely ever meet the few conditions for peerage he discusses, let alone the additional conditions that have arisen in the literature. If they are merely sufficient, then peerage looks more like a perception or a matter of one speaker's regard of another, and King is right to point out that we rarely have the requisite reasons to believe that a dissenter is our peer. The situation is bleak and inspires little confidence either way, but things are at least more realistic if we regard peerage conditions as sufficient, and even more so as perceptual.<sup>260</sup>

### Possible or Actual?

The third matter that arises out of questioning the attainability of peerage and the prevalence of peer disagreements is whether all peer disagreement (and, perhaps, non-peer disagreement) is epistemically significant or if it is only actual, genuine peer disagreement that matters. Opinions on this are divided:

Whether we find the possibility of disagreement intellectually threatening, I suggest, will and should ultimately depend on our considered judgements about *how rational* the merely possible dissenters might be in so dissenting... But if this is correct, then the extent to which merely possible dissent should be seen as intellectually threatening effectively reduces to questions about the strength of the reasons that might be put forward on behalf of such dissent.<sup>261</sup>

---

<sup>260</sup> One point I have not considered and cannot recall seeing within the discourse but is relevant here, is whether each condition for peerage ever varies in its significance or prominence between matters or domains. I contend that suggesting each peerage condition can vary in significance creates the misleading implication that the conditions are necessary, and leaves room for too simple an explanation of peer disagreements.

<sup>261</sup> Kelly 2005; 181.

When I find that others disagree with me on a certain question, this gives me, *ceteris paribus*, reason to be less confident than I was that I am right... disagreement from epistemic peers gives one reason to suspend judgment entirely on the disputed question. It is clear, moreover, that the mere possibility that someone might disagree with me does not have the same epistemic significance... since there might always be people who disagree with one on any question at all, treating merely possible disagreement as on a par with actual disagreement would result in total skepticism.<sup>262</sup>

For the fact of disagreement to take on a compelling degree of urgency, it must be acknowledged by at least one party to the dispute. If both peers were blissfully unaware of their doxastic differences, the mere counterfactual observation that they *would* disagree with one another *if* quizzed on the point in question, can hardly be expected to have an effect on their beliefs and credences.<sup>263</sup>

Generally, the consensus is that actual disagreement creates stronger or more pressing reasons to doubt or reconsider one's attitude or judgment than the mere, distant possibility of disagreement, or that merely possible disagreements simply aren't weighty at all. As well as the authors quoted above, David Christensen argues that possible disagreement is merely a consequence of being non-ideal thinkers, whereas actual disagreement realises the possibility that we have made a genuine error.<sup>264</sup> Brandon Carey argues that possible disagreement leads to one of two claims and that neither provides the same defeating evidence as an actual case of disagreement: *either*, of something that actually exists, I could be in a disagreement with that thing, *or* there could be a thing that I would be in a disagreement with.<sup>265</sup>

Some take these sorts of arguments as contrasting or conflicting with claims like those made by Thomas Kelly. However, it is important to note that Kelly is not saying *all* merely possible disagreements are as worrisome as all actual disagreements. What he *is* saying is that, in some cases, actual disagreement is no more significant than possible ones, and that the reality of disagreement may be a contingent matter determined by factors that are easily dismissible as irrelevant. How much the possibility of a dissenting judgment bears on our own is determined by how rational we believe a would-be dissenter is, so if we believe that the contrary view is at least partway rational then we should acknowledge its weight. The possibility of disagreement

---

<sup>262</sup> Kornblith, H. 2010; 34.

<sup>263</sup> Gelfert 2011; 511.

<sup>264</sup> See Christensen, D. 2007; 208.

<sup>265</sup> See Carey, B. 2011; 379.

with speakers who are clearly irrational or a-rational will therefore offer no reason to doubt our view.<sup>266</sup>

The question of whether merely possible disagreement matters any more than actual disagreement derives from the more general question of what, if anything, matters about disagreement in general. I addressed this matter in the previous chapter and considered N. Gabriel Martin's claim that disagreement is significant as a signal of a speaker's possible error and, in the case of peer disagreement, our own possible error as well. Moreover, we have seen that Nathan King maintains that peer disagreement only matters when it is acknowledged, that the mere fact that there is or could be some possible peer somewhere who happens to disagree, is irrelevant. Along similar lines, Axel Gelfert contends that disagreements are "at least *prima facie*, occasions to reflect on the fallibility of the methods by which we acquire beliefs", and that only acknowledged peer disagreements can bear epistemic weight. Peers who are unaware of dissent, or who merely suspect the counterfactual possibility of would-be disagreement, are under no compulsion or requirement to change their beliefs. For Gelfert, therefore, ascertaining that a speaker is one's epistemic peer has at least some priority over ascertaining how to respond when disagreeing with a peer and whether such disagreement compels a change of attitude.<sup>267</sup>

Given my scepticism about the possibility of peerage, it is tempting to think that I would be dismissive of Gelfert's position. I grant that, if I regarded peerage as a factual possibility, I would agree that establishing it takes precedence over deciding how to respond to a peer disagreement. However, Gelfert claims that we can disregard merely possible disagreements and actual disagreements with epistemic inferiors, which I find utterly unconvincing. The possibility or presence of peerage should supplement the baseline significance of both disagreement *and* agreement. To say that possible peer disagreements and actual asymmetric disagreements are of no significance would be incredibly presumptive and almost epistemically hubristic. For Gelfert, disagreeing with an epistemic inferior is not significant, but this focuses *only* on the significance for the speaker occupying the position of the superior, when such disagreement will undeniably matter for anyone occupying the position of the epistemic inferior. I will say more on this in the next chapter when I consider the importance and implications of both agreement and disagreement with and between epistemically asymmetric agents.

I have argued that epistemic peerage is *unattainable*, that the conditions for it are impossible, or at least highly difficult to meet, and that it would be incredibly rare to encounter a genuine epistemic peer, let alone one with whom we disagree on some matter. I have argued that

---

<sup>266</sup> See Kelly, T. 2005; 181.

<sup>267</sup> See Gelfert 2011; 510-2.

there are three facets to this unattainability, or rather that the challenge raises three issues and implications. These are that it is not clear whether the conditions for peerage ought to be regarded as necessary or sufficient, that it reiterates comparisons between the significance of actual vs merely possible disagreements, with respect to their epistemic significance, and it raises a distinction between peerage as fact and peerage as a regard or perception. Peerage as a fact is difficult to ascertain, but we need not dismiss the importance of regarding or perceiving someone as our peer or equal, even if such perceptual judgments can be mistaken or misinformed in the same way that we might mistakenly judge that someone is our epistemic inferior or superior.<sup>268</sup> When we regard or think about others in the context of inquiries, dialogues, or disagreements, we attend to certain features which include things like their intelligence and evidence, intellectual character, and underlying biases. Sometimes we will make mistakes in judging each other some way or another, but other times we will judge each other accurately because we attend to the right features.

Regardless, we are far more likely to come across an epistemic superior who is better equipped and better off regarding the matter at hand than us, or an epistemic inferior who is worse off, such as a layperson who will likely be less interested or equipped. The asymmetries that are relevant here are more numerous, meaning there are many ways in which a speaker could be epistemically inferior or superior, and I will address the matter directly in the next chapter. For now, it should be enough to posit that the conditions for peerage are not necessary or fundamental requirements, and that they are not as easily met as they seem, at least not collectively and simultaneously for the same topic. Embracing this need not commit us to claiming that merely possible disagreements are always just as significant as actual disagreements, and I agree that, with respect to at least some matters, the mere possibility of disagreement should not worry us. However, we should always be wise to our fallibility; to assume that there are no dissenters whatsoever to some view, attitude, conviction, or opinion that we hold would be fundamentally dogmatic, naïve, and arrogant. The fact that dissent always lingers on the horizon should not push us towards scepticism or doxastic hesitancy, but neither should we take the faraway possibility for granted by ignoring it and living in blissful ignorance. We will be inevitably confronted with actual disagreements at some point, and it will give us reason to reconsider our judgments and attitudes.

Setting this aside, peerage being unattainable only scratches the surface. There is a greater concern that it signals, namely that regarding the degree of *idealisation* on which the peer

---

<sup>268</sup> See Ballantyne 2019. Again, Ballantyne describes epistemic trespassers as agents who make judgments or claims or who profess to be sufficiently or expertly informed about matters which (at best) border on their own areas of expertise and competence.

disagreement discourse relies. The fact that peerage relies on unrealistic conditions is not enough of a problem, many authors assume without argument that these conditions are met often enough and that we can clearly identify our peers. I now turn to this assumption, and the various problematic implications it raises. This will lead naturally to a more pressing concern that the peer disagreement discourse relies too heavily on a specific form of idealisation that goes beyond the mere altering of models to make sense of an abstract but plausible phenomenon. This idealisation not only distorts reality and risks placing unrealistic expectations on human agents (e.g., to be equally rationally considerate, epistemically virtuous, and intellectually informed and capable), but it does so to achieve the misguided aim of showcasing disagreement's epistemic import which, as I've argued, is both fruitless and futile.

### 7.2: The What, How, and Why of Idealisation

I have discussed the unattainability of the conditions for epistemic peerage and argued that we ought to think of peerage more as a matter of regard or perception, and of the peerage conditions as merely sufficient, rather than a necessary requirement that determines peerage as a fact. When we consider peerage and peer disagreement, we should be concerned with what it is rational for one speaker, A, to do in light of their disagreement with another speaker, B, who A regards as their epistemic peer. Much of the peer disagreement literature concurs with this thought but allows it to dominate and draw focus away from the question of why disagreement is epistemically significant. Moreover, much of the foundational literature treats peerage as a very real, discernible relation, despite having nearly impossible conditions for it. This assumption sets a precedent for idealisation within the peer disagreement discourse. Authors present peerage as a tweaked model of agents to try to make better sense of disagreement. While they might not assume that all (or enough) speakers are capable of meeting the conditions for peerage, the implication is that peerage is common-enough for peer disagreement to stand at the forefront of the epistemology of disagreement. As we will see, the discourse engages in a typical instance of idealisation presenting the ideal case as a benchmark for better understanding the real life, messy cases. In reality, humans are limited, fallible creatures, with varying capacities, dispositions, and resources that are neither infinite nor equally distributed.<sup>269</sup>

Let me say a little more about what idealisation refers to, both generally and specifically with regard to peer disagreement. Broadly speaking, idealisation describes any kind of practice that involves exaggerating or falsifying information intentionally so that it misrepresents how

---

<sup>269</sup> My concern, then, is that focusing only or primarily on ideal cases will not get us very far in learning about the significance of non-ideal cases, but also that focusing only on non-ideal, messy cases, will only prove to be a non-starter.

things are. The phenomenon is typically discussed by philosophers of science, and there are two kinds of idealisation that usually get considered: *Aristotelian* idealisation strips away features of a model or a system that are deemed inessential or not of interest, whereas *Galilean* idealisation occurs when features of a model are deliberately distorted or when a model is knowingly based on false assumptions. There is also a third kind – *Wittgensteinian* idealisation – that is not explicitly discussed by philosophers of science but nonetheless picks out a familiar phenomenon because it refers to having a strict limitation of examples and cases.<sup>270</sup> Of these three, the first two are often done to simplify models or make them more relevant and fitting, since they generally concern objects, events, notions, or phenomena that are deeply complicated. Angela Potochnik describes this kind of idealisation as:

assumptions made without regard for whether they are true and often with full knowledge that they are false... we often resort to lying a little bit: we artificially simplify the parts of accounts that we are not interested in to improve our access in a variety of ways to the parts we are interested in.<sup>271</sup>

Potochnik distinguishes idealisations from *idealised representations*, which refers to representations or models that incorporate idealisations. The value of idealised representations comes from the fact that they aid our understanding of the world and help us to more easily make sense of it. But the idealisations themselves (i.e., the false assumptions or misinformation) are perhaps more contentious. After all, we would otherwise resent and condemn misinforming others as deceptive and dishonest, but if they are a means to a desirable and otherwise unattainable end and epistemic good, it generates something of a dilemma. There are familiar examples of this that we would likely dismiss as unproblematic – adults might use simple explanations or half-truths to educate or correct children (e.g., “You’ll only get big and strong by eating your vegetables!”) It is only as children mature and become able to grasp the reality that adults will give them the full story, and arguably this is comparable to the kind of idealisation that Potochnik describes:

Understanding has a dual nature: it is both an epistemic achievement and a cognitive state. Understanding thus must achieve both the proper relationship to the world, to the object of understanding, and the proper relationship to the cognitive agent, to the subject of understanding...

---

<sup>270</sup> See Kinzel, K. and Kusch, M. 2018; 49-50.

<sup>271</sup> Potochnik, A. 2020; 934-5.

If grasping a causal pattern can constitute understanding of a phenomenon, then in the right circumstances, idealisations – false posits – can directly contribute to scientific understanding and, thus, are of direct epistemic value.<sup>272</sup>

Potochnik offers an account of epistemic acceptability, on which a claim is unproblematic *iff* its divergence from the truth is not significant, bearing in mind its role in the overall representation and the purpose for which that representation is used. The thought is that different specific aims of understanding can motivate or warrant different kinds of idealisation.<sup>273</sup>

However, when I say that the peer disagreement discourse is heavily idealised or relies heavily on idealisation, this is not quite what I have in mind. The phenomenon of peerage is not incomprehensible, or difficult to grasp or hypothesize, so it is not as though deliberate misrepresentation is essential for explaining why disagreement matters in simple terms. Nonetheless, the discourse operates as though peerage is common enough or relies sufficiently on hypothetical cases of disagreement between ideal agents or within ideal circumstances. It is, I contend, a false assumption to think that peerage truly is common.<sup>274</sup> Moreover, this leads to a fourth sense of idealisation that has recently arisen out of considerations from moral and political philosophy, and to a lesser extent discussions of epistemic injustice. This sense of idealisation is discussed most prominently by Charles Mills who claims that it occurs when there is a substantial, and perhaps deliberate, gap between what agents or groups actually have, do, or experience and what is imagined or hypothesized of them. Among other things, this could refer to our capacities, resources, experiences, history, or culture, but if a theory exaggerates or substantially misrepresents them, then it is idealisation in this sense that is occurring.<sup>275</sup> As Mills describes it:

Factually, idealization involves the attribution to the agents (as conceived of in the theory) of human capacities significantly deviant from the norm (for example, their degrees of rationality, self-knowledge, ability to make interpersonal cardinal utility comparisons, and so forth). Morally, idealization involves the modelling of what people should be like (character), how they should treat each other (right and good actions) and how society should be structured in its basic institutions (justice)...

---

<sup>272</sup> Ibid; 936.

<sup>273</sup> Ibid; 937-8.

<sup>274</sup> I have in mind here peers in the factual, literal sense – recall that I am sceptical that speakers tend to meet *all* the conditions for peerage. That is what I consider to be the relevant false assumption.

<sup>275</sup> One way of aiding the distinction between this kind of idealisation and the three other kinds is that one is idealisation about humanity and the other is about the world.



Now what distinguishes ideal theory is not merely the *use* of ideals, since obviously non-ideal theory can and will use ideals (certainly it will appeal to the *moral* ideals, if it may be more dubious about the value of invoking idealised human capacities). What distinguishes ideal theory is the reliance on idealization to the exclusion, or at least marginalization, of the actual.<sup>276</sup>

The thought is that certain moral or political theories will overlook or obscure the systematic marginalisation or oppression of certain communities and generate representations that are thoroughly unrealistic. These representations might be of those communities, or of human experience in general; sometimes it can be intentional, other times it will be accidental and unforeseen.

While idealisation may be unintended in some cases, it might nonetheless reflect historical or socio-cultural attitudes, or it might be motivated from similar reasons to the earlier kinds of idealisation and be used to simplify complicated or alien phenomena. Again, I am unconvinced that this sufficiently explains why the peer disagreement discourse is so idealised, partly because the discourse is relatively recent, but also because I do not believe that the idealisation in peer disagreement is maliciously intended. The prevailing thought is that by focusing on highly ideal instances of disagreement, any factors that would otherwise interfere or obscure what is epistemically significant about disagreement are stripped away. Some authors are open about this and their deliberate discussion of rare and ideal instances of disagreement, such as Jon Matheson:

S1 and S2 are in an *idealized disagreement* about p at t just in case:

- i. S1 and S2 genuinely disagree about p at t,
- ii. S1 and S2 are qualified [well-informed] epistemic peers about p at t,
- iii. S1 and S2 each have access to their own evidence, processing of it, or resultant attitudes relevant to p at t, and this access is no better or worse than their access to the other individual's evidence, processing of it, or resultant attitudes relevant to p at t,
- iv. At t, S1 and S2 are not aware of the attitudes of any 'third-parties' regarding p.<sup>277</sup>

Or Jennifer Lackey:

---

<sup>276</sup> Mills, C. 2017; 75.

<sup>277</sup> Matheson, J. 2015; 29.

*Idealized disagreement:* A and B disagree in an idealized sense if and only if, relative to the question whether  $p$ , (1) A and B are aware that they hold differing doxastic attitudes, (2) prior to recognizing that this is so, A and B take themselves to be epistemic peers with respect to this question, and (3) A and B are epistemic peers.<sup>278</sup>

Perhaps it is a common trait among peer disagreement authors that they implicitly believe their idealisation commits no harm, especially in the earliest discussions wherein peerage only requires equality in intelligence and intellectual virtues or dispositions. Another possibility is that early accounts of peerage were simply designed to fit around the arguments that were being offered about how to respond to peer disagreements; many authors take themselves to be offering the “intuitive” response, after all.

### 7.3: Peer Disagreement is Idealised

Further reflection on what motivates the idealisation in peer disagreement will not shed light on why it is an issue, so I will now consider some examples from the peer disagreement literature in which authors acknowledge idealisation or openly critique it. I have already noted that I regard the conditions for peerage as unrealistic and unattainable. However, what I believe to be potentially problematic is the risk that people will treat peerage as a live possibility, a frequent reality *and* a non-negotiable element of core interest. Most authors in the discourse openly acknowledge that they are assuming ideal circumstances as a background condition or feature of the discourse. Some others maintain that only a few elements or aspects are idealised, it is only more recently that authors have discussed idealisation openly and offered criticisms.<sup>279</sup> Ultimately, if we assume that peerage is a common relation, one may infer that establishing the epistemic worth of their opinions is only possible if they can identify epistemic peers. They may even feel unduly pressured to meet conditions that are impossible even for those who are in the epistemically best position. The fact that agreement and disagreement between non-peers receive comparatively less attention and peer agreement none at all may only worsen this lack of clarity. Peer disagreement is presented as a hypothetical consideration, but one that is given a central role, but I am not convinced that this is entirely the best approach. However, as I have already noted, I fear that focusing exclusively on messy, real-world cases will not help us fare much better either.

---

<sup>278</sup> Lackey, J. 2010; 302-3.

<sup>279</sup> See, for example, Tim Kenyon's 2020. Kenyon notes that peerage and peer disagreement could be used as a way of comparing the epistemic standing of speakers, but that it is rarely ever used in this way. He argues that equal epistemic standing and intellectual character is often assumed as a background condition for the phenomenon of interest but is generally set aside or dismissed after being initially flagged or noted.

Within the peer disagreement literature, discussion of idealisation tends to fall into one of three categories: some authors, such as Thomas Kelly, maintain that only one or two conditions or features of peerage and peer disagreement are ideal. Kelly argues that the only idealisation in peer disagreement is the assumption that peers arrive at their opinions and attitudes independently. This assumption maximises the force of a peer's opinion compared to one's shared first order evidence, and thus generates a stronger pressure to heed the view of one's peers. If opinions are formed with less independence, then they have a far weaker force, so the fact that a small number of epistemic peers determine that  $p$  independently from each other counts for much more than the fact that lots of peers concluded that not- $p$  collectively.<sup>280</sup> The second approach, which I regard as the most common, is to acknowledge idealisation or ideal circumstances as a general feature but not regard it as problematic, instead considering it fundamental or essential for the purposes of identifying disagreement's epistemic significance. For example, Ted Warfield and Richard Feldman:

In the stipulative sense of “peer” introduced, peers literally share all evidence and are equals with respect to their abilities and dispositions relevant to interpreting that evidence. Of course, in actual cases, *there will rarely, if ever, be exact equality of evidence and abilities*. This leaves open questions about how exactly conclusions drawn about the idealized examples will extend to real-world cases of disagreement.<sup>281</sup>

In this case, the idealisation is openly stated (albeit limited only to a few peerage conditions), but what to do about it and why it should carry on is left completely unaddressed.

There are other examples of this as well. Jennifer Lackey makes a careful and explicit distinction between idealised and everyday or ordinary disagreements. Disagreements are idealised if and only *iff* it meets the following three conditions:

- 1) that speakers are *aware* that they hold different beliefs or doxastic attitudes,
- 2) that they *take themselves* to be epistemic peers, and
- 3) that they *are*, in fact, epistemic peers.

In other words, ideal cases of disagreement occur between speakers who truly or correctly regard each other as peers. In contrast, ordinary disagreements are those that occur between speakers who take themselves to (roughly) be peers, but it lacks the third requirement that

---

<sup>280</sup> See Kelly 2010; 205.

<sup>281</sup> Feldman, R. and Warfield, T. 2010; 2. My emphasis.

speakers actually meet the conditions for peerage or genuinely be epistemic peers.<sup>282</sup> Lackey considers both these kinds of disagreement to be discussed in the literature, but inconsistently and rarely ever explicitly or using idealisation to identify or pick out the notion she has in mind. Lackey's overall claim is that there are at least some cases of disagreement that cannot and should not be idealised, because doing so either makes the disagreement inexplicable or disconnected from real life, familiar cases of disagreement, including the kinds that initially motivated the epistemology of disagreement.<sup>283</sup>

Lackey notes the idealisation, but not any great issue or fault with it, nor does she try to steer the interest in ordinary cases over and above ideal disagreements. Like many other authors, her aim is to try to make sense of the response question and show that neither conformity nor dogmatism explain why disagreement is epistemically significant, but that both deliver “intuitively correct” results.<sup>284</sup> She recognises that there are fundamental asymmetries between speakers, but she contends that we can only assess the epistemic significance of disagreement when we abstract these differences away, or else disagreement could be attributed to them.<sup>285</sup> Perceptual disagreements are the least plausibly idealised, it would be difficult to maintain speakers are perceptual equals, so such disagreements can only really be explained by someone's error or malfunctioning. Likewise, most arithmetical disputes such as the Restaurant Check are almost certainly due to errors on at least one side, so it is far less plausible that they are idealised.<sup>286</sup>

Lackey's argument is that regardless of whether a case is idealised or ordinary, it is not clear that conforming to an opponent's view or compromising and suspending works for every instance. But, like myself, she is adamant that remaining steadfast or dogmatic is no better or preferable. We will always have certain disparities or privileges in information or evidence, including personal information about our own cognitive and mental well-being that we won't have about anyone else. As such, if you and I disagree about the location of a building, for example, even in ideal circumstances our disagreement will give us each reason to suspect that the other has suffered some cognitive malfunction or error, and thus warrants us both reconsidering whether we are epistemic peers.<sup>287</sup> Lackey acknowledges that there is some idealisation in peer disagreement, but she does not contend that it is as prevalent as I do and she does not believe it is problematic at all.

---

<sup>282</sup> Lackey, J. 2010; 303-4.

<sup>283</sup> Ibid; 314.

<sup>284</sup> Ibid; 304-5.

<sup>285</sup> Ibid; 298.

<sup>286</sup> Perhaps they do not motivate conformity because they are relatively mundane or low-stakes cases. I'm not sure that this follows, but I won't argue for this here. I'm more convinced by Lackey's conclusion that such cases can't often be idealised, and that they seem to make it reasonable to cease considering a dissenter my peer. Alternatively, it might be that such cases can be ideal, but the disagreement is plainly the result of one or all parties' arithmetic error.

<sup>287</sup> See Lackey 2010; 309-10.

Finally, the third, and more recent approach to discussing idealisation in peer disagreement is to openly critique the reliance on it. We have seen one example already from Nathan King's observation that the likelihood of meeting the conditions for peerage are low, and that there is practically no chance of being in a peer disagreement. King notes that peer disagreement cases are presented as epistemically significant in virtue of the information they yield, but argues that because we rarely enter into those cases, so we rarely get the relevant information.<sup>288</sup> This is not so much a point about idealisation, it is more about the lack of realism in the discourse; however, it suffices as a cursory example, and other discussions are offered by authors who are more explicitly critical, or who recognise more examples of idealisation than just the conditions for peerage. Thomas Grundmann concurs that much of the peer disagreement discourse is idealised and cites the Restaurant Check example as both prominent and clearly idealised, in the following respects:

1. There are no semantic intricacies or verbal miscommunications, both participants are entirely clear about what they are disagreeing over, and there is no opportunity for (semantic or doxastic) relativism.
2. No permissiveness of evidence, there is only one correct and rational answer to what everyone owes, and either both participants are mistaken or one of them is.
3. The participants are justified in believing that they are epistemic peers – most messy, real-world cases aren't such that participants have good reason to think of each other as peers.<sup>289</sup>

Simon Kirchin and Steve Pethick also note further assumptions within the peer disagreement discourse. Firstly, it is often assumed that there is a definitive answer to most examples in the literature, even if it is unclear. Regardless of whether the answer has already been determined or has yet to be revealed, one is reasonable provided that one accepts a certain answer as definitive.<sup>290</sup>

Second, agreement on all the relevant terms is assumed as necessary for peers to disagree. That is, all the relevant terms of the disagreement are assumed to be (implicitly or explicitly) shared or clarified. There are no semantic differences or terminological ambiguities because the central terms are, at some point or another, clearly and sharply defined. Thirdly,

---

<sup>288</sup> See King, N. 2012; 267. King then notes that disagreement typically does make us aware of the higher-level possibility that our circumstances are less than ideal, with respect to our evidence or disposition.

<sup>289</sup> See Grundmann, T. 2019; 130.

<sup>290</sup> Kirchin, S. and Pethick, S. (ms); 8.

speakers in most examples of peer disagreement also have prior clarity or agreement on what features and evidence are relevant to their disagreement. The diners in the Restaurant Check case are not only clear on the meaning of terms like ‘diner’ and ‘owe’, but on other features such as ‘agreeing to pay’ and numerical concepts.<sup>291</sup> If peers were to disagree on what evidence is relevant to their dispute, or if they were to have fundamentally different understandings of what a key term refers to, then what appears to be a peer disagreement becomes either a higher order disagreement about evidence, or a merely verbal dispute. The same goes for the features we each regard as constitutive for the situation; some disagreement may be acceptable, but *too much* difference over which features are considered relevant may mean that the speakers simply won’t regard each other as peers.<sup>292</sup>

Idealisation in the peer disagreement discourse generates some worrying implications that are rarely considered or discussed. They are sometimes alluded to or underemphasised before being summarily dismissed or promptly explained away, other times it is merely alluded to amidst references to there being a scarcity of peers. Typically, when idealisation *is* discussed, there is no follow-up comparison to non-ideal instances, no move to generalise to more realistic settings, and little consideration of disagreement with non-peers, nor agreement whatsoever. More often, however, the idealisation is overlooked entirely. Still, one might wonder, what’s the big deal, why does this idealisation matter so much, and what might some of the uncritical authors say in response? I turn to these questions next.

#### 7.4: The Problem of Idealisation in Peer Disagreement

Fundamentally, the peer disagreement discourse does not simply begin with a mildly distorted picture and add in further details to make it realistic. It presumes (both implicitly and explicitly) the possibility that some agents can and will be epistemic peers and will encounter or disagree with peers with regards to various matters, that speakers can be equals in various respects, perhaps with little effort. But there is very little explanation within the discourse about how this is supposed to happen, given that humans are fallible beings with limited capacities, unpredictable dispositions and evidence that varies wildly between domains and groups. In other words, the discourse either misrepresents human agency as better epistemically positioned overall than it actually is, or it glosses over the real-life limitations and obscures them by presenting common examples in an unrealistic way. This is justified on the grounds that it is the only way we can identify and understand why disagreement matters. However, the arguments offered that purport to explain why disagreement matters are poor and unconvincing, and there is little

---

<sup>291</sup> Ibid; 9-10.

<sup>292</sup> Ibid; 13.

discussion of the transition from ideal to non-ideal cases such as those between various kinds of epistemic non-peers, nor whether agreement matters in any comparable way.

To be clear, I do not think we should abandon peer disagreement altogether, since purported cases still make for interesting examples that are worth considering when they are clearly *hypothetical* considerations. If, however, we proceed on an assumption that literal, factual peers exist and that peer disagreements are pervasive, we rely too much on an unrealistic picture and potentially demand too much of those who take the discourse sufficiently seriously. The discourse is gradually coming to realise and acknowledge that greater attention should be given to more familiar and realistic disagreement cases. Peerage is an ideal concept, but we need not regard it as trivial, I argue. We should aspire and want to seek out speakers who are our epistemic equals with whom we can have productive conversations and spirited debates; it should motivate us towards greater intellectual virtuosity or expertise, even if peerage is unattainable. As Lackey puts it:

No conversation is entirely free of differences in the distribution of epistemic goods, status, power, psychology, cultural expectations, practical constraints, or some combination thereof. Speaking up against others almost always involves a calculation – whether conscious or not – that is based on one’s position and the costs and benefits of dissent on *this* topic at *this* time with *this* conversational participant.<sup>293</sup>

Or, from John Hardwig:

... each researcher is forced to acknowledge the extent to which his own work rests on the work of others – work which he has not and could not (if only for reasons of time and expense) verify for himself. Thus in very many cases *within* the pursuit of knowledge, there is clearly a complex network of appeals to the authority of various experts, and the resulting knowledge could not have been achieved by any one person.<sup>294</sup>

Neither of these claims are specifically about peerage or peer disagreement, they are observations about our epistemic asymmetries, inequalities, and interdependence.

---

<sup>293</sup> Lackey, J. 2018; 90. My emphasis. Lackey is here specifically talking about pressures to remain silent and to suppress one’s dissent, with the background consideration that objecting is an imperfect duty. I will address this discussion and the duty to object later on this thesis.

<sup>294</sup> Hardwig, J. 1985; 348.

There is no reason for the epistemology of disagreement not to be sensitive to these differences, and while it is misleading to presume that peerage is possible, there is nothing wrong with instead presenting it as a hypothetical state akin to a thought experiment and using it for intuition testing. Doing so can coincide with treating peerage as a matter of regard with sufficient, but not necessary conditions. There are plenty of situations where speakers will fall short of being peers, either substantially or just slightly, there many other circumstances where speakers will regard each other as equals in some respect, and with good reason. As such, we should surely consider how peer disagreement compares to non-peer cases, which I will do in the next chapter. There is little sense in reducing or limiting the conditions for peerage, since that only offers more ways to explain away disagreement too easily. What I mean by this is that reconsidering the conditions for peerage generates a trade-off between authenticity and an ease of explanation. Either the established trends of the peer disagreement discourse can continue as it has done so far, focusing largely on unrealistic cases with a less discernible source or explanation but a clearer or more pressing compulsion to respond whichever way. The alternative is to focus primarily on disagreement cases that are more familiar and common, but also more easily grounded in a speaker's error, which generate a weaker motive to respond one way or the other. If we reduce the grounds for peerage but maintain they are all necessary then we can quite easily say, of any given disagreement, something along the lines of "Oh, but they're only disagreeing because *that* speaker is better informed, so it doesn't matter if *they* stick to their guns!"<sup>295</sup>

How might one respond to what I have said? One possible line of argument would be to maintain that there is nothing wrong with thinking about ideal agents and circumstances, *iff* it is only for the sake of argument. More specifically, there is no other way to make sense of peer disagreement or perhaps the discourse is (chiefly) engaging in hypotheticals and speculation. I would respond by saying that mistaken conclusions could still be drawn from this sort of thinking, which is perhaps an unintended intellectual harm, but a harm all the same. If those who are unfamiliar with the discourse read the arguments about peerage and internalise the thought that peers are only those speakers who meet the criteria, then they may feel pressured to try to improve themselves and seek out epistemic peers. Doing so would only waste their waste time, energy, and resources unnecessarily. Alternatively, they may feel that they have no epistemic peers or equals and disengage from various debates and dialogues, essentially leading to intellectual and epistemic isolation or self-exile.

---

<sup>295</sup> It need not be one's degree of information or gathering of evidence, any asymmetry could be used to explain the disagreement between two non-peers or not-quite-peers.



Jonathan Matheson has attempted to defend and justify the use of ideal examples of disagreement in the peer disagreement literature. He notes that the question of disagreement's epistemic significance is/was motivated largely by observations about there being widespread disagreement over matters of religion, politics, philosophy, and the like. However, the focus is frequently shifted to idealised disagreements with epistemic peers that we simply don't encounter. Matheson makes a similar claim to my own that the epistemology of disagreement should tell us what to do in everyday and familiar instances of disagreement and that the impact of *those* cases is largely unclear.<sup>296</sup> He endorses the Equal Weight View as the rational response to peer disagreement and maintains that beginning with ideal disagreements provides a baseline for ascertaining responding to non-ideal disagreements:

The people we find ourselves disagreeing with are simply never exactly as intelligent as us, exactly as informed, exactly as open-minded, etc. Epistemic peerhood simply does not obtain. That said, can we learn anything about how to rationally handle everyday cases of disagreement from how idealized cases should be handled?<sup>297</sup>

Matheson is more open about relying on idealisation; he contends that by stripping away idealised factors or suppositions such as equality of evidence, we can establish how to respond to non-ideal cases. He settles on a kind of epistemic particularism, taking influence from W.D. Ross's notion of *prima facie* moral duties. The epistemic significance of everyday, non-ideal disagreements, and our response to them, should be determined by how our evidence and our reasons to believe one way or another balance. In other words, nothing universal can be said about everyday disagreement, we can only make claims based on specific cases and details regarding the precise asymmetries between speakers. Ideal disagreements are still important and "intrinsically interesting" but as a (mere) route to non-ideal disagreements.<sup>298</sup> While I agree with this sentiment, if the epistemic significance of disagreement is nothing more than a mere balance of *prima facie* defeasible considerations, it is rather hollow and lacking. The implication of Matheson's argument is that disagreement matters only as much as it gives us reason to either retain our judgment, suspend judgment, or concede to our dissenter's judgment, and the extent to which it matters depends on asymmetries in our physical and mental faculties, intelligence, evidence, and epistemic factors. I have discussed two alternative accounts of disagreement's epistemic significance: it either signals a speaker's error or is an opportunity for epistemic

---

<sup>296</sup> See Matheson, J. 2014; 315-6.

<sup>297</sup> Ibid; 321.

<sup>298</sup> Ibid; 327-8.

betterment. Both could easily accommodate Matheson's simple account *and* still say more about disagreement's significance. Since it is essentially an adapted version of Ross's moral theory, its merits as an account of rational action depend on how compelling one finds the original moral case. To be clear, I do not take issue with the arguments as an account of rational action in non-ideal circumstances, but as an account of the epistemic significance of disagreement it leaves much to be desired. Matheson is one of the few authors in the peer disagreement discourse to openly discuss idealisation in any great length, but his discussion is limited to the use of ideal examples or the presumption of epistemic symmetry in its various forms.

While its aim is to identify what is epistemically significant about disagreement, I read the peer disagreement discourse as placing greater emphasis and focus on rational action, on how we rationally ought to respond to peer disagreements. Moreover, most authors acknowledge the ideal circumstances of peer disagreements, and thus may simply implicitly regard peer disagreements as ideal representations. My issue is not with the very presence of idealisation, but on what I regard as a worrying overreliance on it. If arguments of the sort I am offering had been raised in the early stages of the peer disagreement discourse, I suspect the likes of (e.g.) David Christensen and Richard Feldman may well have bitten the bullet. I am more confident that they might now reconsider or at least be amenable to a non-idealised epistemology of disagreement. However, discussions of idealisation in epistemology are novel, they have only recently seen development, so the matter is comparatively younger than the peer disagreement discourse.<sup>299</sup> As discussions of ideal and non-ideal epistemology advance, I contend that so too will we come to recognise the extent to which the peer disagreement discourse is idealised, as well as its fundamental lack of necessity. Regardless, peer disagreement relies on distorted, potentially harmful assumptions that are dressed up with mystery and intrigue to give it some degree of philosophical glamour. This illusion needs to be broken, and the epistemology of disagreement given more substance to match its faux style.

### 7.5: Peer Disagreement Overlooks Peer Agreement

I now change tack, turning to my final argument against the peer disagreement discourse. Idealisation is prevalent but many writers in the discourse acknowledge that it is there. In contrast, a second question I have – how might the discourse on peer disagreement apply to *peer agreement*? – might turn out to be less worrying but more fruitful. As such, the discourse may be regarded as *inattentive* with regards to peer agreement. As far as I am aware it has received no

---

<sup>299</sup> For example, Jennifer Rose Carr: “ideal epistemologists are concerned with questions about what perfectly rational, cognitively idealized, computationally unlimited believers would believe... Often this involves presupposing or defending epistemic norms that, arguably, no actual humans can satisfy.” 2022; 1132.

substantial consideration and is generally confined to cases involving groups of peers and how *their agreement* bears on our judgment, almost akin to questions regarding group or expert testimony. We can find one example, again from Jon Matheson:

Evidence that one has regarding the existence of other agreeing parties will give additional evidential force that pushes away from [doxastic attitude] D1 (possibly in the direction of D2, but possibly in the opposite direction) ... Plausibly, the weight to be given to an individual's opinion corresponds to how likely that individual is to be right on the matter. Clearly, the more individuals whose opinion on the matter S1 is aware of, the more complicated the evidential story will be. In particular, the independence of the various opinions (or lack thereof) will make a difference as to how much weight each opinion is to be given.<sup>300</sup>

The speakers Matheson has in mind may be epistemic peers, or agents who are not peers but are nonetheless rational, reliable assessors of evidence. The fact that peer agreement has not been (I argue) sufficiently discussed is not indicative of idealisation. However, it is still a glaring oversight that implies a lack of substance to the epistemology of disagreement and arguably a tunnel visioned discourse. How we should treat group beliefs is also an open discussion in epistemology, with the question typically being over whether we should regard the group as having a unified identity with a single belief, or whether we should regard the belief of each member individually and independently.

The questions I now want to raise are whether the same worries I've raised about peer disagreement apply to peer agreement and whether the theories of response fare any differently between peer agreement and peer disagreement. As will become apparent, I argue that the problems I have discussed do not apply just to peer disagreement, and that the views still leave much desired when applied to peer agreement, perhaps even less than with peer disagreement. One might be curious as to why I am interested in peer agreement, given my strong doubts about peerage. To reiterate, I am sceptical that *literal* epistemic peers are as common as the discourse supposes. If we instead relax the conditions so that they are merely sufficient for peerage, and treat peerage itself as a matter of perception, we will plausibly encounter and engage more epistemic peers, perhaps even agreeing or disagreeing with them. Initially, I will just focus on cases of peer agreement, but I will address non-peer agreement and disagreement in the next chapter.

---

<sup>300</sup> Matheson 2014; 326.

While it is worrying that there is mostly silence about peer agreement, there might be mere suspicions that the positions offered as replies to peer disagreement cannot cope well when applied to peer agreement. That, as will become apparent, is my intuition. The peer disagreement discourse was largely motivated by the matter of whether speakers rationally ought to revise their attitudes or beliefs once they acknowledge a disagreement with their peer and equal. We have seen that the discourse largely maintains, and sometimes *assumes*, that focusing on ideal cases of disagreement with all asymmetries abstracted away makes it easier to see why disagreement is epistemically significant. We have also seen that I am doubtful that the current literature addresses the matter sufficiently, and these doubts will extend to peer agreement. I will now revisit the four main views about peer disagreement: Equal Weight, Extra Weight, Right Reasons, and Total Evidence, and discuss each in turn.

### Total Evidence and Right Reasons

Total Evidence, you will recall, stated (roughly) that epistemic peers should base their beliefs about a given matter not only on their shared first order evidence, but on the total body of evidence that includes the fact that each peer believes as they do. Two preliminaries need to be stated. The first is that I will consider how this view applies to speakers both when peerage is a factual matter and a matter of perception or regard. Secondly, I will also assume that each speaker's belief or attitude is an individual, independent piece of psychological evidence or first order evidence about the distribution of opinions but *not* first order evidence for the subject on which there is agreement. Imagine two speakers – Roger and David – and suppose they are discussing music. They each assert that “Paul Simon wrote the song America”, having shared all their respective evidence for it. The point of interest here, remember, is whether either or both speakers should revise or otherwise change their attitude once their agreement has been established.

In the first version of this case, Roger and David are literal epistemic peers, I assume that they each meet the conditions for peerage and have no discernible or disclosed epistemic asymmetries. In this instance, their agreement or shared belief shouldn't really count for much. They have mutual propositional evidence and two pieces of psychological evidence, the fact that Roger believes that Paul Simon wrote “America” and the fact that David also believes this, which they can and likely will take to strengthen their overall reason to believe as they do. While each of them has two pieces of psychological evidence, as it goes with peer disagreements, it would be illegitimate bootstrapping to treat their own belief as a reliable or good reason to believe as they (already) do. Therefore, they may only regard each other's belief as epistemically weighty or

evidence to heed. Suppose also that the total evidence (first and higher order evidence) supports believing that Paul Simon wrote “America”. Roger and David already believe this with justification, which suggests that their agreement does not count for very much except perhaps that they have another reason to retain their shared belief, or reason to increase the degree of confidence with which they believe.

Now suppose that Roger and David merely *regard* each other as peers, regardless of whether they are peers on matters such as folk music, songwriting history, or the works of Paul Simon. Given their high esteem of each other, they might anticipate agreement on the matter. Again, Total Evidence does not make much difference other than confirming their expectations and offering a slightly stronger overall reason to believe as they already do. Two things are worth bearing in mind though. Firstly, in both versions of the case I just considered, I assumed that the total body of evidence did in fact support belief in Paul Simon being the songwriter of America (that is, Roger and David not only had a justified belief but a justified and *true* belief.) Suppose instead that the total evidence was incredibly misleading, or that Roger and David had both confused Paul Simon and Paul Weller. In that case, they would still agree but they would also have both believed irrationally, with neither their agreement or endorsing Total Evidence making any difference at all, since they would need to know that their total evidence is misleading or has been misinterpreted. This leads to my second point, which is that Total Evidence is intended to make sense of what epistemic peers ought to do once they have shared their beliefs, it delivers no outcome of note in peer agreement because their agreement is already established. It would be absurd to compel agreeing speakers to then change their mind *unless* some decisive evidential defeater appears. We can imagine this happening but, again, it would not be the result of Total Evidence; the view suffers the same issues whether applied to peer agreement or to peer disagreement. It does not tell us what the evidence actually points to in any given case (nor could it), it merely proscribes believing as all the available evidence directs us, to the degree of confidence that it warrants.

I want to consider a third case. The discourse largely assumes that symmetry is fixed, but if someone I regard as my peer asserts that P and I later find out that they actually knew that P, are they my peer or my epistemic superior? Factually, they are my epistemic superior, their knowledge is a clear advantage. But I may still regard them as my peer, and still have very good reason to think as much. Imagine Roger and David in this situation; they still regard each other as peers and at time t1 they ascertain that they agree that Paul Simon wrote “America”. At a later time t2, it comes to light that Roger knew that Paul Simon wrote “America” whereas David merely had a justified belief. Between time t1 and t2, there does not seem to be any major

difference in what it is rational for either speaker to believe. If the evidence supports belief that Paul Simon (not Paul Weller) wrote “America”, then it would be rational for David and Roger to believe it. After Roger discloses his knowledge at  $t_2$ , the subsequent total evidence plausibly supports believing more strongly than it would if Roger and David both *merely* believed it. However, only David benefits from the evidential upgrade, since Roger already has knowledge and justification which, presumably, is 1) stronger than mere belief, 2) immune to evidential defeaters, and 3) impossible to upgrade or build on. He cannot cite his own knowledge as reason to increase her confidence, which is likely at its maximum because of her knowledge. Total Evidence delivers asymmetric results in cases where an asymmetry between speakers unfolds, only one party’s epistemic standing is affected when the other has the advantage of knowledge.

Are there any other notable or significant differences between the belief cases and the knowledge case? Arguably, learning that Roger knows that Paul Simon wrote America gives David testimonial knowledge, *iff* knowledge is transmissible via assertions and testimony, otherwise Roger’s assertion may be apt to give David testimonial belief and/or evidence. Whether assertions are fit to be both evidence and knowledge simultaneously is a matter that exceeds the focus of this chapter and detracts from my argument, although I am content to grant knowledge via testimony, and thus knowledge via agreement.<sup>301</sup> At any rate, I am not sure that the matter bears much on peer agreement or the Total Evidence view. If I agree with a peer who knows that P, then I will come to know that P. Fine. Ultimately, the Total Evidence view is a non-starter. It doesn’t really give us an account of why peer agreement is significant because peer agreement either turns out completely mundane or it bottoms out in testimonial justification or knowledge. Endorsing Total Evidence in peer agreement cases simply reiterates what peers take their evidence to indicate, and perhaps reaffirms that they have at least some reasons to believe as they (already) do.

Moreover, as we saw with peer disagreement, it offers no indication of what attitude is fitting or what belief is true, only what is justified from one’s available evidence. Total Evidence requires objective, omniscient rational foreknowledge which neither it nor peer agreement can provide. As with peer disagreement, we do not have the God’s eye perspective. Peer agreement and disagreement are both, I contend, evidence concerning the distribution of opinion about a matter, not first order evidence related to the matter at hand. We may take broad consensus or

---

<sup>301</sup> It is tempting to ask whether the two differ. The epistemology of disagreement literature occasionally identifies agreement and disagreement as testimony, but there is little explicit discussion of any relation, so it is unclear if we are to treat the relevant assertions in peer and non-peer interactions as testimony or some other phenomenon or speech-act. I offer some cursory thoughts on the matter at the end of the thesis, but the relation deserves greater consideration.

consensus from a trusted individual as strong reason to believe as they do, but it is no guarantee of the truth or certainty of that individual's view. Total Evidence does very little, then.

How does the Right Reasons view fare? This was the view that a peer's rationality is simply a product of their pre-disagreement reasoning, and whether they had evaluated the evidence correctly and responded to it rationally in the first place. In other words, once peers ascertain that they agree or disagree, they are confined to their original reasoning; in the disagreement case, either one peer was rational and remains rational, or neither of them are rational. With peer agreement, it is all or nothing – both peers are rational or neither of them are. As with Total Evidence, their agreement has no impact on their rationality. (Peer) disagreement regarding a non-permissive matter is no proof that either belief is false, and agreement over some belief is no proof that the shared belief is true. Two epistemic peers could each have the same well-reasoned, evidence-backed belief and recognise that they agree, and they could also still be mistaken and be blissfully ignorant about it. Once again, Right Reasons requires an agent to have omniscient foreknowledge for them to be rational, which cannot be provided by peer agreement and, again, is incompatible with epistemic peerage, because one agent's foreknowledge provides them an advantage that precludes any prospect of epistemic peerage. If the peers in question merely believe, but still agree, then their respective rationality will linger in uncertainty until they collectively gain some additional confirming or defeating evidence.

### Equal and Extra Weight

The Total Evidence and Right Reasons views offer nothing of note when applied to peer agreement; the only way to ensure that a shared attitude is rational is to know that it is rational, in which case one's interlocutor must also have such foreknowledge to be a peer. However, if both speakers have rational foreknowledge, then agreement is surely inevitable and the interest is lost. The Equal Weight view states that epistemic peers ought to value their opinions equally, and when disagreement between epistemic peers occurs, each peer ought to revise their attitude so that they are more closely aligned or alike. In peer agreement cases, Equal Weight cannot tell us anything that is not already obvious, nor can it compel any change of attitude. All epistemic peers involved, *ipso facto*, value their own opinion and their interlocutor's opinion equally because they share the same opinion. Roger and David are epistemic peers, they have the same evidence regarding the authorship of the song "America" and, based on that evidence, they agree that Paul Simon wrote the song. The only difference between their opinions is that Roger's opinion is Roger's and David's opinion is David's, they cannot hold each other's opinion, but their

respective opinions have the same (semantic *and* pragmatic) content, and epistemic status (justification and truth.)

Equal Weight compels disagreeing peers to value each opinion equally either by claiming that neither view is weighty at all or by splitting the difference. It is strange to think that there are any grounds for agreeing peers to alter their views, let alone to think, “My peer and I agree on P and must *therefore* each revise our belief that P.” Unless we have credences, degrees of confidence, or some other fine-grained notion of belief in mind, there are no differences to be split, and suspending judgment would be simply absurd *and* irrational. Again, neither Equal Weight nor the peer agreement itself offers any indication of whether the peers are rational or accurate in their shared belief, although the fact that they agree may strongly indicate that the agreed-upon belief is justified, so it would be odd to ignore it. It delivers nothing but inevitable results at best, or at worst it is a non-starter.

Extra Weight turns out much the same. Applied to peer disagreement, it maintains that one must at least partially value the opinion of one’s interlocutor while still being permitted to value their own opinion higher. But in peer agreement the only difference between each peer’s opinion is the agent that holds it.<sup>302</sup> Their attitudes have the same contents, and they are based on the same evidence and reasoning. Extra Weight turns out no different than a view of testimonial self-trust, at least in a weaker sense of *merely* relying on one’s own faculties.<sup>303</sup> Ultimately, between the dogmatic belief that comes from isolated self-trust and the result of applying Extra Weight to a case of peer agreement, the only difference is the fact that we get an item of psychological evidence from the latter that we do not get in the former. While not entirely worthless or inert, I contend that this evidence could be comparable to ordinary testimonial evidence. Both are equally incapable of proving decisive or of confirming the truth, fittingness, or justification of the attitude in question. The grounds for prioritising our view are unclear when our interlocutor has just the same view, and they are, from our perspective, as intelligent, informed, disposed, and accurate.

Agreement is epistemically significant because it signals the possibility that one’s belief is rational, likely true, and well-supported; it gives us some reassurance or comfort in our abilities

---

<sup>302</sup> For any given opinion, proposition, belief, or attitude, *p*, we might even wonder if belief that *p* is the sort of thing one can own, or whether there is any difference between ‘being of the opinion that *p*’ and ‘believing that *p*’. Roger and David both believe that Paul Simon wrote “America”, but does it make sense to separate Roger’s opinion, David’s opinion, and their (shared, agreed-upon) belief?

<sup>303</sup> Katherine Dormandy discusses a richer notion of self-trust that is on par with trust in others. The latter accommodates a relationship which is governed by interpersonal norms, and in which one supposes that others care about their needs. Merely relying on faculties (one’s own or other agents) fails to accommodate this. See her forthcoming.



as reasoners. However, this makes it neither proof of the truth of a belief nor higher order confirmation that a speaker is reliable or rational, just as peer disagreement does not prove either the falsity of one's belief or their irrationality and unreliability as an agent. Regardless, I contend that *peer agreement* is largely epistemically inert or valueless, even applying the views from the peer disagreement discourse makes no difference to this. Applying these views generally leaves us compelled either to maintain our belief after establishing that we agree with an epistemic peer or to absurdly value our own opinion over our peer's opinion, which delivers little that is different. If we agree with someone we regard as an epistemic peer, we will likely be somewhat reassured, perhaps greater peer agreement will reassure us more, although I suspect there will be a point at which it will level out. Peer agreement does not give us any reason to doubt or reconsider our evidence, reasoning, and judgment, even though we are fallible agents with evidence that occasionally misleads, potentially poor reasoning and mistaken judgments, and that we might think it reasonable to double check our beliefs even when we have strong reasons to be confident in them.

While the conclusions I have drawn are pessimistic, the fact remains that peer agreement has not been discussed anywhere near as substantially and as comprehensively as peer disagreement. The replies to peer disagreement do not fit well when applied to peer disagreement, but they offer even less of note when applied to peer agreement, failing to offer any account of why agreement matters except in those cases in which peer agreement speaks for itself, I argue. Others in the peer disagreement discourse may have different judgments to my own but the fact remains that the matter has been given little consideration. I will not speculate on whether this sidelining is intentional or not, except to reiterate that much of the peer disagreement discourse assumes that disagreement matters most (or most clearly) when we focus on ideal cases and setups. It is not clear whether this is also assumed without argument about agreement, but I believe I have shown that focusing only on ideal cases involving peers does not show why agreement or disagreement are epistemically significant. I argue that we need to reconsider the approach to the epistemology of disagreement that has been adopted and focus on agreement and disagreement between epistemically asymmetric agents. That is, between non-peers.

I turn to this in the chapter that follows, and as will become apparent, there are several different respects in which one can fall short of being epistemically equal or symmetric. 'Non-peer', 'epistemic inferior', and 'epistemic superior' are vague umbrella terms that can pick out many asymmetries between agents. Subsequently, identifying what the implications are for agreement and disagreement between epistemically asymmetric agents will not be as simple as

merely comparing or contrasting with epistemic peers. While we might have some intuitions, I am confident that they can and will be challenged because the different combinations of interaction and speaker-types I have in mind may generate varying questions and implications. Many of these will concern belief-dependence and deference, overlapping with considerations from the epistemology of testimony and the philosophy of expertise.

Chapter 8: Epistemic ‘Non-Peers’: Agreement, Disagreement, and Deference

I have argued that the peer disagreement discourse is neither helpful nor significant, at least not to the extent that the current philosophical literature presents it. The main views are unconvincing, have problematic implications, and rest on strict assumptions about evidence and evidential defeat. Moreover, the discourse assumes highly unattainable peerage conditions, and occasionally idealised capacities and capabilities of human agents, which distorts reality to emphasize the faux importance of peer disagreements. In short, I have argued there is a great deal wrong with the discourse and, for all the attention it has received, the outputs and achievements are nil. This inattentiveness within the discourse is particularly worrying, and having considered peer agreement in the previous chapter I will now address non-peer agreement and disagreement.

It might be tempting to think that we can straightforwardly partition different scenarios (e.g., agreement with an inferior, disagreement with an inferior, etc.) and adapt or extrapolate our intuitions. I contend that there are both general and specific implications that follow from considering non-peer agreement and disagreement. No single scenario is straightforward or simple to understand, contrary to what some might assume. “Epistemic non-peer”, “epistemic superior”, and “epistemic inferior” are vague and imprecise terms which cover a variety of relations between agents or speakers. Narrowly focusing on scenarios involving (e.g.) experts or the ignorant is no different from focusing just on epistemic peers, both miss some wider perspective. I criticised the peer disagreement discourse for its emphasis on the matter of response, this was in part because of the discourse’s reliance on idealised conditions and circumstances. The scenarios I discuss here are more realistic, familiar, and common. As such, there are three implications I believe arise regardless of which of the partitioned scenarios are being considered:

Firstly, is the urgency of collaborative and *collective inquiry*. This echoes similar claims I made about responding to peer disagreement, that neither conciliating nor remaining steadfast, nor suspending judgment are apt or optimal responses to peer disagreement. I argue that regardless of whether we are agreeing or disagreeing, all speakers should have an open mind to reinforcing their reasoning, evidence, and judgment by co-operating and clarifying. Importantly, this need not mean a blind or blanket acceptance of a given claim or argument, but seeking counterarguments for assertions or attitudes which, as Jeremy Fantl points out, are not misleading and with which we disagree.<sup>304</sup> Fantl regards open-mindedness as both the willingness to change one’s mind in light of relevant counterarguments, *and* doing so if and only *iff* one can find no issues or flaws with that counterargument. To be closed-minded, consequently, is to

---

<sup>304</sup> See Fantl, J. 2018.

spend sufficient time with an argument, find no flaws or issues with it, but to still be unwilling to change one's mind or reduce one's dissenting position. For Fantl, it is appropriate to engage with arguments and counterarguments in some situations, but not always; when you know they are misleading or you know you will be unable to sufficiently reflect on them, there is a decisive reason to not engage with them open-mindedly. Sometimes we are vulnerable or blind to arguments and therefore disposed to take it seriously, perhaps too seriously such that we are misled or jump to accept a position for no considered reason.<sup>305</sup>

The second implication that applies to each case of agreement or disagreement with an epistemically asymmetric agent is the role played by intellectual character and *motive* in relation to one's intellectual standing. In particular, the need for co-operative inquiry between asymmetric agents goes hand in hand with there being *prima facie* epistemic duties, not merely for epistemically superior agents but for epistemic inferiors as well. Those who occupy the position of the epistemic superior have other regarding duties to ensure the improvement and development of epistemic inferiors, not simply by acting as sources of true belief but by fostering understanding and by mitigating epistemic injustices. Epistemic inferiors are responsible for enabling and committing to their own improvement by engaging with and sometimes challenging their superiors. Those who occupy the position of superior are not guaranteed to be perfectly virtuous or well intentioned, they can be both knowledgeable or experienced *and* self-interested or deceptive, so we should be willing and able to challenge them.

Thirdly, following on from these two points, we can identify a reconfiguration of what inquiries may aim at and what constitutes *progress* in a domain or subject. Progress can either come from the top down, when it involves the pursuit or acquisition of new questions, knowledge, and true beliefs, or it can come from the bottom up by improving the epistemic standing of epistemic inferiors and narrowing the gap between inferiors and superiors. Both, I argue, are relevant regardless of whether one is agreeing or disagreeing, and no matter how one regards their interlocutor. With regards to each individual scenario, I will make the following specific claims:

1. Disagreement with epistemic superiors might generate some concern or challenge to our reasoning, but fixating on this only creates a hindrance. We should not be afraid to challenge our superiors and press them to clarify their reasoning and evidence.
2. Agreement with epistemic superiors ought to be pursued as a plausible aim of inquiry, but it is far from our chief or sole priority. When we do agree with our epistemic

---

<sup>305</sup> Ibid; 7, 12-13.

superiors, we should still avoid taking it for granted because we can still be deceived, misled, or mistaken.

3. Likewise, we ought not to disregard disagreement from epistemic inferiors when they could plausibly have evidence or perspective we lack, and may be dependent on us for support, guidance, or correction.
4. Agreement with epistemic inferiors might raise doubts about our reasoning or standing, but we should ascertain that we both have comprehended and assessed the matter correctly while avoiding isolation and remaining open to re-evaluating, if necessary.

Regardless of the scenario, the upshot is clear: we ought not take our interactions with those we regard as epistemically better or worse off at face value, we ought to be mindful of our mutual epistemic interdependence and fallibility as non-negotiable and inescapable parts of open inquiry. I argue that dismissing these scenarios out of hand presumes too much. In sections 1 and 2, I discuss certain accounts of epistemic asymmetry from the literature, as well as discussing others that have not been discussed. Sections 3-6 address each partitioned scenario and their respective differences and unique considerations. I summarize and reiterate the urgency of collective inquiry in section 7.

### 8.1: Epistemic Superiority

I begin by outlining the various possible respects in which speakers can be epistemically asymmetric. So far, I have focused on peers who are equal in the following ways:

*Evidence* – speakers' (access to) evidence and their ability to process it.

*Cognitive* – physiological and mental capacities (e.g., perception), and readiness (i.e., whether they are attentive or distracted, and inhibited or enhanced in some way).

*Dispositional* – strength of possessed intellectual virtues, vices, and underlying biases, or likelihood to respond rationally to evidence.

*Accuracy* – time spent considering the matter the matter at hand, likelihood of having a true belief on it, or having more true beliefs than false beliefs in the wider subject domain.

These offer some starting points, but we will soon see that asymmetry extends well beyond them. Broadly speaking, we would likely think of an epistemic superior as anyone who is more than an epistemic equal or peer. The distinction between regarding someone as a superior and

being someone's superior by a matter of fact echoes the same distinction as it occurs in the peer disagreement debate.<sup>306</sup> I have already said that I think we are better off thinking of epistemic peerage as a matter of regard rather than fact, so for the ease of discussion I will assume through this chapter that all epistemic standing relations are also matters of regard rather than facts. Since the question of what makes someone epistemically superior (or inferior) comes apart from that of responding to agreement or disagreement, I will consider both matters separately.

Taking the list of peerage conditions first, we can say that a speaker is our epistemic superior if we regard them as exceeding beyond the marks of peerage or equality, being better than us in some or all of the following respects:

Having stronger or greater *evidence* and/or more competently *understanding* it.

Having greater background *knowledge* and/or greater overall intelligence.

Having more *true beliefs* than false, and/or a greater track record of *justified* beliefs.

Having greater intellectual *virtues* and/or fewer intellectual *vices* and *biases*.

Being more *likely* to form true and/or justified beliefs from the available evidence.

### Experts

We would typically regard the *expert* as the prototypical epistemic superior for any given domain, discipline, or subject. Many early accounts draw on the conditions above, i.e., experts are those who have the greatest and most relevant skills, experience, and resources, and usually the most knowledge or justified beliefs. One of the earliest accounts of expertise comes from Alvin Goldman, who distinguishes three accounts of expertise:

*Cognitive expertise*: possessing more true beliefs or higher degrees of confidence in true propositions compared to others in a domain.

*Objective expertise*: possessing cognitive expertise and the disposition to use it to develop new questions and inquiries, and attain or ascertain true beliefs or knowledge regarding them.

*Reputational expertise*: being regarded as an objective expert by some external party (with or without justification).<sup>307</sup>

---

<sup>306</sup> E.g., Bryan Frances, in his 2013, offers an account of a *recognised* epistemic superior – someone who I acknowledge has all my evidence and more besides, is smarter overall and has thought about the matter at hand more than I have, and who I would otherwise (before finding out their opinion) insist is more likely to have a true belief.

<sup>307</sup> See Goldman, A. 2001; 91-2.

Goldman's distinction provides a natural foundation, with many of his conditions coinciding with those above. I take them to emphasize the importance of a speaker's motive or disposition, which can surely come apart from their degree of knowledge, one can acquire expert-level knowledge or true beliefs by good fortune rather than any effort at inquiring. Likewise, we can imagine experts with great knowledge and familiarity regarding a subject and a malicious disposition to deceive others.

The objective sense of expertise implies that experts contribute to the development or progress of their domains, but Goldman only has a top-down approach to progress in mind. Ascertaining new questions and generating further knowledge or true beliefs advances a domain in one way, but it widens the gap between epistemic superiors and inferiors. Narrowing that gap and improving the epistemic positions of one's epistemic inferiors is, I argue, another (bottom up) approach to progress or advancing a domain. This is done by not merely communicating knowledge or true beliefs to those who occupy the position of the epistemic inferior, but by helping them to gather more relevant evidence, develop their understanding and pursue intellectual interests and relevant inquiries. An epistemically superior agent who exercises their superiority in this way would be very different from someone who merely aims for further knowledge.

Goldman is not the only one to discuss or suggest an account of expertise; Adam Elga contends that superiority is determined by how we respond to a speaker's opinion. An expert, for Elga, is any speaker whose opinion we defer to completely and unconditionally. If we defer to the opinion a speaker would have if they possessed evidence that we have, we treat them as a *guru*. Elga regards it as unwarranted to dismiss anyone's opinion, even if they themselves are behaving irrationally and render themselves exempt from consideration as an expert or a guru.<sup>308</sup> Importantly, since we are never fully certain or able to identify reliable speakers, we may only regard a speaker as an expert in ideal circumstances, which is why Elga's notion of a guru is conditional and why he regards a peer in looser terms than most.<sup>309</sup>

### Authorities

Not all notions of epistemic superiority are grounded in expertise; some instead maintain a distinction between experts and *authorities* of various sorts. For example, Linda Zagzebski:

---

<sup>308</sup> See Elga, A. 2007; 479-80. Since Elga thinks a guru is someone whose conditional opinion is what we would defer to, one might think that a guru is, in essence, a *merely possible epistemic superior*. Or, indeed, that they are actually an epistemic inferior, since they lack evidence that we possess.

<sup>309</sup> Ibid; 483-4. For Elga, an epistemic peer is merely someone who has access to all the same evidence and who we regard as being equally good as ourselves at evaluating (factual) claims.



The general point is that an *epistemic authority* is someone who does what I would do if I were more conscientious or better than I am at satisfying the aim of conscientiousness – getting the truth.<sup>310</sup>

Or, in response to Zagzebski, Christoph Jäger:

If someone is a Socratic authority for a subject in a given domain, he not only has a higher ratio of true to false beliefs in the domain than the subject does. He also displays superior methodological skills and insights which enable him properly to assess evidence, reasons, methods of thinking and investigation, and so on, and to communicate such insights to others.<sup>311</sup>

Taking each in turn, Zagzebski maintains that *epistemic conscientiousness* is the quality to use our cognitive capacities as best we can to get at the truth. She contends that there will be self-reflective agents who desire truths but do not always look for reasons to believe, or at least not for all their beliefs. As one becomes more conscientious, so too do they become more epistemically virtuous.<sup>312</sup> Authorities are the exemplar in this regard, they are someone we regard or recognise as more conscientious than ourselves. For Zagzebski, we ought to defer to the reasons an authority possesses for believing as they do and replace our own reasons for believing as we do with them. Jäger responds to Zagzebski's account of authority and contends that one may instead be an authority by being a reliable source of understanding, which some inquiries prioritise over truth or knowledge. Socratic authorities demonstrate their authority by showing us where and how we have misjudged or reasoned wrongly.

### Hybrid Accounts

Michael Croce attempts to offer distinct accounts of expertise and epistemic authority while drawing on both Jäger's and Zagzebski's accounts and addressing what he regards as their shortcomings. For Croce, the distinction depends not merely on their epistemic goods but on abilities and character traits as well:

*Expert.* A subject A is an expert in domain D (for a subject S) iff:

- (1) A has more accurate information than the majority of people do in D;

---

<sup>310</sup> Zagzebski, L.T. 2012; 109. My emphasis.

<sup>311</sup> Jäger, C. 2016; 179.

<sup>312</sup> See Zagzebski, L.T. 2012; 48-9.

- (2) A possesses *expert-oriented abilities*.

*Authority of belief (AofB)*. A subject A is an AofB in domain D for a subject S iff:

- (1) A is more conscientious than S – who considers her to be an EA – in D;  
(2) A possesses and makes use of *sensitivity to S's needs*.

*Authority of understanding (AofU)*. A subject A is an AofU in domain D for a subject S iff:

- (1) A is more conscientious than S – who considers her to be an EA – in D;  
(2) A possesses and makes extensive use of *novice-oriented abilities*.<sup>313</sup>

Experts have greater overall subject-relevant evidence and expert-oriented abilities, inquiry directed intellectual virtues such as intellectual curiosity, creativity, courage, and autonomy. They are committed to what I regard as the top down approach to progress in inquiries by identifying new questions, sufficient evidence, and correct evidence. Authorities prioritise fellow inquirers more, they are personable and exhibit sensitivity to the needs of all members of a domain. There are two kinds of authority, both are ascribed externally in virtue of some audience perceiving the relevant traits in the authority. We regard someone as an authority *of belief* if they are sensitive to our intellectual needs and deficiencies, and as an authority *of understanding* if and only *iff* they are sensitive to our needs and exhibit intellectual empathy and generosity by helping us to properly and fully address those needs.<sup>314</sup>

## 8.2: Epistemic Inferiority

Much of the epistemology of testimony asks whether epistemic goods can be acquired with testimony, e.g., whether it is possible for a creationist teacher to convey knowledge of evolution to her students when her creationist convictions mean she does not herself believe evolutionary theory. Cases pertaining to moral or aesthetic testimony are typically set up such that the subject either approaches someone they regard as reliable and sincere or reflects on the claims from someone who is held in high intellectual esteem, usually groups and/or experts. The subject either has no firm belief or conviction on a matter, is uncertain of what to believe between two or more possibilities, or is simply seeking reassurance that they have judged things right. For example, Sarah McGrath contrasts three different cases. In the first case, the subject independently reflects on the morality of eating meat before judging that it is morally wrong to do so. In the second case, the subject believes that eating meat is morally wrong purely on

---

<sup>313</sup> Croce, M. 2018; 23-4.

<sup>314</sup> See Croce, M. 2018; 22-4.

someone else's say so, with no knowledge of what relevant non-moral information they have or how much they have considered the topic. In the third case, the subject acknowledges their similar moral sensibilities to an interlocutor, as well as other relevant asymmetries, such as the amount of relevant non-moral information each possesses, and how much each has reflected on the question.<sup>315</sup> Since the interlocutor outperforms the subject in these regards, the subject defers to their judgment.

We get an idea of epistemic inferiority from testimony cases, but no formal conditions for occupying that position. One can doubt what to believe about a matter regardless of whether they are well positioned to form a belief, so we can rule out doxastic uncertainty. Just as epistemic superiors are those whom we regard as exceeding beyond the merits of peerage, we may regard someone as our epistemic inferior because they fall short of being our epistemic peer for any of the following reasons:

Having weaker or less *evidence* and/or less competently *understanding* it.

Having less background *knowledge* and/or less overall intelligence.

Having more *false beliefs* than true, and/or a limited track record of *justified* beliefs.

Having fewer intellectual *virtues* and/or more intellectual *vices* and *biases*.

Being less *likely* to form true and/or justified beliefs from the available evidence.

Our regard of someone's epistemic standing does appear to be something of a balance of factors, which can be outweighed such that one might be inferior in one or two respects but superior in considerably more. Perhaps we know we have some advantage over a speaker, that their reasoning is compromised, or that their evidence is flimsy.

### Testimonial Disregard

The named kinds of superiority I discussed in the previous section provide some corresponding notions of inferiority, although they are not discussed by the authors themselves. For example, Elga argues that we treat a speaker as an expert when we defer to their opinions unconditionally, but only in ideal circumstances and never in real life. McGrath's second case describes just the same kind of deference, applied to a moral matter, and calls it *pure moral deference*.

---

<sup>315</sup> See McGrath, S. 2009; 321-2.

In cases of pure moral deference, one in effect treats the person to whom one defers as having purely moral information that one lacks; in this respect, one treats the other person as a moral expert in the strong sense.<sup>316</sup>

Whether we consider the interaction in terms of what we are doing (deferring) or who it involves (e.g., experts), the corresponding case would be that we treat someone as our inferior in *some* sense if we entirely and unconditionally disregard their opinion. Elga thinks gurus are speakers whose opinions we conditionally defer to, so it is plausible that disregarding opinions comes with a degree of flexibility, perhaps even certain constraints. In this sense, a speaker's epistemic standing is determined by the extent to which we (would) consider or disregard their testimony, rather than how we regard them and their ability to judge.

### Novices and Subordinates

Jäger and Croce's accounts of epistemic authority largely correspond to the peerage conditions, so we already have notions of epistemic inferiority that mirror them. Given the emphasis Croce places on sensitivity and support for *novices*, we might suppose that a novice is someone who is epistemically dependent on others. It is certainly used to describe people that lack subject-relevant knowledge, evidence, resources or abilities, sometimes as synonymous with being a layperson.<sup>317</sup> Recall that Zagzebski regards epistemic authority as dependent on one's epistemic conscientiousness. If we are an epistemic authority for another speaker who is less conscientious or worse at getting at the truth than us, then they are epistemically inferior in a corresponding way. I contend that we might think of such speakers as our epistemic subordinates but, again, this is my reading rather than Zagzebski's.

### Competent Inferiors

To the best of my knowledge, Maura Priest provides the most dedicated discussion of epistemic inferiority. She contends that there are at least some epistemic inferiors whose opinion bears epistemic weight and import which we ought to heed. Her account of epistemic standing is relatively limited and only specifies evidential and cognitive asymmetry, while distinguishing one's epistemic standing from their *competence*, which is their likelihood of forming more true beliefs about a subject than false. A speaker is epistemically inferior if they are cognitively or evidentially worse off than us (or both) but they are a competent inferior if they are still more likely to have a true belief concerning the matter at hand than a false belief. High competence is not equivalent

---

<sup>316</sup> Ibid; 322.

<sup>317</sup> E.g., Goldman 2001; 89. Croce 2018; 2.

to being an expert because expertise is context-dependent, it varies between subjects. Competence merely indicates a speaker's epistemic merit and their track record or veridical history.<sup>318</sup> Not all speakers in a domain are epistemically incompetent, and those speakers who are incompetent are not equally incompetent. However, there are different explanations or grounds for incompetence, some which have worrying implications if gone unchecked or taken for granted; Priest distinguishes named examples of epistemic inferiors based how their inferiority is explained:

### Incompetent Inferiors

Some speakers can be incompetent because they are genuinely *ignorant* about a subject, unfamiliar with the prominent questions, arguments, or evidence. These speakers are well meaning and sincerely aspire for true beliefs but are fundamentally incapable of acquiring or forming them. Other agents are willingly incompetent and form beliefs arbitrarily, neither caring about the truth of their beliefs or intentionally being deceptive or misleading. These speakers are *bullshitters*, the arbitrariness of their opinions comes from their indifference and apathy, and the likelihood that their views are right is equal to the likelihood that their views are wrong. Finally, some inferior speakers are *maliciously* incompetent, they weaponize their inferiority and deliberately perpetuate falsehoods and deceive others, often consistently asserting beliefs or opinions that are the result of false higher order beliefs or biases.<sup>319</sup>

### 8.3: Disagreement with Epistemic Superiors

Epistemic standing is the product of a balance of various factors and reasons which all contribute to our overall judgment and regard for our interlocutors. If I have more discernible reasons to regard someone as better positioned to make a judgment on a matter (myself or someone else), then they are epistemically superior with regards to that matter, from my perspective. The asymmetry between speakers may be slight or substantial, and this will determine our response to agreement and disagreement with each other to some extent. I contend there is rarely, if ever, any immediate justification for disregarding the opinions of other speakers and that we should be inclined towards collective inquiry and clarificatory dialogue by default. Doing so contributes to inquiries, advances domains, and potentially improves the epistemic standing of inferior speakers. Having examined the available and plausible accounts of epistemic superiority and inferiority, I now turn to partitioned scenarios of agreement and disagreement.

---

<sup>318</sup> See Priest, M. 2016; 265.

<sup>319</sup> Ibid; 279-80.

I was unconvinced by there being any all-encompassing, uniform notion of inferiority or superiority, in the same way that I was unconvinced that all the conditions for epistemic peerage must be considered necessary. We might regard it as intuitive to respond to certain scenarios in a particular way; my aim is to put pressure on those intuitions and ascertain whether they hold weight. I think we can surmise certain implications of agreeing or disagreeing with superiors and inferiors just by considering general, broad-brush cases (i.e., by using “inferior” and “superior” as catch-all terms). This is not to say that the finer details, such how we specifically regard a speaker’s epistemic standing, will never have any impact, but a general overview of each scenario should suffice for the purposes of my discussion. Much of what I consider in the remainder of this chapter will reiterate or present in new ways claims and considerations that are less often discussed in the disagreement literature but do appear in the epistemology of testimony.

I begin with disagreement with epistemic superiors. That is, any circumstance in which we find that we draw different judgments about a matter than someone we regard as better placed to make a judgment or form a true and/or justified belief about that matter. Suppose someone who denies the existence of climate change is discussing global warming with a climate scientist who has substantial experience in researching the phenomenon and its effects, more than the sceptic by far, as well as plenty of evidence which he comprehends well. The climate sceptic would, to most, be regarded as the epistemically inferior party. I am interested here in what the people in the same position as the sceptic should do, i.e., when they regard someone as epistemically superior and ascertain that they disagree with them about the matter on which they consider them superior.

Importantly, I leave open the matter of whether the sceptic regards the scientist as their superior. Because the climate denier doubts the possibility of climate change and global warming, they may not believe there is any way to compare themselves with the scientist, or they may simply think there is no specific common ground on which to regard the scientist as a superior. I will say more on this later, for now this should be sufficient to illustrate that I have in mind primarily cases in which one speaker regards themselves at an obvious intellectual disadvantage, without necessarily specifying which respect or the extent to which the asymmetry extends. I believe three concerns arise from disagreeing with someone we regard as our epistemic superior:

- 1) That it implies or signals an error in our reasoning or judgment.
- 2) That it signals vulnerability, specifically that that comes from inferiority.
- 3) That deferring to a superior’s dissent compromises one’s intellectual autonomy.

The first is the clearest and most intuitive worry, disagreeing with our epistemic superior will likely cause us to believe we have made some rational or intellectual error. I say it is intuitive because in any group of epistemically asymmetric agents, someone is most likely to have judged things right and someone else is most likely to have judged wrongly. Occupying the position of the inferior means being closer to the latter.

A superior speaker will likely have all our evidence and more besides, so we might regard disagreement with superiors as clear defeating evidence, and a reason to concede or defer to their judgment, if not to suspend judgment on the matter.<sup>320</sup> Linda Zagzebski regards unconditional deference to epistemic superiors as intuitive and unquestionable; we ought to regard them as our best sources of information, evidence, or belief:

... in certain cases the authority stands in for me in determining whether *p*. I have reasons for thinking that the authority is more trustworthy than myself in some range of beliefs, based on aspects of myself that I trust, but that is compatible with letting the authority's determination of the truth within that range replace my own determination of the truth within the range.<sup>321</sup>

Zagzebski justifies trust in epistemic authorities and their beliefs with two conscientious judgements that she believes any individual would grant or possess:

1. That we are each more likely to form a justified belief by believing as the authority does, than by trying to figure out what to believe by ourselves.
2. That the authority's belief would be more likely to survive conscientious self-reflection than any belief we would reach by our own efforts.

She acknowledges that these two claims are insufficient to justify the authority of a speaker's belief by themselves, that minor differences may lack the requisite weight and that we can make mistakes in identifying authorities.<sup>322</sup>

Zagzebski appeals to an argument from Joseph Raz to justify this, with Raz maintaining that we are more likely to generate a track record of false beliefs if we treat the beliefs of putative authorities merely as ordinary reasons to believe, without ascribing them some special

---

<sup>320</sup> See, for example, Bryan Frances's 2014; 55-6.

<sup>321</sup> Zagzebski 2012; 107.

<sup>322</sup> Ibid; 109-111.

significance. We ought to pre-emptively favour the beliefs and reasons of authorities and allow them to screen off whatever alternative beliefs we might come to from making up our own minds.<sup>323</sup> While I acknowledge that disagreeing with a superior will cause hesitation and uncertainty, I consider it too hasty to automatically defer to someone's view just because we regard them as better than us regarding the inquiry, subject, or matter at hand. Our perceptions can be mistaken and taking a belief or assertion for granted risks deception and error all too easily. Such disagreement is a mere fact of circumstance and the result of human infallibility. No-one is omniscient or has perfect knowledge and understanding so, in other words, epistemic superiority is not the same as epistemic perfection. However, everyone will be epistemically superior in some domain or discipline, and inferior in plenty of others.

The second concern – that epistemic inferiority carries a degree of vulnerability – follows naturally from this. In addition to not being (as) well placed to make judgments that are accurate or justified, we may not be especially discerning of those judgments that are so, leaving us liable to deception or exploitation from superiors with malicious intentions. Pre-emptively deferring to a superior's opinion does very little to mitigate this vulnerability. If there is a substantial asymmetry between two speakers, then the one who is worse off may be very easily led or gullible, which any superior could exploit (for good intentions or ill). Epistemic superiority does not guarantee good intentions or virtuous character, so anyone that occupies the position of the inferior could find themselves locked into accepting as rational the view of whichever superior they encounter first. As Jennifer Lackey notes:

... once I identify an expert, her testimony is authoritative for me: all of my evidence on a given topic is normatively screened off, replaced by only what the expert tells me. Given this degree of reliance, picking reliable experts is of paramount importance epistemically. But as noted earlier, widespread disagreement in a domain seriously calls into question our ability to identify reliable sources in that domain.<sup>324</sup>

The fact that we risk deferring to a potentially pseudo-authority should compel us to always be on our guard. We ought not to treat superiors as oracles, even those who we confidently believe are well-intentioned and virtuous or well-informed and experienced superiors. Again, this is a brute fact that follows from the non-ideal circumstances in which we inquire. Thus, as Lackey suggests, we should avoid deferring if there is reason to doubt the accuracy of an authority's

---

<sup>323</sup> See Raz, J. 1988; 68-9.

<sup>324</sup> Lackey, J. 2018; 233-4.



assertion, or we should treat the testimony of authorities as *advice* or as guidance to reflect on, rather than bare assertions to endorse by default.<sup>325</sup>

The third concern from disagreement between asymmetric speakers is that inferiors face a dilemma of either having to relinquish their intellectual autonomy by deferring to the superior's opinion or retaining it and taking responsibility for their rationality by disregarding the disagreement altogether. However great a failure of rationality (if any) is risked by ignoring disagreement with an epistemic superior depends on the extent of the asymmetry between the inferior and superior. I discuss each of these in turn. Some of the criticism against acknowledging the import of higher order defeat was that it compels agents into epistemically akratic decisions, to believe against their evidence or change their mind for weak reasons. What Zagzebski regards as a necessary outcome from the authority-inferior relation could have just the same implications. Alternatively, we may instead change or downgrade our perception of a speaker's standing. Downgrading a peer's status has been given some consideration, subject to certain constraints, so I do not consider it too absurd that these constraints hold for re-evaluating non-peers as well:

However, if person A is to be justified in downgrading person B, it is not enough that A and B are not equals as a matter of fact. Rather, in order for A to be justified in downgrading B, A must have a good independent reason to suppose that B is her epistemic *inferior* in the context of their dispute.<sup>326</sup>

Although I suggest that there is a probable trade-off between intellectual autonomy and deference with epistemic superiors, this only applies for particularly strong notions of intellectual autonomy. At least some believe that deferring has far greater benefits than independently pursuing one's own research and inspecting the first order evidence for oneself. For example, Neil Levy notes, where a subject or topic has established experts and expert opinions (whether a unified consensus or a range of opinions), layfolk and non-experts stand to more reliably acquire knowledge by deferring than by going it alone. Experts are more obviously compelled to do their own research, since they possess the relevant competences and resources, whereas non-experts typically cannot discern trustworthy sources or weigh evidence. Even if deferring brings the cost of the right to intellectual self-governance, there are only a rare set of cases where one gains

---

<sup>325</sup> Ibid; 238-9.

<sup>326</sup> Simpson, R.M. 2013; 572. Interestingly and importantly, Simpson takes this as reason to doubt any notion of peerage that relies on evidential and cognitive equality, since there is little chance that parties to any complex dispute will be strict equals in these regards, or that either agent can justifiably downgrade the other.

knowledge more reliably by not deferring.<sup>327</sup> Levy does note a further trade-off; he contends that engaging in our own research secures better understanding of a subject matter but does not ensure any knowledge for the independent researcher, whereas deferring brings knowledge but not understanding. Moreover, even ‘official’ stories or lines from expert consensus can be false or misleading. Deferring to that official line will not enable a non-expert to recognise this, but lay research may clarify the deception.<sup>328</sup> Elsewhere, Levy argues that intellectual autonomy requires us to exhibit a virtue, in some cases through ‘respectful listening’ and in others by shouting down. We exercise this virtue by refusing to listen to speakers whom we know to be acting in bad faith and preventing them from being heard by third parties, and engaging with or deferring to speakers whom we know are reliable and virtuously motivated without being excessively reliant or servile.<sup>329</sup>

Other concerns about deference have been discussed, with a prevailing worry being that outright deference is an intellectually irresponsible act that compromises one’s epistemic autonomy. Again, responses vary and tend to cite the imperfection of epistemic agents and circumstances. Some, including Zagzebski, maintain that deferring is rational because it will typically give us a better belief than any we would form or acquire independently.<sup>330</sup> Therefore, we should outsource our inquiries to epistemically superior agents and focus our attention to achieving what we know we can do well. Against this are those that contend epistemic autonomy gives us a degree of self-governance and self-direction, even if it means we are fundamentally responsible for any failures in judgment. Failed inquiries humble us as agents and demonstrate our shortcomings, enabling us to cultivate epistemic virtues even if we don’t acquire epistemic goods such as knowledge or understanding.<sup>331</sup> I have argued that conciliatory replies to peer disagreement lead to defeatist and indefinite judgment suspension, and that such views underestimate the capabilities of many agents. I contend the same can be said of deferring outright to epistemic superiors, that a blanket or uniform response of conceding whenever one draws a different judgment to their epistemic superior offers very limited benefits for both parties.

We can rely on others without completely relinquishing our self-governance and without being utterly dependent. The potential gains of collective inquiry are far greater than those of

---

<sup>327</sup> See Levy, N. 2022; 5.

<sup>328</sup> Ibid; 10.

<sup>329</sup> See Levy, N. 2024. Levy goes on to propose an alternative emphasis on ‘intellectual independence’ as being disposed to think *interdependently* rather autonomously and *independently*. He borrows the term intellectual independence from Heather Battaly’s 2021.

<sup>330</sup> See Zagzebski 2012; 108-9. See also Worsnip, A. forthcoming; 3.

<sup>331</sup> See Matheson, J. 2022; 13.

complete deference, and while epistemic autonomy means self-direction, it need not mean intellectual solitude or isolation, as Jon Matheson and Nathan King have each noted:

... the independent epistemic life looks impoverished if not impossible. We routinely rely on others for language, ideas, information, as well as a number of other epistemic tools. In addition, we rely on others to develop our skills in using these tools. None of this is in conflict with autonomous deliberation.<sup>332</sup>

Here's something less obvious: as narratives of intellectual autonomy, the stories as told above are woefully incomplete. They paint their protagonists as solitary rebels, and omit the ways in which these thinkers relied on others' intellectual capital. But for all its independence, intellectual autonomy is not a matter of going it alone... We can put it like this: Autonomy requires thinking *for* ourselves, but not *by* ourselves.<sup>333</sup>

Again, we stand to gain more from treating superior testimony as advice and considering it alongside our current evidence and reasons than from merely repeating it verbatim or dismissing it out of hand. Moreover, if superiors are disposed to aid inferiors (as I will argue they ought to be), then disagreement will be less of a hindrance to inquiry than it might seem; it will certainly be more civil and help to advance inquiries. Nonetheless, we ought not to worry every time we disagree with an epistemic superior. We gain more and lose little by challenging them and pressing them for reasons and arguments than we do from simply taking what they say for granted. If we think there is something wrong with what someone says, or if we are unsure of what they mean then, regardless of their position or standing, we have the choice to accept it or dismiss it, but also to consider it at length first. Agreement with someone is not a requirement, even if we regard that someone as more likely to be correct than us or more likely right than wrong. Disagreement with an epistemic superior should give us reason to pause and reconsider, but we should be prepared to ask the superior to clarify and explain themselves as well.

I will consider one final point. When I introduced the scenario of disagreeing with an epistemic superior, I used the example of a climate scientist and a climate change denier disagreeing over the possibility and impact of global warming and suggested that while we as an external third party might regard the two as epistemically asymmetric, the climate denier might not. While they might grant that the scientist is a general superior with regards to the natural sciences, with substantially greater background knowledge of the subject and experience in the

---

<sup>332</sup> Ibid; 3.

<sup>333</sup> King, N.L. 2021; 88.

relevant sub-disciplines and inquiries, they might not be convinced that there is a specific topic on which they can compare themselves. That is, because the climate change denier is fundamentally sceptical about climate change and various related phenomena, he may be disposed to think of those who affirm or accept climate change as a real or possible phenomenon, as neither his inferior nor his superior, but as simply incomparable. Bryan Frances considers the same possibility using the example of religious disagreements between theists and atheists:

If you are an atheist, then in one pertinent sense you will consider any theist to be much less informed about C and any other religious matters compared to yourself, since you will think that almost all of her interesting religious beliefs are false (or based on false presuppositions). But in another sense you will judge her to be your epistemic superior on religious matters, and even particular claim C, as she knows a lot more about religion and various theories of the trinity than you do.<sup>334</sup>

While I think we can confidently say that the theist and the atheist are liable to disagree over the existence of God, and climate scientists and sceptics can differ over matters such as the existence of global warming, whether there is any intellectual disparity is very much an open matter.

As Frances notes, the same speakers disagreeing over the same matter might regard themselves as intellectually asymmetric in some respects, and intellectually incomparable in others. From the outside looking in, we might have a clear perception, but I am interested here in how the participating speakers regard each other, which is somewhat obvious in some cases, but not in all. My overall argument that reflecting on non-peer agreement and disagreement scenarios lends support to collective inquiry would not, I suspect, be undermined or defeated by the consideration of cases where one's comparative epistemic standing is uncertain. While the uncertainty some speakers may have about how they stand in relation to others regarding certain subjects does complicate the discussion further, it does not occur in every subject. There are many subject matters which admit of clear epistemic asymmetries, and they are the kinds of scenario I am chiefly interested in here, but I acknowledge that cases such as the ones I have just mentioned deserve further attention elsewhere.

#### 8.4: Agreement with Epistemic Superiors

---

<sup>334</sup> Frances, B. 2013; 129.

I have argued that disagreement with epistemic superiors should not be taken for granted or disregarded without cause. I now consider whether agreement with epistemic superiors warrants the same reaction. The scenario has not often been discussed in either the testimony or disagreement literature, but the intuitive response might be relief and reassurance about our reasoning. After all, epistemic inferiority typically means the odds of having justified or true beliefs are stacked against us, and that we might think such agreement is highly unlikely. Of the matters that are discussed in the literature, the nearest alternative cases are those involving superiors that agree with each other and the question of whether an inferior ought to let this consensus determine or direct their belief. Broadly, I regard the scenario of agreement with an epistemic superior as raising the familiar question of response, even though we might intuitively think that there is nothing more that needs to be done when we actually encounter such agreement. There are two dimensions to the question here, which I consider in turn:

- 1) What, if anything, is the aim or end of inquiry, or what is the inquiry stopper?
- 2) Is it rational to continue inquiring beyond a point of inquiry cessation?

Closing or stopping inquiries has received plenty of attention, if we can ascertain that inquiry aims at something specific and agreement with epistemic superiors is conducive or equivalent to that aim then there is an obvious reason to pursue it.

I contend that we need not settle on a specific aim of inquiry or account of the inquiry stopper to demonstrate why such agreement matters, since it is conducive to various epistemic goods and forms of progress in a domain. Knowledge and understanding might be the ideal ends of inquiry, but they are not always attainable or discernible in some domains, so agreement or consensus between those occupying the positions of superiority might be the best alternative. I have already argued that advancing a subject or domain, or contributing to progress within it, is a plausible aim of inquiry, and thus for agreement and disagreement, which is not acknowledged within the peer disagreement discourse. However, as Jane Friedman notes, the question of what the aim of inquiry is can be understood in two different ways: there is the “General Question” of what it is for inquiry to have an aim, and the “Specific Question” of what, if anything, that aim is.<sup>335</sup> The latter admits of further distinctions pertaining to whether we are talking about the practice of inquiring as having an aim, or of individual inquirers or groups having personal or shared aims. Many who write on the aim of inquiry, says Friedman, do not specify or clarify which question they are interested in, although many usually defend their preferred account of

---

<sup>335</sup> See her 2023; 1.

the aim of inquiry, sometimes granting a plurality of candidates and/or prioritising one as essential.<sup>336</sup> Ultimately, Friedman settles on a de-facto pluralism about the aim of inquiry:

All inquirers may well have some aim qua inquirers, but it doesn't look as though they all have, of necessity, one of the acquisitive aims that are the focus of the debate over the specific aim of inquiry. And it doesn't even look as though every inquiry needs to have any sort of acquisitive aim at all. It may well be true that we typically inquire because we want to know more but that sort of desire or personal goal is not essential to the activity.<sup>337</sup>

She contends it is implausible that all inquirers share a specific aim such as knowledge or true belief and suggests that this can be empirically verified.<sup>338</sup> There is no consensus on a single, uniform aim or end of inquiry, and therefore of the inquiry stopper. I argue that agreement and consensus with epistemic superiors is a viable possibility, it is *an* end of inquiry, but it does not stand as *the* end of inquiry.

Regarding the second dimension, many might maintain that there is something strange about continuing to pursue questions or matters that are long closed, or that have remained dormant without development for some time. That is, supposing that there is a discernible inquiry stopper, some regard it as odd or absurd to keep going once one has reached that stop point or attained the relevant good. I think there is something intuitive about this hesitation; after all, what sense is there in trying to prise open a closed question? However, most domains and matters are rarely truly closed or regarded as settled, so I am not convinced that it is entirely irrational to continue to inquire into certain matters. Even those questions that are regarded as resolved will likely attract some interest, if only for the sake of learning about it or gaining some degree of competence in it. Some participants will acknowledge that a) they are far from occupying the position of the epistemic superior, and b) that they are highly unlikely to close any gaps in knowledge or certainty for the entire domain but will still desire to improve their epistemic standing. This is done by inquiring, consulting, considering and, indeed, sometimes by agreeing and disagreeing. Even if we find our inquiries lead us to agree with a long-established or well supported view, and even if we think further questions or consideration will drag things out unnecessarily, I regard it more problematic to take agreement for granted and let it pass unchecked.

---

<sup>336</sup> Ibid; 12, fn 24.

<sup>337</sup> Ibid; 14-15.

<sup>338</sup> Ibid; 12-14.

This, I contend, gives us an answer to the question of response for agreement with epistemic superiors. Just as we should not let disagreement with superiors compel us into defeatism or indefinite suspension, neither should we doggedly pursue agreement at the cost of other epistemic goods. To be clear, agreement with superiors should reassure us, but it can be apparent and misguided, it does not ensure that we are in any way accurate or justified in our shared judgment. Epistemic superiority is not epistemic perfection and, as should now be clear, comes in various forms. There can be self-interested experts and incompetent but well-meaning authorities. Collectives made up of superiors may not exercise any intellectual independence, resulting in *groupthink*, the tendency to share beliefs or converge on common judgments with varying degrees of independence. By taking agreement at face value, any speaker risks overlooking this.

Recall the distinction between being in the *state* of disagreement and engaging in disagreement as an *activity*. I contend the same distinction applies to agreement, and that some states of agreement are minimally valuable. Agreement from deference involves no effort to assess what a superior's testimony means or amounts to, it risks deception or misinterpretation, and it is simply echoing their judgment and offers little of epistemic worth to the agent, except perhaps a potentially true belief. Agreement through individual or co-operative inquiry demonstrates intellectual curiosity and humility; if it comes from a joint pursuit between inferiors and superiors, it shows the superior's intellectual sensitivity and empathy, especially if they try to foster the inferior's knowledge or understanding. As such, we do not need to engage an epistemic superior in dialogue to agree with them. Just as we can simply echo a speaker's assertion, we can (e.g.) read what they have claimed in a publication and endorse that without any attempt to scrutinize it.<sup>339</sup> The upshot is that we can still gain more from actually reflecting on testimony and thoroughly scrutinizing it and not merely taking it at face value.

I have argued that we do not need any natural inquiry stopper to maintain the importance of agreeing with our epistemic peers, and that not committing to any specific account of the inquiry stopper need not diminish the importance or significance of agreeing with epistemic superiors. However, again, the scenario of agreeing with a regarded epistemic superior is not typically discussed. Plenty of attention is paid to whether subjects ought to defer to the testimony of agents they trust or regard as superior (or both), but the subject in such cases is not necessarily or explicitly an epistemic inferior but are typically just uncertain about what to believe about the matter at hand. Granted, inferiors can and sometimes will be uncertain in just this way, but just as Maura Priest regards epistemic standing and competence to come apart,

---

<sup>339</sup> E.g., of the many US citizens who claimed that Donald Trump won the 2020 presidential election, at least some would have heard the claim indirectly, not from Trump himself but from media reports or third parties asserting it.

surely so too do degrees of doxastic certainty. We might be poorly placed to make a judgment (epistemically inferior), more likely to have false beliefs than true (incompetent), but this need not mean we are uncertain or lack confidence in what to believe. Even epistemic superiors who have a good track record can have such doubts, so the phenomenon is not unique or exclusive to epistemic inferiors.

What does any of this tell us about how to respond to agreement with epistemic superiors? Again, the intuitive response might be to be reassured by it in general, but otherwise to acknowledge it and move on. I am less convinced by this, and instead contend that, at least sometimes, we ought to be willing to challenge agreement and consensus. Imagine a version of the Restaurant Check case in which I am generally poor at making quick calculations and miscalculate at least as often as I get things right, whereas you are far more reliable and consistent with mental arithmetic. Now suppose that in this instance, you and I calculate our shares of the bill and we both come to the same conclusion, with neither of us making any absurd calculations that are greater than the total bill or doubting that we have got it wrong. Agreement over a case as simple as this, given our asymmetry to the few factors that contribute to our judgment, will likely come as a pleasant surprise but otherwise be of little consequence or significance. That is, unless the bill or party of diners is incredibly high, or all the diners are particularly poor at arithmetic.<sup>340</sup> We could easily recalculate if we drew different conclusions or if one of us thinks we have both made a mistake.

Now consider a more complicated example of agreement between asymmetric agents: imagine that I am a junior economist just starting out, whereas you are highly experienced and particularly adept at predicting fluctuations in interest rates. Imagine that we both independently make our predictions about how the interest rates are going to go over the next financial quarter and we both reach the same conclusion. Compared to the Restaurant Check, there are many more factors to consider when making a prediction, the outcome has far more significant implications, and we are making a prediction about an as-yet unsettled matter rather than determining a fact that has already been determined. As such, any inferior in such a scenario will likely be far more surprised at independently reaching the same judgment as their superior. Complex cases matter more because they are i) less common, ii) more labour-intensive and perhaps iii) require specialised understanding, but agreement with a superior in complex cases likely tells us something most other instances of agreement do not. That is, that we have

---

<sup>340</sup>There is an important difference between agents who use the wrong methods but get the right answer by sheer luck, and those who use the right method but arrive at the wrong answer. The latter agent clearly understands the methods but has likely made a simple error, whereas it's not possible to determine the extent of the former agent's understanding.



reasoned correctly with our available evidence but, more importantly, that we have demonstrated some degree of understanding or that our judgment is not merely the result of a lucky guess. Many topics are within networks of deep and highly complex subject matters, so it is far from certain that our superiors are always in full and cohesive agreement with each other, on all related matters in a domain, let alone that they are correct in their shared judgments. Again, epistemic superiority and epistemic perfection are far from identical, no expert or authority is guaranteed to get things right every time, but we stand to gain more from challenging them, even just for the sake of clarity and double checking, than if we let it lie or ignore it altogether.

### 8.5: Disagreement with Epistemic Inferiors

I have considered scenarios in which the subject occupies the position of the epistemic inferior and I now turn to scenarios from the opposite perspective, beginning with instances of disagreement with an epistemic inferior. It may be tempting to maintain that epistemic superiors may permissibly *disregard* disagreement with superiors as irrelevant and mundane, or to dismiss disagreement with epistemic inferiors as a knee-jerk response when one occupies the position of the superior. However, this judgment is too quick and dogmatic, says Priest; instead, our response to disagreement with epistemic inferiors will depend on the kind of inferior involved:

We can then see that disagreement from competent inferiors can give superiors reason to adjust credence and reevaluate evidence, and hence that epistemic inferiority alone is insufficient grounds for dismissing opinion. Next come more nuanced difficulties. Disagreements between superiors and *incompetent* inferiors should sometimes boost the superior's justification. On the other hand, *agreement* from incompetent inferiors can defeat justification. Either way, inferior opinion carries epistemic weight. But, this still fails to cover all ground; uninformative inferior disagreements remain.<sup>341</sup>

I have already argued that intellectual vulnerability accompanies epistemic inferiority. I now contend that certain *prima facie* responsibilities follow from epistemic superiority and are made salient by disagreeing with inferiors, along with the lingering risk of epistemic injustices and harms. It may seem mundane or irrelevant if we disagree with someone we regard as either poorly placed to make judgments in general or unlikely to be correct on a relevant matter at hand. However, in doing so, we may be doing an interlocutor a disservice and harming them in their capacity as a knower and a source of knowledge, evidence, or belief. Therefore, I argue that

---

<sup>341</sup> Priest, M. 2016; 264.

any epistemic superior can and should aid inferiors to make sense of things by offering more than just their assertions or testimonial beliefs, they should be sensitive to the needs and deficiencies of epistemic inferiors, and the risks of epistemic injustice and harm. As with the previous sections, I take each aspect of the concern in order.

Firstly, we might think disagreement with our epistemic inferiors is irrelevant and unlikely to hinder or undermine our view, that it will present no argument or evidence we have not yet considered. Recall that Maura Priest discusses interactions with epistemically inferior agents, contending that we have reason to take the opinions of competent inferiors seriously, but also that we should also heed the opinions of inferior speakers who are not arbitrarily incompetent. For some, there may be a clear and simple upshot: inferior speakers should concede or suspend judgment, whereas superior speakers may remain steadfast. Bryan Frances endorses just this, describing a situation where one speaker is clearly superior in every factor that determines epistemic standing (the “Disagreement Factors”). He calls the judgment that obviously superior speakers may retain their beliefs the “Superior Rule”:

**Superior Rule:** If, before the discovery of disagreement, you thought that a certain person is a genuine expert on the question as to whether belief B is true, you are definitely not an expert, she surpasses you in some Disagreement Factors, and you do not surpass her in any Disagreement Factors – so she is your epistemic superior on the question – then, upon realizing that the two of you disagree, you should adopt her view on B or at least suspend judgment (and she should keep her belief).<sup>342</sup>

Again, this might appear intuitive in some cases, but an expert or some other epistemically superior speaker remaining steadfast following a disagreement with an inferior speaker need not mean disregarding the opinion of that inferior. At the very least this should not be assumed without due consideration, certainly not for all cases. Occupying the position of the epistemic superior, I argue, means having certain *prima facie* duties to one’s epistemic inferiors. This goes beyond merely inquiring for whatever reason, and those duties maintained by John Stuart Mill regarding the responsibility of political authorities to uphold freedom of expression for all. Mill argues that occupying a position of authority and refusing to hear an opinion or actively censoring it and hindering others’ free expression falsely and bombastically presumes one’s own infallibility.<sup>343</sup>

---

<sup>342</sup> Frances, B. 2014; 55-6.

<sup>343</sup> “All silencing of discussion is an assumption of infallibility. Its condemnation may be allowed to rest on this common argument, not the worse for being common. Mill, J.S. [1859] 2015; 19.

While free speech is not the same as an entitlement to an audience, I am sympathetic to Mill's position, I contend that it is counterproductive to disregard the opinion of our epistemic inferiors and make no effort to advise, guide, or correct them. However, I have in mind separate, *prima facie* other regarding duties to actively aid other agents and support their intellectual growth; *these* duties are derived partly from the various aims of inquiry, but chiefly from our mutual intellectual interdependence. A clear example of this can be ascertained by imagining students in a mathematics classroom. There will quite likely be a diverse range of competences and degrees of theoretical understanding between the students, such that some will better and more easily comprehend certain questions or subjects than others. In cases like this, it matters much less whether a student actually gets the answer to any given question correct, and more that they can demonstrate their understanding of the relevant theory, which is why many exam questions in mathematics have marks available for showing how one reaches their answer. It is far easier for a superior to correct a simple mistaken application or extension, whereas if a student does not comprehend the theory and merely gets correct answers from lucky guesses, guiding or developing their understanding may prove to be a far greater challenge.

I contend that aiding the development of inferiors are among the responsibilities that epistemic superiors bear, both for its own sake and as a form of intellectual progress. Progress can either come from the top down by identifying new questions and increasing the knowledge (and therefore knowledge gaps) in the domain, or from the bottom up by improving the epistemic standing of inferiors. If we can aid others intellectually and improve their epistemic standing, then we *prima facie* epistemically ought to do so, even if it means that we sacrifice something we value but *especially* if we face no great cost for doing so.<sup>344</sup> Again, intellectual asymmetries and epistemic interdependence are facts of the matter; we are all someone's superior and someone else's inferior, so we have reason to support one another which at least sometimes confers a duty or responsibility. If we disagree with the judgment of someone we regard as our epistemic inferior, we should not dismiss it but instead engage with them by considering their judgment, even if we think it is mistaken or we do not like it all that much. In other words, we should check that either of us actually has things right, that we understand each other correctly, and we should each clarify our position if we do not, explaining how we reached it and why we think the other has gone wrong.

---

<sup>344</sup> Casey Johnson (2020) offers a care-based account of obligations very similar to what I have in mind. She also posits a distinctly epistemic vulnerability but identifies it in a general notion of epistemic interdependence and the need for certain epistemic goods, rather than the epistemic inferiority I am discussing here. Because no one can acquire every necessary epistemic good by themselves, any dependence one has on others who are well-positioned to meet our needs, generates non-contractual obligations for those who are well-placed.

Moreover, I have raised a familiar concern that in at least some instances, it would commit an epistemic injustice or harm to disregard a speaker's testimony, experiences, or claims, so the mitigation and suppression of such injustices are also plausibly among the duties that apply (more) to epistemically superior agents.<sup>345</sup> Priest has in mind the inquiry-related or intellectual issues of disregarding inferior disagreement but says little about the possibility that doing so might be epistemically unjust or harmful. The sort of scenarios I have in mind are like the following, which describe both a general and a specific subject matter:

**Susan's story:** Susan has been profoundly deaf since childhood. She is a hearing aid wearer, and likes to use induction loops built into some public spaces such as theaters and cinemas, to help cut down the background noise that can make hearing speech very difficult. But this depends on the building having an induction loop fitted and properly maintained. Like many other induction loop users, Susan frequently finds that the advertised loop system is either working poorly or not working at all. Almost as often, she then has the experience of making a complaint about it only to have the problem denied. ("Of course there is a loop. Of course it's switched on, and of course it's working.")

If she persists, she is often then met by the suggestion that she doesn't know what she's talking about: Susan, who's been wearing aids for 45 years, is told that she misunderstands how to use the equipment properly, or asked whether, perhaps, she's just forgotten to switch on her hearing aid.<sup>346</sup>

In this case we have the general matter of using hearing aids and induction loops, and the specific question of whether *this* loop is working. Susan is epistemically superior on the matter of using hearing aids and induction loops and of regarding her own personal experiences, since she has decades of experience on the matter. However, she falls short of knowing whether *this* specific induction loop is functioning and can only hold suspicions or beliefs based on educated guesswork, so she *is* an inferior in this respect. While her enquiry about the hearing loop may be read or regarded as a disagreement, I do not think she has committed any harm. In fact, I contend it is just the opposite, I regard her as being the wronged party because of her clear

---

<sup>345</sup> To be clear, I contend that we *all* plausibly have *prima facie* duties to prevent injustices and harms of all sorts, but the duty falls more pressingly on those who are in a better position to do so; in the case of epistemic injustices, those who occupy the position of the superior, either in the intellectual or socio-economic sense, carries the lion's share.

<sup>346</sup> Scully, J.L. 2018; 1.

familiarity with the technology and her experience enabling her to clearly identify when something is not functioning.

Things are a bit clearer in the following case, wherein we have the general subject matter of medicine and the specific matter of an ill person's health, wellbeing, and pain.

An ill person can suffer testimonial injustice in one of several ascending ways. At the simplest level, ill people's testimonies can simply be ignored by healthcare professionals, perhaps being heard but neither acknowledged nor considered... Patient testimonies may be heard and acknowledged, but judged to be insufficiently articulate, and so once again excluded from epistemic consideration... Or those testimonies may be acknowledged but *subordinated to the authority of health professionals*, so what the patient has to say is heard, but coupled to the conviction that it is the view of health professionals (and in particular physicians) that is authoritative or primary. So the ill person may be judged to be a useful informant but not a participant in the collective practice of interpretation and understanding of their medical situation.<sup>347</sup>

The health professionals – doctors, physicians, surgeons, and nurses – clearly occupy the position of epistemic superior since they all have extensive training and experience in being able to understand and treat various illnesses and injuries. However, again, there is a specific question of what an individual patient feels and reports, of which the patient is clearly the epistemic superior. Assuming our experiences of illness and pain are exclusively and uniquely our own, the professionals cannot confirm any patient's report as legitimate, but they can help the patient to understand and make sense of it, even if they do not agree with a patient's self-diagnoses. Alternatively, the professional can disagree with the patient's reports and disregard it, perhaps believing the patient is exaggerating the severity of their illness or injuries, or that they are prone to hypochondria.

Ignoring a patient's testimony among the most common examples of testimonial injustice, one of the epistemic injustices Miranda Fricker posits. In conversation with each other, speakers can either regard each other with an *excess of credibility* and cause them to develop "an epistemic arrogance that a range of epistemic virtues are put out of his reach, rendering him closed-minded, dogmatic, blithely impervious to criticism, and so on."<sup>348</sup> Or we can regard each other with a *deficient credibility*, holding false beliefs about their expertise or motives. Underestimating a speaker can be accidental, in which case the moral culpability may come apart

---

<sup>347</sup> Carel, H. and Kidd, I. 2014; 531-2. My emphasis.

<sup>348</sup> Fricker, M. 2007; 20.

from the epistemic injustice, but it can be intentional as well. Fricker maintains that there are hermeneutic injustices that are the result of systematic, structural flaws such as embedded prejudices or biases:

From different points in history one might draw on many depressing examples of prejudices obviously relevant to the context of credibility judgement, such as the idea that women are irrational, blacks are intellectually inferior to whites, the working classes are the moral inferiors of the upper classes, Jews are wily, Orientals are sly... and so on in a grim catalogue of clichés more or less likely to insinuate themselves into judgments of credibility at different moments in history.<sup>349</sup>

While I do not think what is happening with Susan or the professional-patient case is necessarily the result of structural biases, we can certainly imagine related cases with structural causes.

The point I want to emphasize is that, while not all instances of disagreement between epistemically asymmetric agents are necessarily harmful, we should be cautious and acknowledge that there is at least a possibility. Disagreement with an epistemic inferior may merely coincide with an epistemic injustice or harm, sometimes the disagreement may cause such harms, and other times disagreement with an inferior agent may involve no injustices or harms at all. However, even if we ascertain that we will commit no harm by disagreeing with an epistemic inferior, we are not thereby entitled to completely disregard our epistemic inferiors' opinions. Occupying the position of the epistemic superior carries other regarding responsibilities to epistemic inferiors, mitigating epistemic injustices and harms may be reasonably regarded among these duties.

#### 8.6: Agreement with Epistemic Inferiors

I have argued that occupying the position of the epistemic superior confers *prima facie* other regarding epistemic duties and accountability towards epistemic inferiors. This is consistent with many arguments endorsing an imperfect duty to object, that we at least sometimes have a duty to voice our dissent or disagreement.<sup>350</sup> This is not to say that, for example, members of marginalised groups always have a duty to correct racists, sexists, homophobes, etc., but that they have just such a duty subject to certain constraints such as whether it is safe for the dissenter to do so. Disagreeing with epistemic inferiors should not be dismissed as insignificant but instead embraced as an opportunity to limit or prevent epistemic injustices and harms (and arguably to

---

<sup>349</sup> Ibid; 22-3.

<sup>350</sup> See, for example, Lackey 2020; Johnson 2018a, 2018b;

promote epistemic justice also.) I now consider the final scenario, that of agreement with an epistemic inferior. I have in mind any instance of agreement with a speaker who we consider substantially less well placed to make a judgment than ourselves, but also with whom we would least expect to agree with, perhaps because we perceive them as less intelligent or worse equipped than us. Agreeing with inferiors may challenge how we perceive ourselves, calling into question the extent to which we are genuinely intellectually superior. After all, if we regard an interlocutor as our epistemic inferior in some or many respects, we plausibly regard ourselves as more likely to get things right regarding the matter at hand. If the divide between the superior and inferior is minor, then this worry is perhaps less salient and prominent, whereas if the divide between them is substantial then it would naturally raise worries for the epistemic superior. We might not just think that we have got things wrong on this occasion because our inferior has reached the same conclusion, but it may cause us to think that we aren't in fact as well placed to judge as thought/think we are. Given all this, we might ask, should we think we are actually epistemically superior and, if so, why?

Regardless of the worries it might incite, the upshot of agreeing with an epistemic inferior is no different than that of disagreeing with inferiors, and arguably disagreeing with equals: we should not be so concerned by the prospect that we abandon all our convictions and suspend judgment altogether. Again, epistemic superiority is not epistemic perfection, and no one can be a truly ideal agent. Those who are more likely to be wrong can get things right just as those more likely to be right can get things wrong, but we should not be so put off by the prospect of our own errors that we avoid forming beliefs at all. Instead, we should engage in inquiries with an open mind to getting things wrong and learning from those mistakes and recognising that agreement with an epistemic inferior sometimes does not mean that all parties are mistaken, but it should nonetheless be double checked.

Recall that Priest believes the opinions of arbitrarily incompetent speakers are insignificant and have neutral epistemic import and bearing, meaning there is no strong reason to consider or disregard them. However, the opinions of consistently and maliciously incompetent speakers must be flagged because malicious inferiors deliberately deceive or have beliefs that are directed by underlying biases; ignoring or overlooking their assertions risks some kind of danger:<sup>351</sup>

Prejudice begets malicious inadequacy. Imagine superior A disagrees with inferior B over the merit of initiative I. Moreover, B is a Klansman. A could learn more from this disagreement than he could from inadequate inferiors. Because B's prejudice leads him to

---

<sup>351</sup> See Priest 2016; 279-80.

consistently support initiatives A opposes and oppose initiatives A supports, when A learns of B's dissent, A should be reassured that he is taking the correct stance...

On the other hand, A would learn little from inadequate inferiors. For instance, disagreement from low-information voters with arbitrary political positions would speak neither for nor against the merit of A's stance on I.<sup>352</sup>

Priest regards agreement with speakers who are consistently epistemically inferior as higher order evidence that speaks against our shared belief and our respective belief-forming processes. Such agents are malicious and track falsity, their opinions are no more reliable than any fabricated or satirical piece in *The Onion* or *The Daily Mash*. If we recognise or become aware of this maliciousness, then we are in a position to avoid any epistemically dangerous situations such as agreeing with them. If instead we agree with maliciously inferior agents and only come to learn of it afterwards, this gives us reason to reduce the degree of confidence we have in our belief because it strongly indicates that we have reasoned poorly. Speakers who arbitrarily form their beliefs, such as those disposed to bullshit and those who are genuinely ignorant, are not tracking falsity, so agreeing with such speakers does not generate the same kind of evidential defeat.<sup>353</sup>

However, I contend that this diagnosis does not get things entirely right. I do not agree that all bullshitters are necessarily inferior in the respect that Priest regards them as. Presumably there can be speakers who have superior cognitive capacities or evidence, but could still be disposed to bullshit, to make arbitrary assertions that do not track truth or falsity, and I would say the same of devil's advocates, contrarians, and conformists. Bullshitting is as much about the sincerity of a speaker's belief as it is about their motivation, and anyone could be in the dark regarding either matter. These two kinds of uncertainty should, I argue, make us more wary of known bullshitters, and more careful to ascertain both the motives and sincerity of a speaker, at least more so than we need to be about those we know are well meaning but insufficiently equipped to make a judgment. I would say the need for caution is most salient when we are engaging with a stranger. Some degree of familiarity with a speaker will offer at least a prospective idea of one's standing and intentions, but it is far from clear if the beliefs and assertions of strangers ought to bear on us at all.

Moreover, I do not agree with Priest's claims regarding epistemic maliciousness. While I grant it is a bad sign to find that we agree with someone who is deliberately deceptive, Priest sets them apart from bullshitters and genuinely ignorant agents, which I believe is a mistake. Anyone occupying the position of the epistemic superior would (and plausibly should) be able to

---

<sup>352</sup> Ibid; 281.

<sup>353</sup> Ibid; 279-81.



recognise when a speaker is trying to deceive them, at least most of the time. I contend, firstly, that the prospects of agreeing with them are surely fairly slim and, secondly, that a speaker's maliciousness counts for more and surely stands out more than their incompetence or the falsity of their beliefs. Priest thinks it is neither here nor there if we agree with a bullshitter or a genuinely ignorant agent because their beliefs are completely arbitrary, whereas agreeing with a malicious deceiver should motivate us to change our belief.

While it is bad to agree with someone we regard as worse off who is malicious, I argue instead that it is much worse to agree with speakers who are insincere and apathetic bullshitters or who are consistently ignorant. If we know beforehand that someone is deceptive in nature, as Priest says, we have reason to disregard their opinion. Epistemic superiority should mean we are equipped to engage with all kinds of inferiors, but it does not ensure any perfect safeguard against misinformation or deception. If we naively accept an inferior interlocutor's assertion without giving it due consideration or scrutiny, we have clearly failed to adhere to due process, but we can surely change our minds or suspend judgment. Agreeing with a genuinely well-meaning but ignorant speaker is a far stronger blow to our supposed superiority, and if we agree with someone who we later find out is a bullshitter who does not care about the truth, we should be willing to reassess our beliefs as standard. Once again, we can anticipate these implications by simply striving for clarity and disambiguation. One should not be led astray by a speaker's charm or charisma, since one could very easily present themselves as knowing more than they do about a subject or topic. Nathan Ballantyne describes a related phenomenon that he calls *epistemic trespassing*, which occurs when a speaker entreats on a discipline beyond the areas of their expertise and falsely professes to have expert knowledge, understanding, and experience:

Epistemic trespassers are thinkers who have competence or expertise to make good judgments in one field, but move to another field where they lack competence – and pass judgment nevertheless. We should doubt that trespassers are reliable judges in fields where they are outsiders.<sup>354</sup>

In other words, if a speaker can be convincing enough, their epistemic inferiority might fly under the radar, and they might more easily be able to disseminate false beliefs and misinformation. An epistemic superior would need to be sufficiently discerning or perceptive and informed about the subject matter, if not especially untrusting of others, to be able to spot a trespasser in advance. Some agents will be sufficiently positioned (i.e., a strong enough superior) to do this without

---

<sup>354</sup> Ballantyne, N. 2019; 1.

much issue, but if one only has a limited degree of knowledge, experience, competence, and evidence, such that many will regard them as only slightly superior, they may be far less discerning and more easily misled.

### 8.7: Closing Thoughts

In this final chapter, I have reflected on the nature of epistemic asymmetry, arguing that vague references to ‘non-peer agreement’ and ‘non-peer disagreement’ are not enough, and that these labels are vague and imprecise. Several factors contribute to our regard of a speaker’s epistemic standing and whether we perceive someone as our epistemic superior or inferior. The disagreement literature offers very limited discussion of agreement and disagreement between inferiors and superiors, focusing generally on cases involving experts and groups. The epistemology of testimony offers some further in-roads, but again the consideration is limited.

While I have partitioned scenarios and kept agreement and disagreement cases separate, I have only offered general arguments about what I believe are the implications of agreeing and disagreeing with superiors and inferiors, broadly construed. Further work must be done, specifically elaborating more on the various applicable notions of superiority and inferiority, and directly examining the consequences of agreeing or disagreeing with each kind of superior or inferior. I have discussed three claims that I believe are not simply unique or limited to one partitioned scenario but can in fact be inferred from most or all of them. The first of these was the urgency of collective and collaborative inquiries, and the need for asymmetric speakers to work together rather than take agreement or disagreement at face value and then go their separate ways. Clarifying any ambiguities or uncertainties that are present because of disagreement and double checking that agreement is genuine and accurate are, I argue, rational and intellectually productive responses that avoid indefinite dogmatism or suspension. Being prepared to challenge our interlocutors, regardless of where we perceive them as standing and whether we agree or disagree, cultivates and exhibits intellectual virtues and typically brings about greater benefits than merely accepting a situation.

Secondly, it is important to recognise or determine intellectual character traits and motives, and that epistemic superiority is neither perfect nor ideal epistemic agency. We should aim to find speakers who recognise that intellectual capabilities can be diverse, display sensitivity to the needs of all participants in a domain or discipline and who are both able and willing to help us. Likewise, we ought to encourage those who occupy the position of the epistemic inferior towards epistemic improvement and self-development, cultivating intellectual curiosity and creativity, as well as critical thinking and argumentative reasoning. These are all indicative of

prima facie epistemic duties we possess in virtue of occupying the position of an epistemic superior or inferior, and acknowledging this when we agree or disagree will ensure greater epistemic civility, humility, and fraternity.

Thirdly, the two previous points indicate that progress in inquiries and dialogues is multifaceted. There can be a top down approach to progress where those who occupy the position of the superior attempt to produce greater knowledge or justified belief, offering new questions and true beliefs regarding those questions. Alternatively, there is a bottom up approach to progress wherein the distance or gap between those who occupy the positions of the superior and inferior is narrowed, and improving the capabilities of those who are less intellectually well off takes precedence. Following on from my observations about motive, epistemic superiors can be self-regarding, and focus only on advancing a domain insofar as it improves their reputation or material gains. Or they may be invested in the development of a domain and/or the needs of their fellow inquirers, which I contend should take priority in high-consensus domains where new knowledge is a rarity. Regardless of where we regard ourselves as standing in a domain, I believe it is entirely possible to advance it, and I can think of no reasons not to try.

I do not think the conclusions I have drawn are controversial or implausible, but they almost certainly require further consideration. Partitioning scenarios in the way I have done provides us with a good starting point, but it certainly does not offer conclusive evaluation of agreement and disagreement beyond the extant discourse. I will reiterate that the epistemology of disagreement focuses chiefly on rare, idealised, hypothetical scenarios between disagreeing epistemic peers, neglecting cases of peer agreement and more realistic and grounded scenarios of agreement and disagreement between epistemically asymmetric agents. Greater accommodation of agreement and non-ideal instances need not be an especially radical change, but it is certainly a necessary one.

Chapter 9: An Afterthought on Agreement, Disagreement, and Testimony

This thesis addresses agreement and disagreement, a pair of phenomena which are highly pervasive and common. Their commonality may make them appear simple, but they are in fact incredibly complicated, multifaceted, and continually shifting. There are many accounts of what it is to agree or to disagree, and I have argued that each succeeds at capturing only part of the picture, sometimes conflicting with each other if not generating disagreement over what agreement and disagreement are or involve. No one, to my knowledge, has offered as detailed an analysis as the one I have provided. I have maintained that there is very little that we can say about the nature of agreement and disagreement that covers all bases, except that they are largely indefinable in singular, precise terms. I have also argued that the epistemology of disagreement has over-emphasized idealised cases of peer disagreement, neglecting peer agreement and agreement and disagreement cases between a variety of non-peers. I want to close this thesis by outlining a few thoughts which tie together these two themes, and to suggest that agreement and disagreement are at least partly or occasionally grounded in testimony.

The epistemologies of testimony and disagreement both share at least one guiding question: how ought the speech acts, beliefs, or assertions of others impact or guide oneself and one's own attitudes or beliefs? When we agree or disagree, we typically assume, however minimally, that we are each sincerely asserting or expressing our attitude, and thus offering what can plausibly be regarded as testimony in some informal sense. Why, then, have the two debates around disagreement and testimony near enough proceeded as if the two sets of phenomena do not interact, overlap, or engage with each other? We need not think that testimony is necessary or sufficient for all kinds and cases of agreement or disagreement (at least not if testimony is envisaged as a verbal utterance or speech act) but there is also something unusual about maintaining they are entirely separate from testimony. There is, I contend, at least a prospective causal or normative connection between acknowledging purported testimony and agreeing or disagreeing with a speaker or agent. The nature of this prospective link or relation between testimony, agreement, and disagreement is one pertaining to *acknowledgment*. More specifically, I believe that at least some instances of agreement and disagreement involve the acknowledgment of *purported* testimony, of a speaker's intention to offer testimony or their belief that they are offering testimony. This need not commit a speaker to hearing or deferring to a speaker's utterances or speech acts, nor does it give them a license to disregard the content of testimony, but it means that a hearer recognises a speaker has a particular intent, aim, or desire. I want to make the case that there is at least a prospective relation between testimony, agreement, and disagreement. I will begin by spelling out the claim in greater detail, before suggesting some

questions or concerns that one might anticipate, as well as an epistemological implication that follows from it.

### 9.1: Acknowledgment and Purported Testimony

The following claim is, I argue, bold, if perhaps relatively simple:

ACKNOWLEDGE: To verbally agree or disagree is to *at least* acknowledge a speaker's purported testimony.

At least a few questions may arise from such a claim, such as what I mean by “purported testimony”, why I emphasise acknowledgment over affirmation, denial, or acceptance, and whether *uptake* plays any role or relevance. These are far from exhaustive, and there are other questions I believe arise for this claim which I will not address in detail here but raise towards the end as a point of further consideration. For now, we can simply take my claim as being that testimony plays *some* role in at least *some* instances of agreement and disagreement, and that it is strange to say that these two concepts or phenomena entirely come apart. In line with my earlier thesis arguments, I do not wish to say that testimony or acknowledgement of purported testimony is a necessary or sufficient condition. However, I will concede that testimony at least appears to be a noticeably salient feature, and agreement and disagreement can quite plausibly be conceived broadly enough that testimony becomes something of a default necessary or sufficient condition. However, I am not convinced that the relation can ever be considered so strong that testimony is necessary or sufficient for *all* cases. The question of whether there is any relation between testimony and agreement/disagreement has not, to the best of my knowledge, been discussed in either the testimony or disagreement literature but, if it has, the consideration has been sparing at best. The one exception comes from Nicholas D. Smith, who in discussing the peer disagreement discourse, proposes that the rational response to peer disagreement mirrors the rational response to any testimonial evidence, provided that it is within our degree of personal justification:

... the perception that a peer disagrees can be treated on a par with the appropriate epistemic response to testimonial evidence. I do not, of course, intend to claim that peer disagreement *just* is or *must always be presented in the form of* testimony... But even so, I claim, the way in which we appropriately assess the evidential value of peer disagreement

will follow the same approach as we apply to cases of testimonial evidence, since there, too, we assess the evidence provided by recognition of someone else's opinion.<sup>355</sup>

In other words, for Smith, peer disagreements need not necessarily be regarded as *grounded* in testimony, but any rational agent's response to them will be comparable. I extend the argument further by arguing that some instances of agreement and disagreement (whether with an epistemic peer or not) can be conceived of as a kind of testimony, or at least that they involve the recognition of attempts to offer testimony. Like Smith, I am careful not to vehemently rule testimony in as a necessary or sufficient condition for any instance of agreement or disagreement; I wish only to offer a mere suggestion and consider its implications.

The prospect of a relation need not be regarded as controversial and need not run contrary to the arguments I have made in the earlier parts of this thesis. My aim is simply to offer my intuition on the matter, and to sketch out some of the implications, all of which can be regarded as requiring further consideration and discussion. The first of the questions I raised above asked what I mean by *purported* testimony and how it may differ from (just plain) testimony or 'actual' testimony. In one of the earliest distinctions, Anthony Coady describes 'formal testimony' as testimony in some inquiry or commission, such as that offered in court proceedings, the invitation to regard one's claims or report as true because one says it is.<sup>356</sup> More specifically:

- a) It is a form of evidence.
- b) It is constituted by persons *A* offering their remarks *as* evidence so that we are invited to accept *p* because *A* says that *p*.
- c) The person offering the remarks is in a position to do so, i.e. he has the relevant authority, competence, or credentials.
- d) The testifier has been given a certain status in the inquiry by being formally acknowledged as a witness and by giving his evidence with due ceremony.
- e) As a specification of (c) within English law and proceedings influenced by it, the testimony is normally required to be firsthand (i.e. not hearsay).
- f) As a corollary of (a) the testifier's remarks should be relevant to a disputed or unresolved question and should be directed to those who are in need of evidence on the matter.<sup>357</sup>

---

<sup>355</sup> Smith, N.D. 2015; 143.

<sup>356</sup> Coady, C.A.J. 1994; 27.

<sup>357</sup> Ibid; 32-3.

These conditions are intended to demarcate formal testimony from ‘natural testimony’, which is just any ordinary testimony encountered in everyday circumstances, such as by giving directions, reporting an incident or affirming an account. Testimony in its natural, typical sense is any speech act or statement that *p* regarding some disputed or unresolved question, offered as evidence that *p*, from someone with relevant competence, authority, or credentials to state truly that *p* to someone in need of the relevant evidence. Natural testimony meets the conditions a, b, and f, and, to some extent, c, above, but it is not assumed to be firsthand, and the testifier is not acknowledged in any formal position of authority or witness. For Coady, then, testimony and the act of testifying involves one *vouching* for a connection between how one says things are and how they are actually, whether within a ceremonial context or not.<sup>358</sup>

Testimony, so understood, is putative or *potential* evidence. I contend that, when a speaker makes an utterance, they will typically have the desire or intention that it be perceived as such (setting aside devil’s advocate cases or some performance of personae.) However, this intention being present need not mean that a hearer or audience will perceive a speech act as testimony, as substantial information, evidence, or as a reason to believe. A speaker treating a speech act as testimony in the way Coady describes or acknowledging the content of their utterance as testimonial content is different from recognising that the speaker *believes* they are offering testimony. In this sense, recognising a speech act as testimony requires the *acknowledgment* of the speaker’s intention or desire; it need not mean accepting or believing the content of their testimony, so one can acknowledge that a speaker intends to offer testimony without agreeing with what they actually say, or acting on it. If I tell you the bus station is two miles down the road after a sharp right, you might grant that I genuinely believe it without being compelled to endorse it yourself or to follow my directions. Many ordinary cases within the epistemology of testimony take it as standard that the speaker provides *bona fide* testimony and that the audience or hearer perceives it as such, but this need not always be the case. If a hearer does not perceive or accept a speech act as testimony, they may still acknowledge the speaker’s attempt, desire, or intention to provide it. *Purported* testimony, I argue, is what is offered by the speaker regardless of whether a hearer or audience regards it as actual testimony. A hearer may have reason to avoid agreeing immediately or without consideration, perhaps to reflect on what the speaker has said or intended.

## 9.2: Acknowledgement and Uptake

---

<sup>358</sup> Ibid; 43.



I emphasize *acknowledgment* of purported testimony rather than other attitudes such as affirmation, denial, or acceptance. We might think that agreement and disagreement require both the acknowledgment of the attempt at testifying and either the affirmation or denial, respectively, of its content. In keeping with my earlier arguments, this is clearly true for some cases of agreement or disagreement, but it need not be true for them all, certainly not so for those cases that involve non-doxastic attitudes or states. If I were to contend that all instances of agreement and disagreement involve affirming or denying testimonial content, it would be no different to endorsing acknowledgment of an utterance as testimony. Not only is this unrealistic, but it raises some contentious implications. We surely recognise that a speaker thinks they are offering testimony (or perceives themselves as doing so), for to do otherwise would simply be dogmatic through and through. If our attitude towards the contents of a speaker's speech act comes apart from this recognition, then we can have agreement or disagreement without affirming, endorsing, or denying the contents. Moreover, if we recall the account offered by Max Kölbel wherein one agrees or disagrees with a speaker *in content* by making utterances with identical or contrary pragmatic contents, we can imagine at least some cases where such agreement and disagreement come apart from acknowledgment of purported testimony. Suppose a speaker and I do not speak the same language and we must rely on a translator; they might say something that I do not understand, and the translator may simply give me an affirmative response without explaining what they have said. We may verbally agree in content, but I think it is implausible to say that I have fully acknowledged their testimony or attempt at providing it. Alternatively, one may simply say that at least some instances of agreement and disagreement have a stronger requirement for the comprehension of a speaker's attempt at testimony. To agree or disagree with a speaker, is not enough to simply acknowledge that that speaker intends to offer testimony; rather, one must also be aware of and comprehend the contents of that purported testimony.

Some of the accounts of agreement and disagreement I have explored (and conceded as genuine) specify or prioritise non-doxastic states, so it is entirely possible for agreement to come apart from belief and disagreement to come apart from disbelieving what a speaker says. In other words, I am saying *here* that at least some instances of non-doxastic agreement or disagreement involve the recognition that a speaker intends or is attempting to provide testimony – the acknowledgment of purported testimony. I believe that we can say the same of at least some doxastic agreement and disagreement, but the non-doxastic case makes the idea easier to motivate. Recall the case of credal disagreements and the atheist/theist/agnostic case. Suppose we said that agreement obtains between two speakers if and only *iff* they share the same belief with the same degree of confidence. That would suggest that, for example, two atheists who

believe that God does not exist to differing degrees decisively do not agree. But this is absurd, unrealistic, and an overly narrow conception of agreement, certainly far from unifying. Nonetheless, it is surely the case that atheists who differ in the strength or degree of their disbelief, will regard each other as purportedly testifying, and that they will say the same of theists and agnostics despite more obviously disagreeing with them.

The third question I raised in response to ACKNOWLEDGE was whether *uptake* had any significance, by which I mean the effect or force that J.L. Austin contends arises as a result of certain speech acts. When a speaker makes an utterance with the aim and intention of using it in some way, such as by expressing an attitude, making an assertion, appraising or expressing some emotional appeal, they perform an *illocutionary* act. For Austin, the success of the performance of such acts depends on the hearer's uptake, or whether the desired effect is achieved. If an audience does not recognise or understand that I've made a particular kind of utterance, then it does not have the force I intend for it, whereas if they do anticipate, recognise, or acknowledge the intention and act accordingly, my speech act has gained uptake.<sup>359</sup> Austin regards affirmation, denial, conjecture, and concession all as the same kinds of speech act, but crucially he also groups them with both *testifying* and *agreement*. He describes them all as expository acts "involving the expounding of views, the conducting of arguments, and the clarifying of usages and references".<sup>360</sup> This seems to bode well for ACKNOWLEDGE, since the kind of acknowledgment or recognition it picks out seems at least somewhat similar to Austin's idea of uptake.

However, it is not clear that illocutionary force always and necessarily requires uptake to succeed or hold. Rae Langton raises a worry that certain speech acts might fail to count as illocutionary acts on Austin's account if the uptake is not granted, or the force of the act is not recognised by a hearer. Langton has in mind the worrying moral implications that an act of refusal regarding sexual advances might fail to count as an act of refusal unless the relevant hearer or audience recognises it as such and acts upon them accordingly.<sup>361</sup> We have, then, Austin's claim that the force of an illocutionary act depends on a hearer's uptake against Langton's concerns that this has morally undesirable implications or consequences. I grant Langton's concerns and concede that not all illocutionary acts require uptake to stand, they may not even require an audience or hearer. But as to whether agreement and disagreement fall within the class of speech acts that require uptake or not, if they do then it seems that there is a rather limited range of cases available that we can count as legitimate or genuine. If uptake is not

---

<sup>359</sup> See Austin, J.L. 1962; 115-6.

<sup>360</sup> Ibid; 160.

<sup>361</sup> See Rae Langton's 1993 and Alexander Bird's 2002.

required, then ACKNOWLEDGE is false, but it raises a risk that too many cases are available which we could call agreement or disagreement.

Again, I am not convinced that we can narrowly specify what agreement and disagreement consist in, and I regard them as too complex and multifaceted to capture with necessary or sufficient conditions. But if uptake is not required for agreement and disagreement, which seems to be the case for at least some cases (recall the distinction between personal disagreement with a speaker's uttered belief or attitude and propositional disagreement with the attitude itself), then it seems that agreement and disagreement may well be too broad. ACKNOWLEDGE is false, then. That need not mean that there is no relationship between testimony, agreement, and disagreement at all though. I think a small revision is in order in light of these considerations. We can grant that some illocutionary acts can pass, even if a hearer does not recognise them, so a speaker may take themselves to be asserting, testifying, agreeing, or disagreeing, regardless of whether a hearer acknowledges this. Rather than saying agreement involve "at least acknowledging purported testimony", something along the lines of the following is more apt:

PART ACKNOWLEDGE: at least *some* instances of verbal agreement and disagreement involve the acknowledgment of purported testimony.

I contend that there are, in at least some cases, overlap between purported testimony and agreement or disagreement, or at least that testimony plays some kind of role. However, given the lack of consensus within the epistemology of testimony about whether testimony confers or carries warrant or justification, or whether it is a legitimate source of belief and evidence, PART ACKNOWLEDGE raises some problems that need to be addressed. I will turn to these now.

### 9.3: Epistemological Implications

In light of concerns about the need for uptake, I have made the claim that in at least some agreement and disagreement cases, participants (hearers or audiences) acknowledge a speaker's purported testimony, the intention or attempt to provide testimony. With this in mind, from PART ACKNOWLEDGE, I think something along the lines of the following results naturally:

TESTIMONY: If we are to say anything about the epistemic significance of agreement and disagreement, we must grant that testimony has epistemic significance or import.

This recognises that testimonial warrant is not universally accepted, that some maintain there are reductionist or pessimistic constraints on it that, if true, downgrade its importance. By the former, I mean any claim that testimonial justification and warrant are explicitly and exclusively reducible to perception, memory, or inference, rather than being an epistemically basic and independent source of justification.<sup>362</sup> By the latter, I mean something along the lines of Robert Hopkins's *Requirement* for moral testimonial warrant, or Philip Nickel's argument that moral agency requires consideration of moral principles or rules:

**The Requirement:** having the right to a moral belief requires one to grasp the *moral* grounds for it.<sup>363</sup>

Morality aims at guiding action rationally, i.e. *from* a recognition of the relevant moral requirements. A moral agent must be responsive to morality as such. One strong way to make this claim is to hold that it is impossible to act morally if one has never considered and adopted anything like a *principle* or *rule* with moral content... More neutrally, we may say that, in a morally good action, moral considerations must play a role, mediated by cognition or rationality, either in deliberation or in the formation of the psychological states which lead one to act as one does.<sup>364</sup>

The upshot of this is that, if testimony has limited or conditional significance, and one grants that at least some agreement and disagreement cases are a species of, or depend on, testimony, then the epistemic import of *those* kinds of agreement and disagreement are subject to just the same conditions or constraints. In other words, the normative status, philosophical significance, and appeal of agreement and disagreement (or rather any cases that involve testimony or something sufficiently similar) will depend entirely on the significance of that testimony. I contend that this is obvious, but by no means trivial. We need a compelling reason to consider testimony epistemically significant or philosophically interesting if we are to vindicate the epistemology of agreement and disagreement. As it happens, I am unconvinced that the arguments against testimony hold weight if we envisage the testimony in question as a case of agreement or disagreement, but I will not offer a detailed explanation here.

However, one might deny either PART ACKNOWLEDGE or TESTIMONY, or both. Denying TESTIMONY grants that agreement, disagreement, and related exchanges *do* involve

---

<sup>362</sup> See Lackey, J. 2006.

<sup>363</sup> Hopkins, R. 2007; 630.

<sup>364</sup> Nickel, P. 2001; 256-7.

testimony of a sort that is somehow immune to any pessimistic or reductionist constraints. This then raises a burden of proof for demonstrating how such testimony uniquely differs from any other kind of testimony which is subject to these constraints *and* still fits within agreement and disagreement. Otherwise, one must have a compelling argument against any and all reductionist or pessimistic arguments. Alternatively, one could deny PART ACKNOWLEDGE and maintain that agreement and disagreement do not involve any epistemically important testimony, but that we still need a better, either neutral or optimistic view, regarding testimony. This may be fine for a compromise or fallback view, but I still contend that it is strange to consider agreement and disagreement *completely* separate from testimony. Finally, rejecting both arguments yields two kinds of denial, one weaker and one stronger. In the weaker form, one simply denies that agreement and disagreement ever involve testimony or need to be reconciled with reductionism or testimony pessimism. In the stronger form, one adds in a stronger sceptical judgment that neither agreement, disagreement, nor testimony bear any epistemic import or philosophical significance whatsoever. Again, any such pessimist would have to show what agreement and disagreement depend on if not testimony, but also why none of these three phenomena bear on us or matter at all. I hesitate to call this a task that is realistic or achievable, let alone simple.

Perhaps I am being too idealistic, or there is something I have overlooked making my argument glaringly mistaken. I still find unusual, if not absurd, the potential sceptical conclusion that there is no relation or overlap at all between these two kinds of phenomena, but also for both to have received such substantial attention almost entirely separate from each other. Perhaps something like PART ACKNOWLEDGE and TESTIMONY have been implicitly assumed without argument, or simply dismissed as ridiculous. Nonetheless, I am yet to have found a persuasive and compelling argument that neither agreement nor disagreement ever involve testimony of any kind in any case. Again, both the epistemologies of testimony and disagreement have at least one shared guiding question at which they aim – what, if anything, is the import of other speaker's opinions or utterances – so it is all the stranger that the two have been considered almost entirely independently.

#### 9.4: Further Considerations

I have offered, I believe, at least a prospective account of a conceptual and epistemic relation between testimony and (some cases of) agreement and disagreement by way of acknowledgement of purported testimony. I have discussed some of the implications such an account raises and suggested a minor modification in light of it. Nonetheless, I believe there needs to be more rigorous examination of whether agreement, disagreement, and other related

exchanges or phenomena are indeed grounded in testimony or are species of testimony. Whether Austin is correct to group testifying and dis/agreeing together has received little to no attention, but if my brief discussion is even partway accurate then it has both reasons backing my intuition and implications that might call it into question. Ultimately, greater consideration is needed ascertaining what the precise relationship between testimony and agreement or disagreement is, whether it is causal or one of mere appearance and, if there is a causal connection, when it arises. Full elaboration requires more consideration of the central concepts and current literature, whereas I have only offered a cursory glance so far. However, I will conclude by suggesting further considerations and questions that arise from my discussion thus far.

First, there is a question of success, aside from considerations regarding illocutionary force and uptake, and whether a speaker needs to acknowledge agreement or disagreement for there to be agreement or disagreement (which obviously need not be the case for all cases but certainly is required for some). The acts of (merely) communicating, testifying, and agreeing or disagreeing may all have separate attached notions or intuitions regarding success, different criteria for what it is to succeed in performing these acts, and these likely need further elaboration. We might not want to say that a speaker has failed to communicate or testify just because a hearer or audience disagrees with the contents of their speech act. Moreover, it is not always clear what is required for a communicative act or an act of testifying to succeed, we might think that uptake is just one condition or requirement and that more, in fact, is needed. Indeed, communication, testimony, and assertions, may all fail for reasons that have nothing whatsoever to do with the speaker or hearer.<sup>365</sup> Relatedly, more could be made of the question regarding uptake and illocutionary force, especially regarding speaker intentions and implicature. It might be that a speaker only intends that the illocutionary force of their utterance is directed towards certain hearers or audiences, and that only certain indirect forms of speech are successful communications, but whether they facilitate agreement and disagreement is still an open matter.<sup>366</sup>

Second, given the various accounts of agreement and disagreement I have considered through this thesis, and my deliberately broad reading of them, it may be that different kinds of testimony may be relevant to different kinds of agreement and disagreement. If so, then this is all the better both for my contentions regarding agreement and disagreement *and* their (partial) basis in testimony. I have claimed that the type of speaker with whom we are engaging places some apparent constraints on the substantiveness or robustness of agreement and disagreement, and it seems entirely plausible that they may affect the relation between testimony and agreement

---

<sup>365</sup> See, for example, Carrie Jenkins' 2014., and Joey Pollock's 2020 and 2021.

<sup>366</sup> See, for example, Nicholas Asher and Alex Lascarides' 2001.

or disagreement also. We might find that testimony from the likes of devil's advocates, groups, or even three-party exchanges such as Jennifer Lackey's "Eavesdropping" case, or intrapersonal agreement and disagreement may make for interesting cases and possible obstacles for my account.<sup>367368</sup>

Finally, I have focused exclusively on the similarities between testimony, agreement, and disagreement, and have suggested some worrisome implications for the epistemology of disagreement, but there may well be a relation of support between the three phenomena. That is, it might be that the epistemology of disagreement can help us to make sense of the questions raised regarding testimony, rather than the other way around. There are important differences between testimony and agreement and disagreement that will need to be addressed as well. My account is far from airtight or conclusive, but I think that there are at least some strong reasons to entertain them and consider them more rigorously, if only to demonstrate why there is a genuine difference between these phenomena.

---

<sup>367</sup> See Lackey, J. 2008; 233. In the example, a speaker and hearer converse in what they believe to be complete privacy but unknowingly have an additional hearer eavesdropping on them. Both the known hearer and the unknown eavesdropper come to believe the speaker's testimony "which is in fact both true and epistemically impeccable".

My intuition is that unknown hearers still acknowledge the purported testimony in agreement and disagreement, but it presents an interesting case in which the speaker lacks awareness or knowledge of their agreement and disagreement. Nonetheless, I have granted the possibility of propositional, faraway, or merely possible agreement and disagreement, so the prospect surely fits within both of my accounts.

<sup>368</sup> For discussions of group testimony, see Lackey's 2014 and 2020, and Deborah Tollefsen's 2007.

## Bibliography

### Introduction and Part 1

Aristotle and Kirwan, C. (transl.) 1993 *Metaphysics: Books Gamma, Delta, and Epsilon 2<sup>nd</sup> Edition* Oxford: Oxford University Press.

Bailin, S. and Battersby, M. 2016 *DAMned if you do DAMned if you don't: Cohen's "Missed Opportunities"*. OSSA Conference Archive; accessed November 2023  
<https://scholar.uwindsor.ca/ossaarchive/OSSA11/papersandcommentaries/90/>

Berlin, I. [1988] "The Pursuit of the Ideal" reprinted in Hardy, H. and Hausheer, R. (eds.) *The Proper Study of Mankind: An Anthology of Essays*; London: Random House Publishing, 1-17.

Bex-Priestley, G. and Shemmer, Y. 2017 "A Normative Theory of Disagreement" *Journal of the American Philosophical Association* 3 (2) 189-208.

Beziau, J.Y. 2016 "Disentangling Contradiction from Contrariety via Incompatibility" *Logica Universalis* 10 (2-3); 157-170.

Blackburn, S. 1998 *Ruling Passions*; Oxford: Oxford University Press.

Bowker, J. 2000 *The Concise Oxford Dictionary of World Religions*; Oxford: Oxford University Press.

Brennan, A. 2017 "Necessary and Sufficient Conditions" in Zalta, E.N. (ed.) *The Stanford Encyclopaedia of Philosophy*; accessed November 2023 <https://plato.stanford.edu/entries/necessary-sufficient/>

Cappelan, H. and Hawthorne, J. 2009 *Relativism and Monadic Truth*; Oxford: Oxford University Press

Cassam, Q. 2018 "Epistemic Insouciance" *Journal of Philosophical Research* 43; 1-20.

Chalmers, D. 2020 "What is Conceptual Engineering and What Should it Be?" *Inquiry*.

Elga, A. 2007; "Reflection and Disagreement" *Noûs* 41 (3); 478-502.

Enoch, D. 2009 "How is Moral Disagreement a Problem for Realism?" *The Journal of Ethics* 13 (1); 15-50.

Ferrari, F. 2022 "Disagreement and Suspended Judgment" *Metaphilosophy* 53 (4); 526-542.

Frances, B. 2014 *Disagreement*; Cambridge: Polity Press.

Frances, B. and Matheson, J. 2019 "Disagreement" in Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy*; accessed November 2023 <https://plato.stanford.edu/entries/disagreement/>

Frankfurt, H. 2005 *On Bullshit*; Princeton, NJ: Princeton University Press

Friedman, J. 2013; "Rational Agnosticism and Degrees of Belief" in Gendler, T. S. and Hawthorne, J. (eds) *Oxford Studies in Epistemology Volume 4*; Oxford: Oxford University Press; 57-81.

Gibbard, A.  
1990 *Wise Choices, Apt Feelings*; Cambridge, MA: Harvard University Press.  
2003 *Thinking How to Live*; Cambridge, MA: Harvard University Press.

Goldman, A. 2010 "Epistemic Relativism and Reasonable Disagreement" in Feldman, R. and Warfield, T. (eds.) *Disagreement*; Oxford: Oxford University Press; 187-215.

Gregory, A. 2012 "Changing Direction on Direction of Fit" *Ethical Theory and Moral Practice* 15 (5); 603-614.

Grice, H.P. 1975 "Logic and Conversation" reprinted in Grice, H.P. 1989 *Studies in the Way of Words* Cambridge, MA: Harvard University Press.

Hanks, P.W. 2007 "The Content-Force Distinction" in *Philosophical Studies* 134; 141-164.



- Haslanger, S. 2020 “How Not to Change the Subject” in Marques, T. and Wikforss, A. (eds.) *Shifting Concepts: The Philosophy and Psychology of Conceptual Variability*; Oxford: Oxford University Press; 235-259.
- Huvenes, T.T. 2015 “Epistemic Modals and Credal Disagreement” *Philosophical Studies* 172; 987-1011.
- Jackson, E. 2020 “The Relationship Between Belief and Credence” *Philosophy Compass* 15 (6)
- Jackson, F. 1985 “Internal Conflicts in Desires and Morals” *American Philosophical Quarterly* 22 (2); 105-114.
- Kirchin, S. 2017 *Thick Evaluation*; Oxford: Oxford University Press.
- Kölbel, M. 2014 “Agreement and Communication” *Erkenntnis* 79 (1) 101-120.
- Kuhn, S. 2019 “Prisoner’s Dilemma” in Zalta, E.N. (ed.) *The Stanford Encyclopaedia of Philosophy*; accessed November 2023 <https://plato.stanford.edu/entries/prisoner-dilemma/>
- MacFarlane, J.  
2007 “Relativism and Disagreement” *Philosophical Studies* 132; 17-31.  
2014 *Assessment Sensitivity: Relative Truth and its Applications*; Oxford: Oxford University Press.
- Moore, G.E. 1903 *Principia Ethica*; Cambridge: Cambridge University Press.
- Morreal, J. 2020 “Philosophy of Humor” in Zalta, E.N. (ed.) *The Stanford Encyclopaedia of Philosophy*; accessed November 2023 <https://plato.stanford.edu/entries/humor/>
- Palmira, M. 2018 “Disagreement, Credences, and Outright Belief” *Ratio* 31 (2); 179-196.
- Ridge, M. 2013 “Disagreement” *Philosophy and Phenomenological Research* 86 (1); 41-63.
- Schopenhauer, A. and Payne, E.F.J. (transl.) 1966 [1818] *The World as Will and Representation, Volume 1*, Mineola, NY: Dover Publications, Inc.
- Schroeder, M. 2010 *Noncognitivism in Ethics*; London: Routledge.
- Sennet, A. 2023 “Ambiguity” in Zalta, E.N. (ed.) *The Stanford Encyclopaedia of Philosophy*; accessed April 2024. <https://plato.stanford.edu/entries/ambiguity/>
- Shemmer, Y. and Bex-Priestley, G. 2021 “Disagreement Without Belief” *Metaphilosophy* 52 (3-4); 494-507.
- Stevens, K. and Cohen, D. 2021 “Angelic Devil’s Advocates and the Forms of Adversariality” *Topoi* 40 (5); 899-912.
- Stevenson, C.L.  
1944 *Ethics and Language*; New Haven, CT: Yale University Press  
1963 *Facts and Values: Studies in Ethical Analysis*; New Haven, CT: Yale University Press
- Stroud, S. 2019 “Conceptual Disagreement” *American Philosophical Quarterly* 56 (1); 15-27.
- Studtmann, P. 2021 “Aristotle’s Categories” in Zalta, E.N. (ed.) *The Stanford Encyclopaedia of Philosophy*; accessed November 2023 <https://plato.stanford.edu/entries/aristotle-categories/>
- Sturgeon, S. 2020 “The Bayesian Theory of States: Critical Discussion” in *The Rational Mind*, Oxford: Oxford University Press; 61-100.
- Szabó, Z.G. 2005; *Semantics versus Pragmatics*; Oxford: Oxford University Press
- Whiting, D. 2020; “Recent Work: Higher Order Evidence” *Analysis Reviews* 80 (4); 789-807.
- Williams, J.N. 1981 “Inconsistency and Contradiction” *Mind* 90; 600-602.
- Wittgenstein, L. 2009 [1953], Anscombe, G.E.M., Hacker, P.M., and Schultz, E. (trans.) *Philosophical Investigations, 4<sup>th</sup> Edition*; Hoboken, NJ: Wiley-Blackwell.
- Yalcin, S. 2011 “Nonfactualism about Epistemic Modality” in Egan, A. and Weatherson, B. (eds) *Epistemic Modality*; Oxford: Oxford University Press.

## Part 2 and Afterthought

- Asher, N. and Lascarides, A. 2001 "Indirect Speech Acts" *Synthese* 128 (1-2); 183-228.
- Austin, J.L. 1975, Urmson, J.O., and Sbisà, M. (eds.) *How to Do Things with Words* 2nd Edition. Cambridge, MA: Harvard University Press.
- Ballantyne, N. 2019 "Epistemic Trespassing" *Mind* 128 (510); 367-395.
- Ballantyne, N., and Coffman, E.J. 2012 "Conciliationism and Uniqueness" *Australasian Journal of Philosophy* 90 (4); 657-670.
- Battaly, H. 2021 "Intellectual Autonomy and Intellectual Interdependence" in Matheson, J. and Loughheed, K. (eds) 2022 *Epistemic Autonomy*, New York, NY: Routledge.
- Berlin, I. [1978] "The Decline of Utopian Ideas in the West" reprinted in Hardy, H. (ed.), Berlin, I. (2013) *The Crooked Timber of Humanity: Chapters in the History of Ideas*; London: Pimlico; 21-50.
- Bird, A. "Illocutionary Silencing" *Pacific Philosophical Quarterly* 83; 1-5.
- Bradley, D. 2019 "Are There Indefeasible Epistemic Rules?" *Philosophers' Imprint* 19.
- Carel, H. and Kidd, I. 2014 "Epistemic Injustice in Healthcare: A Philosophical Analysis" *Medicine, Health Care and Philosophy* 17 (4); 529-540.
- Carey, B. 2011 "Possible Disagreements and Defeat" *Philosophical Studies* 155 (3); 371-381.
- Carr, J.R. 2021 "Why Ideal Epistemology?" *Mind* 131 (524); 1131-1162.
- Carter, J.A. 2014 "Group Peer Disagreement" *Ratio* 27 (3); 11-28.
- Christensen, D.  
2007 "Epistemology of Disagreement: The Good News" *The Philosophical Review* 116 (2); 187-217.  
2009 "Disagreement as Evidence: The Epistemology of Controversy" *Philosophy Compass* 4/5; 756-767.  
2010 "Rational Reflection" *Philosophical Perspectives* 24; 121-140.
- Coady, C.A.J. 1992 *Testimony: A Philosophical Study*; Oxford: Oxford University Press.
- Conee, E. 2009 "Peerage" *Episteme* 6 (3); 313-323.
- Croce, M. 2018 "Expert-Oriented Abilities vs. Novice-Oriented Abilities: An Alternative Account of Epistemic Authority" *Episteme* 15 (4); 476-498.
- Dormandy, K. (forthcoming) "Epistemic Self-Trust: It's Personal" *Episteme* 1-16. Accessed November 2023
- Dorst, K. 2019 "Higher-Order Uncertainty" in Skipper, M. and Steglich-Petersen, A. (eds.) *Higher-Order Evidence: New Essays*; Oxford: Oxford University Press, 35-61.
- Dutant, J. 2016 "The Legend of the Justified True Belief Analysis" *Philosophical Perspectives* 29 (1); 95-145.
- Easwaran, K., Fenton-Glynn, L., Hitchcock, C., and Velasco, J.D. 2016 "Updating on the Credences of Others: Disagreement, Agreement, and Synergy" *Philosophers' Imprint* 16 (11) 1-39.
- Elga, A. 2007; "Reflection and Disagreement" *Noûs* 41 (3); 478-502.
- Enoch, D. 2010 "Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement" *Mind* 119 (476); 953-997.
- Fantl, J. 2018 *The Limitations of the Open Mind*, Oxford: Oxford University Press.
- Feldman, R.  
2006 "Reasonable Religious Disagreements" in Anthony, L. (ed.) *Philosophers Without Gods: Meditations on*

- Atheism and the Secular Life*; Oxford: Oxford University Press, 194-214.
- 2009 “Evidentialism, Higher-Order Evidence, and Disagreement” *Episteme* 6 (3); 294-312.
- Feldman, R. and Warfield, T. (eds) 2010 *Disagreement*; Oxford: Oxford University Press.
- Frances, B.
- 2013 “Philosophical Renegades” in Lackey, J. and Christensen, D. (eds) *The Epistemology of Disagreement: New Essays*; Oxford: Oxford University Press, 121-166.
- 2014 *Disagreement*; Malden, MA: Polity.
- Frances, B. and Matheson, J. 2019 “Disagreement” in in Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy*; accessed November 2023 <https://plato.stanford.edu/entries/disagreement/>
- Fricker, M.
- 2007 *Epistemic Injustice: Power and the Ethics of Knowing*; Oxford: Oxford University Press.
- 2012 Group Testimony? The Making of a Collective Good Informant *Philosophy and Phenomenological Research* 84 (2); 249-276.
- Friedman, J. 2023 “The Aim of Inquiry?” *Philosophy and Phenomenological Research*; 1-18.
- Gelfert, A. 2011 “Who is an Epistemic Peer?” *Logos and Episteme* 2 (4); 507-514.
- Goldman, A. 2001 “Experts: Which Ones Should You Trust?” *Philosophy and Phenomenological Research* 63 (1); 85-110.
- Greco, D. and Hedden, B. 2016 “Uniqueness and Metaepistemology” *Journal of Philosophy* 113 (8); 365-395.
- Grundmann, T. 2019 “How to Respond Rationally to Peer Disagreement: The Preemption View” *Philosophical Issues* 29 (1); 129-142.
- Hardwig, J. 1985 “Epistemic Dependence” *Journal of Philosophy* 82 (7); 335-349.
- Hills, A. 2009 “Moral Testimony and Moral Epistemology” *Ethics* 120 (1); 94-127.
- Hopkins, R. 2007 “What is Wrong With Moral Testimony?” *Philosophy and Phenomenological Research* 74 (3); 611-634.
- Horowitz, S. 2014 “Epistemic Akrasia” *Noûs* 48 (4); 718-744.
- Jäger, C. 2016 “Epistemic Authority, Preemptive Reasons, and Understanding” *Episteme* 13 (2); 167-185.
- Jenkins, C.S.I. 2014 “Merely Verbal Disputes” *Erkenntnis* 79 (1); 11-30.
- John, S. 2011 “Expert Testimony and Epistemological Free-Riding: The MMR Controversy” *Philosophical Quarterly* 61 (244); 496-517
- Johnson, C.R.
- 2018a “Just Say ‘No’: Obligations to Voice Disagreement” *Royal Institute of Philosophy Supplement* 84; 117-138.
- 2018b “For the Sake of Argument: The Nature and Extent of Our Obligation to Voice Disagreement” in *Voicing Dissent: The Ethics and Epistemology of Making Disagreement Public*, New York: NY, Routledge; 97-108.
- 2020 “Epistemic Vulnerability” *International Journal of Philosophical Studies* 28 (5); 667-691.
- Kappel, K. 2010 “On Saying That Someone Knows: Themes from Craig” in Haddock, A., Millar, A., and Pritchard, D. (eds) *Social Epistemology*; Oxford: Oxford University Press, 69-88.
- Kelly, T.
- 2005 “The Epistemic Significance of Disagreement” in Hawthorne, J. and Gendler, T. (eds) *Oxford Studies in Epistemology*; Oxford: Oxford University Press, 167-196.
- 2010 “Peer Disagreement and Higher-Order Evidence” in Feldman, R. and Warfield, T. (eds.) *Disagreement*; Oxford: Oxford University Press.183-217.

- Kelp, C. 2021 *Inquiry, Knowledge, & Understanding*, Oxford: Oxford University Press.
- Kenyon, T. 2020 “Peer Idealization and Internal Examples in the Epistemology of Disagreement” *Dialogue* 59 (1); 69-79.
- King, N.L. 2012 “Disagreement: What’s the Problem? or A Good Peer is Hard to Find” *Philosophy and Phenomenological Research* 85 (2); 249-272.
- Kinzel, K., and Kusch, M. 2018 “De-idealizing Disagreement, Rethinking Relativism” *International Journal of Philosophical Studies* 26 (1); 40-71.
- Kirchin, S. and Pethick, S. (unpublished manuscript) “Conceptual Disagreement and the Epistemology of Disagreement Accessed November 2023 [https://b71307fb-f69e-4cfc-809a-f1aa56c4307a.filesusr.com/ugd/53f68c\\_dcefdabfabcb401dbe07966fa6a9161d.pdf](https://b71307fb-f69e-4cfc-809a-f1aa56c4307a.filesusr.com/ugd/53f68c_dcefdabfabcb401dbe07966fa6a9161d.pdf)
- Kopec, M. 2015 “A Counterexample to the Uniqueness Thesis” *Philosophia* 43 (2) 403-409.
- Kornblith, H. 2010 “Belief in the Face of Controversy” in Feldman, R. and Warfield, T. (eds.) *Disagreement*, Oxford: Oxford University Press; 29-52.
- Kvanvig, J. 2003 *The Value of Knowledge and The Pursuit of Understanding*, Cambridge: Cambridge University Press.
- Lackey, J.  
2008 *Learning From Words: Testimony as a Source of Knowledge*, Oxford: Oxford University Press.  
2010 “A Justificationist View of Disagreement’s Epistemic Significance” in Haddock, A., Millar, A., and Pritchard, D. (eds) *Social Epistemology*; 145-154.  
2014 *Essays in Collective Epistemology*, Oxford: Oxford University Press  
2018a “Experts and Peer Disagreement” in Benton, M.A., Hawthorne, J., and Rabinowitz, D. *Knowledge, Belief, and God: New Insights in Religious Epistemology*, Oxford: Oxford University Press, 228-245.  
2018b “The Duty to Object” *Philosophy and Phenomenological Research* 101 (1); 35-60.  
2020 *The Epistemology of Groups*, Oxford: Oxford University Press.
- Langton, R. 1993 “Speech Acts and Unspeakable Acts” *Philosophy and Public Affairs* 22 (4); 293-330.
- Lasonen-Aarnio, M. 2014 “Higher-Order Evidence and the Limits of Defeat”
- Levy, N.  
2019 “No-Platforming and Higher-Order Evidence, or Anti-Anti-No-Platforming” *Journal of the American Philosophical Association* 5 (4); 487-502.  
2022 “Do Your Own Research!” *Synthese* 200: 356; 1-19.  
2024 “Against Intellectual Autonomy: Social Animals Need Social Virtues” *Social Epistemology* 38 (3); 350-363.
- Lord, E. and Sylvan, K. 2021 “Suspension, Higher-Order Evidence, and Defeat” in Simion, M. and Brown, J. (eds) *Reasons, Justification, and Defeat*, Oxford: Oxford University Press, 116- 145.
- Martin, N.G. 2019 “What is the Epistemic Significance of Disagreement?” *Logos and Episteme* 10 (3); 283-298.
- Matheson, J.  
2009 “Conciliatory Views of Disagreement and Higher-Order Evidence” *Episteme* 6 (3); 269-279.  
2014 “Disagreement: Idealized and Everyday” in Matheson, J. and Vitz, R. (eds.) *The Ethics of Belief: Individual and Social*, Oxford: Oxford University Press, 315-330.  
2015 *The Epistemic Significance of Disagreement*, Basingstoke: Palgrave Macmillan.  
2022 “Why Think for Yourself?” *Episteme*, 1-19.
- McGrath, S. 2009 “The Puzzle of Pure Moral Deference” *Philosophical Perspectives* 23 (1); 321-344.
- Mill, J.S. [1859] “On Liberty” in Philp, M., and Rosen, F. (eds.) 2015 *On Liberty, Utilitarianism, and Other Essays*, Oxford: Oxford University Press, 5-112.

- Mills, C. 2017 *Black Rights/White Wrongs: The Critique of Racial Liberalism*; New York, NY: Oxford University Press.
- Oppy, G. 2010 “Disagreement” *International Journal for Philosophy of Religion* 68 (1-3); 183-199.
- Pollock, J.  
 2020 “Context and Communicative Success” in Ciecierski, T. and Grabarczyk, P. (eds) *The Architecture of Context and Context-Sensitivity*; Springer.  
 2021 “Linguistic Understanding and Testimonial Warrant” *Erkenntnis*.
- Potochnik, A. 2020 “Idealization and Many Aims” *Philosophy of Science* 87 (5) 933-943.
- Priest, M. 2016 “Inferior Disagreement” *Acta Analytica* 31 (3); 263-283.
- Raz, J. 1988 *The Morality of Freedom* Oxford: Clarendon Press.
- Rosen, G. 2001 “Nominalism, Naturalism, Epistemic Relativism” *Philosophical Perspectives* 15; 69-91.
- Schoenfield, M. 2014 “Permission to Believe: Why Permissivism is True and What it Tells Us about Irrelevant Beliefs” *Noûs* 48 (2); 193-218.
- Scully, J.L. 2018 “From ‘She Would Say That, Wouldn’t She?’ to ‘Does She Take Sugar?’ Epistemic Injustice and Disability” *International Journal of Feminist Approaches to Bioethics* 11 (1); 106-124.
- Simpson, R.M. 2012 “Epistemic Peerhood and the Epistemology of Disagreement” *Philosophical Studies* 164 (2); 561-577.
- Smith, N.D. 2015 “Peer Disagreement, Testimony, and Personal Justification” *Athens Journal of Humanities & Arts* 2 (3); 141-150.
- Steel, R. “Against Right Reason” *Philosophy and Phenomenological Research* 99 (2); 431-460.
- Sundell, T. 2011 “Disagreements about Taste” *Philosophical Studies* 155 (2); 267-288.
- Titelbaum, M. 2015 “Rationality’s Fixed Point (or: In Defence of Right Reason) in Gendler, T.S., and Hawthorne, J. (eds.) *Oxford Studies in Epistemology, Volume 5*; Oxford: Oxford University Press, 253-294.
- Tollefsen, D. 2007 “Group Testimony” *Social Epistemology* 21 (3); 299-311.
- Weatherston, B. (unpublished manuscript) *Do Judgments Screen Evidence?*  
 Accessed November 2023 - <https://brian.weatherston.org/JSE.pdf>
- Whiting, D. 2020; “Recent Work: Higher Order Evidence” *Analysis Reviews* 80 (4); 789-807.
- White, R.  
 2005 “Epistemic Permissiveness” *Philosophical Perspectives* 19 (1); 445-459.  
 2014 “Evidence Cannot Be Permissive” in Steup, M. and Turri, J. (eds.) *Contemporary Debates in Epistemology, 2nd Edition*; Hoboken, NJ: Wiley-Blackwell, 312-323.
- Worsnip, A. (forthcoming) “Deference to Experts” in Dancy, J., Sosa, E., Steup, M., and Sylvan, K. (eds) *A Companion to Epistemology*; Wiley-Blackwell. Accessed November 2023 - <https://www.alexworsnip.com/wp-content/uploads/2023/01/Deference-to-Experts.pdf>
- Zagzebski, L.T. 2012 *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*; Oxford: Oxford University Press.