

Kent Academic Repository

Full text document (pdf)

Citation for published version

Dragulinescu, Stefan (2016) Mechanisms and Difference-Making. Acta Analytica . pp. 1-26. ISSN 0353-5150.

DOI

<https://doi.org/10.1007/s12136-016-0292-1>

Link to record in KAR

<http://kar.kent.ac.uk/55988/>

Document Version

Publisher pdf

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Mechanisms and Difference-Making

Stefan Dragulinescu¹

Received: 13 November 2015 / Accepted: 15 March 2016

© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract I argue that difference-making should be a crucial element for evaluating the quality of evidence for mechanisms, especially with respect to the robustness of mechanisms, and that it should take central stage when it comes to the general role played by mechanisms in establishing causal claims in medicine. The difference-making of mechanisms should provide additional compelling reasons to accept the gist of Russo-Williamson thesis and include mechanisms in the protocols for Evidence-Based Medicine (EBM), as the EBM+ research group has been advocating.

Keywords Mechanism · Difference making · Causation · Medical causation

1 Introduction

It is commonplace that in epistemology and the philosophy of science, the nature of evidence in medicine has become in recent years one of the most researched topics. An important aspect of it is the evidence of mechanisms, and a group of philosophers of science have been urging for some time that mechanisms be included in the evaluation of causal medical claims at a higher level than the one currently afforded by the EBM (Evidence-Based Medicine) protocols (Clarke et al. 2014a), following the general lines of the Russo-Williamson thesis, according to which both mechanisms and population level assessments are necessary in order to draw medical causal inferences (Russo and Williamson 2007, 2011). A crucial part of this project, otherwise known as EBM+,¹ is dealing with the epistemology of mechanisms themselves and with the quality of

¹Where the “+” means “mechanisms + trials”, but also refers to the power of intersecting species of evidence. The analogy guiding the EBM+ group is that of steel-reinforced concrete, where the two materials—one good under compression, the other good under tension—mutually support each other. For an introduction, see <https://blogs.kent.ac.uk/ebmplus/>

✉ Stefan Dragulinescu
s.dragulinescu@gmail.com; sd508@kent.ac.uk

¹ Philosophy Department, School of European Culture and Languages, University of Kent, Canterbury CT2 7NF, UK

evidence for mechanisms used to put forward causal claims at this level. Indeed, Clarke et al. have provided in their (2014a) paper a preliminary protocol that takes into account important elements such as the robustness of mechanisms, the proportion of elements found, the research groups, and the methodologies involved in mechanistic research. The purpose of this paper is to advance some way towards understanding what the *robustness* of mechanisms amounts to, as well as towards making sense of how mechanistic claims are (or indeed, should) to be taken in conjunction with the population level assessments.

Thus, in §1, after presenting briefly the programme and purpose of the EBM+, I will propose that difference-making should be a criterion for judging the robustness of mechanisms, and more generally, for justifying the importance of mechanisms in evaluating evidence in medicine. I will argue in §2 that the root of the mistaken separation between mechanisms and difference making that is commonly embraced in the literature is an implausible, and often tacit, ontic pluralistic view about (medical) causation. As I will show in §3, it is precisely because of missing this feature of mechanisms that EBM theorists were able to argue in favor of a scheme of evidence evaluation that downgrades the importance of mechanisms or ignores it altogether, thus disregarding the path of research advocated by EBM+.

Indeed, among others, this paper aims to provide a response to Jeremy Howick's (2011) criticism of the Russo-Williamson thesis. As one of the leading thinkers of the EBM movement, Howick has provided us with trenchant and illuminating expositions as to the purpose and functioning of evidence-based medicine, and the details and perspective of the Russo-Williamson thesis could of course be criticized, as I did myself from a certain perspective.² At the same time, I agree with the basic insight of the thesis, which is that mechanisms are crucial for medicine, and I will try to explain why Howick could not come to terms with it, as well as why he, as well as the other EBM theorists, should accept the gist of it.

2 The Protocol for the Quality of Evidence of Mechanisms

The EBM+ project stems from work done initially by Fredrica Russo and Jon Williamson into the epistemology of causal claims in medicine. What came to be known as the “Russo-Williamson thesis” (henceforth RWT), states that “the health sciences make causal claims on the basis of evidence both of physical mechanisms, and of probabilistic dependencies. Consequently, an analysis of causality solely in terms of physical mechanisms or solely in terms of probabilistic relationships, does not do justice to the causal claims of these sciences” (Russo and Williamson 2007, p. 157). That is to say, RWT states we need

² Namely from the point of view of the Humean/anti-Humean distinction (see Dragulinescu 2012), and I still stand by the argumentation I have there advanced regarding Williamson's epistemic account of causation. However, when it comes to the problem of extrapolation, I have failed to see that, even if mechanisms as such also have problems of extrapolation, taking mechanisms in conjunction with population studies should be an advance towards mitigating the extrapolation worries. Hence, the critique I address to Howick in section 3 of the present paper is also an acknowledgment of my own misguided view on the gist of the Russo-Williamson thesis. I am focusing on Howick's paper because his exposition of the problem of extrapolation for mechanisms is much more clear and informed with scientific acumen than my own exposition in Dragulinescu (2012).

evidence of both *production* (which comes primarily from a study of mechanisms on a microstructural level and entails knowledge of entities, their spatio-temporal arrangement, etc.) and *dependency* or *difference making* (which comes primarily from a study of the probabilistic relationships on the level of population assessments), in order to formulate adequately medical causal claims. Initially connected with Williamson’s Humean epistemic account of causation, RWT has gained larger adherence within a group of philosophers with variegated perspectives, who have put forward a fully articulated argumentation for placing evidence of mechanisms higher than it currently is in the EBM protocols. Clarke et al. (2014a) argue that evidence of mechanisms should complement evidence of correlation “because taken on its own each sort of evidence has its own limitations, because they need to be taken together to establish causal claims, to transport causal claims to new populations and to apply causal claims to individual patients” (p. 341).

Clarke et al. also list the following preliminary features for grading evidence of mechanisms: the independent methods that confirm (or disconfirm) a feature, the independent research groups that confirm (or disconfirm) a feature, the proportion of features found (larger or smaller), knowing analogous mechanisms as opposed to not knowing analogous mechanisms or even worse, knowing that analogous situations do not exhibit such mechanisms, robustness, i.e., being reproducible across a wide range of conditions, as opposed to fragility of mechanisms, i.e., not being reproducible even in slightly varying conditions.³

<i>Pluses</i>	<i>Minuses</i>
Each independent method that confirms a feature	Each independent method that fails to confirm—or, worse, disconfirms—a feature
Each independent research group that confirms a feature	Each independent research group that fails to confirm—or, worse, disconfirms—a feature
Larger proportion of features found	Smaller proportion of features found
Analogous mechanisms known	The analogy is a weak one, or, worse, analogous situations exhibit no such mechanism
Robust, reproducible across a wide range of conditions	Fragile, not reproducible in slightly varying conditions

The list is extensive and covers a large part of the epistemological means that should be employed for assessing the quality of evidence of mechanisms. But as the authors themselves urge, more work needs to be done. In particular, one would want to know more about the most important of the criteria listed above, namely robustness. It is certainly useful to know that a robust mechanism is reproducible in a wide variety of conditions, but one wonders—what is it that it is reproduced? Some ready-made answers come easily to mind; for instance, one could say that it is the *functioning* of mechanisms that is being reproduced. But the ready-made answers call to mind other questions. What does the functioning of a mechanism consist in? We come thus to an important aspect of evaluating evidence of mechanisms, namely that one needs to know

³ Clarke et al. (2014a p. 357).

or to establish what a mechanism *is*, at least to a certain extent, in order to put more flesh to the bones of such preliminary criteria.⁴ However, a strange feature of mechanistic accounts nowadays is that most of them, while mentioning production, responsibility for events, etc. avoid the terminology of difference-making (in its counterfactual guise, as probabilistic dependency, or whatnot).⁵ Here are some well-known examples:

Bechtel and Abrahamsen (2005) ‘A mechanism is a structure performing a *function* in virtue of its component parts, component operations, and their organization. The orchestrated *functioning* of the mechanism is *responsible* for one or more phenomena.’

Glennan (2002) ‘A mechanism for a behavior is a complex system that *produces* that behavior by the interaction of a number of parts, where the interactions between parts can be characterized by direct, invariant, change-relating generalizations.’

Machamer et al. (2000) ‘Mechanisms are entities and activities organized such that they are *productive* of regular changes from start or set-up to finish or termination conditions.’

Illari and Williamson (2012) ‘A mechanism for a phenomenon consists of entities and activities organized in such a way that they are *responsible* for the phenomenon.’⁶

Furthermore, evidence of mechanisms is often set in contrast to difference-making evidence, where the latter should come from various probabilistic or counterfactual accounts, mostly applied to population level assessments, the main cited reason being the problem of pre-emption—cases of mechanistic production in which the

⁴ In order to reach the largest audience, Clarke et al. do use the broad, non-committal definition of mechanism in provided in Illari and Williamson 2012 ‘a mechanism for a phenomenon consists of entities and activities organized in such a way that they are responsible for the phenomenon’ (Illari and Williamson 2012, p. 120, *apud* Clarke et al. 2014a p.343). It is a very good definition, which I would seek however to particularize, taking of course the risk of going astray and entering into an area of controversy; it is I believe, unavoidable, in order to make the evidence claims sharper.

⁵ For an overview of complex system mechanistic accounts, see Psillos 2004; Illari and Williamson 2012.

⁶ Truth be told, Glennan’s account implies a counterfactual account of the components of mechanisms, albeit not of the final output, whereas Illari and Williamson do mention the counterfactual element as part of their notion of ‘responsibility’, without developing further though. (Illari and Williamson 2012, p. 125). Casini, *et al.* (2011) and Clarke et al. (2014b) have put forward a scheme for modelling mechanisms using recursive or dynamic Bayes nets, which is quite consistent with the difference making view of mechanisms, given that causal modelling looks of course for probabilistic dependencies between variables. The authors do not draw the consequences as to the nature of mechanisms though; Clarke *et al.* (2014b) note that Woodward’s counterfactual approach and the epistemic account is consistent with Bayes nets, but choose to remain agnostic on this side (pp. 1653, 1754). Casini has very recently criticized the interventionist counterfactual aspect of Glennan’s account of mechanisms, proposing instead the dynamic Bayes nets approach (Cassini 2015). Casini notes that other formalisms seems suitable to represent mechanisms depending on various biological contexts (differential equations, neural networks, Petri nets, agent based models) and finds difficult to arrive at an analogy or common ground between them all, suggesting that “some form or other of pluralism” would do, as well as the integration of such pluralism within the epistemic account of causation or his own referential view (Cassini 2015, p. 25–26)

counterfactual approach cannot isolate the causes at stake because other causes would have produced the effects anyway. Now, this pluralistic view of evidence, which is in fact embraced by RWT and is quite innocuous in itself,⁷ has gone hand in hand with a pluralistic approach of causation, which says, in its *conceptual* guise, that we possess different concepts of causation, while in its *ontic* formulation, the more drastic one, says that the multiplicity of concepts of causation reflects the multiplicity of causal relations as such.⁸ The main proponent of this pluralistic approach has been Ned Hall, who has argued in his “*Two Concepts of Causation*” (2004) that we have a concept of causation as dependence (expressed counterfactually), and one as production (which could be expressed for instance in the mechanistic accounts of Salmon and Dowe, and in the complex systems accounts of mechanisms) and that these point to different causal relations, as such.⁹

The pluralistic view of *evidence*, as well as *conceptual* pluralism about causation, are quite acceptable views. That is to say, one can accept that we can have separate evidence of difference-making and separate evidence of production, as well as that our concepts of dependency and production are different from one another. On the other hand, the *ontic* pluralism is a strange and implausible position.¹⁰ Indeed, the reason why this ontic pluralism about causation—which has as a consequence the ontic separation between mechanisms and difference-making—has gained and continued to enjoy so much popularity, is slightly mysterious. Beyond Hall’s work, the (medical) philosophical literature also had at disposal for instance Bird’s 2007 monistic interpretation of capacity-based causation, Psillos’ 2004 argument that counterfactuals are essential for mechanisms and Strevens’ 2013 unificatory approach for Hall’s two concepts of causation. However, no one seems to have drawn the consequences for medical mechanisms and the heated problem of the relation between mechanisms and population assessments, in the sense of adding difference-making as a characteristic of

⁷ Russo and Williamson (2007, pp. 158–162). Clarke et al. 2014a, also set in contrast evidence of mechanism and evidence of ‘correlation’. (Clarke et al. 2014a, p. 343) According to Illari’s disambiguation of the Russo-Williamson thesis (2011), these are not to be taken as two *kinds* of evidence, in the sense in which, for instance, a good RCT should give us also evidence of the existence of mechanism, even if we do not know its details. But the distinction between difference making and mechanisms as *objects* of evidence is maintained by Illari as well (Illari 2011, pp. 141, 146). Illari’s disambiguation is welcome and illuminating, but I would want to go a step further—difference making is a necessary condition for mechanisms. The difference between mechanistic microstructural, and probabilistic evidence from the level of populations is epistemic, in the sense in which they look at the same causal relations, but from ‘inside the box’ and from ‘outside the box’, as it were. More on this will be explored above.

⁸ See de Vreese 2006; Godfrey-Smith 2008, for instance.

⁹ Hall 2004, p. 181.

¹⁰ A strong and convincing rejection of causal pluralism can be found in Russo and Williamson 2007, pp. 165–167, Williamson 2011, pp. 435–437, where it is used as a premise for the epistemic account of causation. I embrace Russo and Williamson’s rejection of ontic causal pluralism and, as a matter of fact, I am following closely their strategy against this type of pluralism, up to the point at which I suggest that difference-making should be construed as part of mechanisms. The question of whether such difference-making is to be cashed out in Humean terms, as it is the case in the epistemic account, or in any other non-Humean terms (i.e., the question of what is it that underlies as ultimate truth-maker the dependencies we take into account when drawing causal claims), lies beyond the scope of this paper, and should be the subject of a strictly metaphysical discussion. But Williamson’s Humean solution is I presume consistent with the difference-making view I am advocating here; see, of course, Hume’s *Enquiry*: ‘[W]e may define a cause to be an object followed by another, and where all the objects, similar to the first are followed by objects similar to the second. Or, in other words, where, if the first object had not been, the second never had existed.’ [1748, Section VII, Part II] 1995, p. 87

mechanisms. I do not mean to say, of course, that the philosophical dispute over causal pluralism and monism (be it evidential, conceptual or ontic), has been completely solved and thrown to the ash heap of (philosophical) history. I mean to say though that the consequences of the ontically monistic account of causation, which is far more natural, intuitively and scientifically,¹¹ have not been given more attention in the philosophical medical literature.¹²

Now, viewing mechanisms as sharing difference making should have drastic, and beneficial effects as to our understanding of causal claims in medicine. First of all, to put it simply—coming back to common sense. The stark separation on an ontic level between the difference-making of population assessments and the mechanistic relations from the microstructural level holds some untenable suppositions. Talk of populations might well be warranted in certain epistemological and evidential contexts, but ontically, it is not populations which produce (each other, to individuals?) gastric cancer or Hodgkin's lymphoma. It is certain cells, within biological organisms, which transform in such and such ways, as part as such and such mechanisms, bringing forth such and such difference-making—a difference-making which is grasped (as one epistemological means) by the population level assessments.

One can usefully think here of the heuristic image of the “black box” and what is outside of it. Causal relations within mechanisms are those inside the “black box”, while the difference-making from population level assessments is analogous to what is “outside the box.” The interactions within mechanisms, on the one hand, and population level dependencies, on the other hand, do not point to different causal relations, but to different epistemological levels at which the difference-making can be grasped. This does not entail that causation in general is or needs to be mechanistic. It entails though that ontically, in biology, and medicine accordingly, causal relations take place inside mechanisms, where the difference-making (realized by mechanistic entities separately and by the mechanisms overall) provides the essential component, the other essential components being the spatial and temporal arrangement of elements, the start-up conditions, the feed-back regulations, etc.

Again, we could very well keep the pluralistic view of *evidence*, in the sense in which we might well have evidence of production and evidence of difference-making. We might also want to keep the conceptually pluralistic view of causation in the sense in which we can separate the notion of difference-making (as entailing dependencies expressed counterfactually, probabilistically, by way of interventions, etc.) and the notion of production (as entailing spatial and temporal arrangement of various elements, etc.). Nonetheless, ontically, we have the same causal relations whose difference-making originates within mechanisms. Differently put, the notions of *difference-making* and *production* are not synonymous, because difference-making can occur outside mechanisms (and can be grasped outside mechanisms themselves, as it happens for populations assessments), and because mechanisms need more than difference-making in order to be mechanisms. But it is on grounds of what happens inside mechanisms that we can speak about difference-making at all at the level of

¹¹ We should bear in mind that in medicine population level assessments and mechanistic evidence are *aggregated*, and are naturally taken by medical researchers to refer to the same causal claims; a point made convincingly in Russo and Williamson 2007.

¹² Again, a notable exception is Russo and Williamson 2007, pp. 158–162; see also Darby and Williamson 2011, pp. 115–117.

populations. It is also the reason why, tacitly, mechanisms are already pervasively present in all medical practice and in the design, implementation and interpretation of population studies, and I shall come back to this at the end of this section.

Second of all—understanding why evidence of mechanisms is important. Again, to put it simply, one would want to know, especially in a science like medicine, what is inside the “black box”, or at least, in certain cases, that the black box exists. Knowing that a mechanism exists, even without knowing its details (in other words, just knowing that the black box is there, but yet unopened) is important to rule out spurious cases of causation,¹³ and it does rule them out by selecting in contrast the one and the same causal relations present within mechanisms and grasped also at population levels. This will be sufficient for explanation and diagnosis, or prevention.¹⁴ On the other hand, in cases of prediction and extrapolation, we can understand why it is that we need to open the black box. A difference-making grasped by population studies can hide different mechanisms exercising contrary or dissimilar effects, can hide a paradoxical behavior of the same mechanism under various feed-back conditions, can hide multiple pathways within the same mechanism which might change their operation under varying conditions. This brings us to the third point.

Third of all—making sharper and more coherent our claims regarding the evidence of mechanisms. In order to know how robust a mechanism is, one needs to know (i) whether it is difference-making (ii) under how different varying conditions, the difference-making is still in place. Moving the accent on difference-making makes us understand the root of the problems outlined by Clarke et al.—the problems of complexity and masking. The problem of complexity is that a mechanism can use multiple pathways in order to achieve a result, whereas, when the problem of masking shows up, the result of one mechanism is canceled out by the activity of a different mechanisms.¹⁵ But what is it that is canceled (on an overall view) or made more complex? It is, the natural answer would go, the difference-making of mechanism components, and of the mechanisms as such. That such problems show up does not indicate at all that mechanisms do not have difference-making as an essential component. It indicates, on the contrary, that difference-making is in place, in the functioning of the mechanistic components, and that overall, on an epistemic level, such difference-making may be hard to track. This brings us to the fourth point.

Fourth of all—facing straight-on the problem of extrapolation. Mechanisms should not be viewed as a panacea to the problem of extrapolation. They can help towards solving it, they contribute greatly towards solving it, but the problem does not disappear

¹³ Russo and Williamson 2007, p. 159.

¹⁴ Darby and Williamson 2011, p. 116.

¹⁵ “Even where a mechanism linking A to B is well established and known in some detail, it can be hard to infer whether A has a positive effect on B, or A prevents B, or indeed whether A has any net effect on B at all. This is particularly true in cases where the mechanism is complicated: where there are several links on a pathway from A to B or where there are several pathways from A to B. It is also a problem where a mechanism is known to be non-robust over time or over other changes in situation. It is typically evidence of correlation that is crucial for determining whether any causation is positive or negative and what the net effect is. Thus, evidence of mechanisms should be used in conjunction with evidence of correlation, not on its own, to infer causal claims.[...]The human body is a complex system, and the more we discover about it the more it seems that it is very common to have multiple mechanisms operating. If there are multiple mechanisms operating, they may impact on each other, and one or more may mask the effects of the mechanism you have discovered.” Clarke et al. 2014a, b, pp. 349–351

completely. However, it is important to realize that the problem is still there. It is with regard to the incapacity of mechanisms to solve entirely the problem of extrapolation that Jeremy Howick and collaborators have written two convincing papers.¹⁶ This line of reasoning, while essentially correct in its letter, has been put to the wrong use, in the sense in which Howick (2011) has taken the arguments against the capacity of mechanisms to completely solve the problem of extrapolation, as arguments against the Russo-Williamson thesis.

But that is completely unwarranted. RWT does not hold that mechanisms are a panacea for our claims of “external validity”, and does not advocate the use of mechanisms *alone* in establishing causal claims. It advocates the simultaneous (and equal importance for the) use of both mechanisms and population level assessments, in contrast to the current view of EBM in which only the population assessments are prioritized. The problem of extrapolation is not a problem *of* mechanisms per se (a problem that supposedly would not show up for population level assessments), just like it is not a problem of population level assessments per se, such that it would not appear for mechanisms, the latter being some sort of panacea. It is a problem of causal relations in medicine, in the sense in which the difference-making, grasped either at the microstructural or the population level, might be modified depending on varying conditions. One needs both population assessments and mechanisms to solve or alleviate it. And using the issue of extrapolation against mechanisms, and against the Russo-Williamson thesis, will just not do, as I will show in §3, after rehearsing briefly in §2 the monistic ontic response that can be gathered, *contra* Hall, from the abovementioned arguments espoused by Psillos and Strevens.

Since the root of the problem of Howick’s views (and the separation between mechanisms and difference making) is the ontically pluralistic view of causation typified by Ned Hall’s perspective, showing how the arguments of Psillos et al. render Hall’s ontic position implausible is a necessary step before approaching Howick. To repeat a point made earlier, the philosophical controversy concerning various types of pluralism is alive and going, and my purpose in the following section is not to provide a knock-out argument against ontic pluralism, but to suggest that the popularity it enjoys in current literature is not justified and that, given its apparent implausibility, the burden of proof as far as medical causation is concerned is on the side of those separating mechanisms from difference-making.

I should stress at the end of this introductory section that, as far as I would understand it,¹⁷ EBM “+” does not militate for any revolutionary change in medical epistemology, but for making explicit what is already tacitly present and currently accepted in medical practice—namely a to-and-fro movement between mechanisms and population studies, in which the difference-making obtained by a study of the former is used to refine and augment the precision of the difference-making obtained by the latter, precisely on grounds that the former and the latter refer to the same causal relations. Mechanisms are tacitly so pervasively present in medical education and practice, and in the design, implementation and interpretation of population studies,

¹⁶ Howick et al. (2013) Howick (2011).

¹⁷ I have no assurance that the members of the EBM+ movement will agree with my understanding; it is however a project that welcomes a plurality of perspectives on how to bring more into focus the role of mechanisms.

that making this tacit presence explicit might even seem superfluous; it is in this way that one can explain the long-standing lack of reaction of the general medical community against the apparently ‘positivistic’, ‘theory-free’ presentation of the purpose and means of EBM, which thereby privileges only population studies.¹⁸ As for as the theorists of EBM themselves, as I said, this lack of reaction also has its source into an unclarified notion of ontic causal pluralism, to whose main philosophic representative I turn next.

3 Mechanisms and Difference-Making

In his sophisticated and widely cited 2004 paper, Ned Hall has argued for a thesis that by and large, might well be true—namely that we have different concepts of production and difference-making. The concept of production entails that causes and effects are *transitive*, in the sense that if c is a cause of e_1 and e_1 is a cause of e_2 then c is a cause of e_2 , *local*, in the sense that causes are connected to their effects via spatiotemporally continuous sequences of causal intermediates and *intrinsic*, in the sense that the causal structure of a process is determined by its intrinsic, non-causal character (together with the laws); on the other hand, the concept of difference-making or *dependency* entails the counterfactual dependency between distinct events as sufficient for causation, and it goes hand in hand with the thesis that *omissions* can both cause and be caused.

However, Hall also argued that these concepts point to different causal relations. *Locality, transitivity* and *intrinsicness* (which make up the concept of production) as well as *dependency* and *omissions* (which are part of the concept of difference-making) are taken to be five true theses about the causal realm. Since there are cases in which only the first three (concerned with production) could be true, but not the fourth (cases of pre-emption), and others in which only the last two, concerned with difference-making, could be held (cases of causation by absence and double prevention—the latter being presented by Hall as a ‘previously neglected’ but decisive set of cases), we need to accommodate different causal relations.

“One immediate result is that counterfactual analyses are doomed to failure (unless, as I think, they are understood to be *narrowly targeted at just one kind of event-causation*): for they need the first three theses if they are to cope with the well-known counterexamples involving overdetermination, but they cannot abide these theses if they are to cope with the counterexamples involving double-prevention (or, for that matter, if they admit omissions as causes and effects). For the five theses I have mentioned are, I claim, all true. Given the deep and intractable tensions between them, that can only be because they characterize distinct concepts of causation. Events can stand in one kind of causal relation—dependence—for the explication of which the counterfactual analysis is perfectly suited (and for which omissions can be perfectly suitable relata). And they can stand in an entirely *different kind of causal relation*—production—which requires an entirely different kind of analysis (and for which omissions are not suitable relata). Dependence and Omissions are true of the first of

¹⁸ For a very interesting survey of the historical evolution of EBM and its presuppositions, see Vandembroucke 1996, especially pp. 1137–1138

these *causal relations*; Transitivity, Locality, and Intrinsicness are true of the second” (Hall 2004, §1, p. 181, italics and bold added).

Hall says directly that the two concepts of causation pick out different causal relations (“events can stand different kinds[s] of causal relations”), and that while these causal relations might “coincide” or not, each of them, on its own, counts as “causal”. These look like unambiguous ontological statements. One might be tempted to think here that Hall restrains himself to conceptual analysis solely, and that his statements of different causal relations being picked out bears no ontological import. In this reading, talk of causal relations between things or events would be reducible to non-committal claims about events of objects being grouped in sets on ground on the relation between them. There are two problems with this reading, however. The first is that Hall does not strain himself to such intricate steps in order to eliminate the ontological consequence of his arguments; on the contrary, in the final part of his paper he concedes in part that all he might have shown is that there are two kinds of causal relations, but not two concepts.¹⁹ The second is that, even if one reads ‘causal relation’ here in terms of sets, this reading, with a definite Humean flavor, of course, still amounts to an ontic or ontological statement, just like the Quine-like characterization of properties in terms of sets is still an ontological claim about the nature of properties—and indeed, it corresponds to the nominalistic interpretation of universals that spans a long history behind. Be that as it may, Hall’s conceptual pluralism has been in fact taken to entail ontic pluralism,²⁰ and we should note that in general, conceptual pluralism normally presupposes ontic pluralism; it is on this line that I am going to analyze it.²¹

Now, as Hall himself admits, it might well be that both of his concepts of causation could be satisfied in the same causal situation, i.e., that the two types of causal relation might well ‘coincide’,²² which brings me to the (quite local) question I am asking in this section—are Hall’s arguments sufficient in order to separate mechanisms in medicine from difference-making? On an attentive look, Hall’s arguments actually allow for this possibility. That is because, as already seen from the passage just quoted, Hall allows that cases of pre-emption could be treated by a counterfactual analysis that would include as support in its analysis the first three theses, that is to say, that would accept production along with counterfactual dependency, in order to have causes as making a difference to the effects.

“One immediate result is that counterfactual analyses are doomed to failure (unless, as I think, they are understood to be narrowly targeted at just one kind of event-causation): for they need the first three theses if they are to cope with the well-known counterexamples involving overdetermination” (Hall, *ibid.*).

¹⁹ Hall 2004, §6

²⁰ See Godfrey-Smith 2008 for instance.

²¹ As Russo and Williamson say “Conceptual pluralists normally hold that each concept of cause picks out a different causal relation—i.e., conceptual pluralism normally presupposes ontological pluralism. However, it is possible to be an ontological pluralist without being a conceptual pluralist by maintaining that there are two causal relations ambiguously picked out by a single concept of cause. Both varieties of pluralism are, we argue, implausible.” (Russo and Williamson 2007, p. 165, bold added).

²² And it should come as no surprise that the distinction between production and dependence has gone unnoticed, for *typically* the two relations *coincide* (more exactly, I think, production typically *coincides* with the ancestral of dependence (Hall 2004, §6, italics added); see also Longworth 2006, p. 55.

Stricto sensu, this is a violation of the **dependency** thesis (thesis 4) as stated by Hall, which says that counterfactual dependency is *sufficient* for causation. However, it suits perfectly an account of mechanisms in which difference-making is a component of the causal relations, along with the production aspects traditionally associated to mechanisms. As Hall himself notes, cases of early pre-emption—i.e., cases in which the actual cause cuts the potential causal process before the actual causal process runs to completion, as in the stock example: Suzy throws a rock at a bottle shattering it, and had Suzy not thrown her rock, Billy would have shattered the bottle by throwing his rock—can be dealt with by bringing in the thesis of **transitivity** (thesis 1) and thereby introducing chains of causal dependency that isolate intermediary events between the actual cause and the actual effect whose counterfactual absence does not involve the actualization of the potential cause.²³ On the other hand, cases of late pre-emption, in which the actual cause cuts the potential causal process only after the actual process completes (e.g., Suzy throws her rock a split second before Billy and her rock gets there first, shattering the bottle) can be dealt with by introducing the criterion of **intrinsicness** (thesis 3), and maintaining that the process linking the actual cause with the actual effect has an intrinsic character such that, given the same laws of nature, the actual cause would have produced the actual effect in scenarios in which the potential cause had not been present (i.e., Suzy would have shattered the bottle in scenarios in which Billy was not present at all).²⁴ Again, in Hall's words:

Perhaps the counterfactual analyst can use the Intrinsicness thesis to handle the problem of Billy and Suzy. After all, in the alternative circumstances in which Billy's throw is absent, it seems correct to say that the causal history of the shattering (back to the time of Suzy's throw) consists exactly of those events on which it depends. What's more, this structure matches a structure that takes place in the actual circumstances, where Billy's throw confounds the counterfactual relations; Suzy's throw, being a part of this structure, will therefore count as a cause of the shattering, thanks to the Intrinsicness thesis. To be sure, this is no more than a suggestion of a revised analysis. But again, what is important is that it is a plausible and natural suggestion to pursue—and it gives a central role to the Intrinsicness thesis (Hall 2004, §3, p. 10).

This solution to the late pre-emption—which might appear as quite unsatisfactory for a 'pure' counterfactual account of causation, but which, again, suits perfectly the account of mechanisms I am adverting here—can also be useful to draw attention on a different way in which counterfactuals find their way into the relations of production, in

²³ “There is an obvious strategy for handling this kind of case. First, we liberalize our analysis, by taking causation to be the ancestral of counterfactual dependence: *c* is a cause of *e* iff there are events *d*₁, ..., *d*_{*n*} such that *d*₁ counterfactually depends on *c*, *d*₂ depends on *d*₁, ..., and *e* depends on *d*_{*n*}. Next, we look for an event *d* (or sequence of events) intermediate between the preventing cause *c* and the effect *e*, such that *e* depends on *d* and likewise *d* on *c*. The strategy works handily in the case before us” Hall 2004, §3.2; the solution was presented by David Lewis, in his 1973, p. 567

²⁴ Again, this was David Lewis's solution, as mentioned by Hall “Suppose we have processes—courses of events, which may or may not be causally connected—going on in two distinct spatiotemporal regions, regions of the same or of different possible worlds. Disregarding the surroundings of the two regions, and disregarding any irrelevant events that may be occurring in either region without being part of the process in question, what goes on in the two regions is exactly alike. Suppose further that the laws of nature that govern the two regions are exactly the same. Then can it be that we have a causal process in one of the regions but not the other? It seems not. Intuitively, whether the process going on in a region is causal depends only on the intrinsic character of the process itself, and on the relevant laws. The surroundings, and even other events in the region, are irrelevant.” Lewis 1986, p. 205., *apud* Hall, §3.

cases of pre-emption and also beyond. In defining the **intrinsic** character of this relation of production in terms of “the causal structure of a process is determined by its intrinsic, non-causal character (*together with the laws*)”, Hall appeals to the laws that govern such processes. However, laws are intimately connected with the problem of the modal character of causation, which is the problem that there should be some way to differentiate accidental successions of events from non-accidental causation. This aspect is noteworthy because the easiest way to express the modal character of causation, i.e., how causes *could* produce their effects, even if they are not actually doing so, is by way of counterfactuals, where counterfactuals in turn could stand as symptoms for the laws of nature at work in the respective circumstances, for causal powers manifesting themselves, for relations of similarity between possible worlds according to a neo-Humean account of causation, etc. Importantly, the main reason why ‘pure’ accounts of production such as Salmon’s seem to have had problems in exhaustively describing causation consists precisely in capturing this modal character.²⁵

Hall does not discuss this aspect, but, as mentioned previously, Stathis Psillos has provided a very interesting discussion of the relation between production and counterfactuals, all the more so since production is treated by Psillos in the framework of the complex systems accounts of mechanisms. *In nuce*, Psillos argues that the complex systems accounts of mechanisms cannot dispense with counterfactuals, because they appeal or *must* appeal to laws in order to back up the relationships between components of mechanisms, as in Glennan’s approach in which the interaction of the parts of the mechanism are characterized in terms of invariant, change-relating generalizations,²⁶ and counterfactuals are the essential ingredient for the modal character of these laws. On the other hand, for complex system theories that ground the functioning of mechanisms on activities (as in the MDC approach, in which we have entities and activities organized such that they are productive of regular changes from start to finish conditions) these activities could only be manifestations or actualisations of modal capacities defined counterfactually (on an anti-Humean approach), or alternatively, *façons de parler* for modal successions (on a Humean interpretation)²⁷; moreover, the interactions between mechanism components need a counterfactual approach as well.²⁸

²⁵ Psillos (2004, p. 293, pp. 36-38).

²⁶ Psillos argues that Glennan’s move to explain the counterfactual feature of laws in terms of mechanisms (in Glennan 2002) stands in tension with his earlier (correct) view in Glennan 1996 that at the level of fundamental physics we do not have mechanisms (Psillos 2004, p. 308). Counterfactuals also show up more directly in Glennan’s approach to mechanisms in the sense that, in Woodward’s sense, the invariance of the relationships among the parts of the mechanism should be established under actual and *counterfactual* interventions. While seemingly allowing for a counterfactual analysis of the interaction between mechanism components, Glennan does not characterize the final output of a mechanism in terms of difference making, as noted in the previous section.

²⁷ ‘The second argument that MDC offer is this: “state transitions have to be more completely described in terms of the activities of the entities and how those activities produce changes that constitute the next change” (2000, p. 5). Here, the emphasis is on the production. As they explain, activities add the “productivity” by which changes in properties (state-transitions) are effected. But isn’t this question-begging? Many would just deny that there is anything like a productive continuity in state transitions. All there is, they would argue, is just regular succession (or some kind of dependence). In any case, the friends of capacities would argue that there is productive continuity in state transitions, but that this is grounded in the natures of the entities engaged in state transitions. If water has the capacity to dissolve salt, and if this capacity is grounded in the natures of water and salt, then all that is needed for the dissolution of salt in water (that is, the activity) is that the circumstances are right and the two substances are brought into contact.’ Psillos (2004, p. 314).

²⁸ Psillos (2004, p. 314).

If production, either in its austere, physical form (as advocated by Salmon and Dowe), or in the form of complex biological systems analysis, cannot arguably dispense with counterfactuals (even in the salient or ‘easy’ cases where production seems to be sufficient to characterize causation), the same could in turn be said about the counterfactual relation, vis-à-vis its putative independence from production. We have touched above upon the difficult cases for counterfactual analysis, namely those of pre-emption, in which dependency and production need to join hands. However, even in the salient or ‘easy’ cases in which the counterfactual analysis appears to go as smoothly as it could possibly be (namely in the cases of causation by absence and double prevention), production shows its head, or it could be plausibly invoked as the necessary attendant of the counterfactual dependency.

Incidentally, medicine is rife with assertions of causation by absence. Take the example of the claim that lack of C vitamin causes scurvy. This claim, with an easy counterfactual rendition, picks up a causal relation for which a description in terms of production is *also* readily available—the formation of spots on the skin, spongy gums, and bleeding from the mucous membranes, defective collagen fibrillogenesis, etc. In fact, the story is much broader, because the whole field of medicine can be viewed as resulting from causation by absence. Aristotle had already defined the disease in *Metaphysics H* as the privation of health, and in contemporary medicine each and every pathological process admits a causal rendition in terms of the dysfunctioning, or the lack of normal functioning of a certain organ, system, tissue, type of cell, etc. Surely, we would not want to say that medical causation consists solely in causation by absence; the entire array of physiopathological processes involving production stands behind any claim that dysfunctioning or lack of functioning has produced such and such harmful effects. Furthermore, any actual or prophylactic medical treatment follows counterfactually the model of double prevention—the harmful causes that prevent or would prevent health are in turn prevented by medical measures—administrating such and such drugs, intervening surgically on such and such organ, etc. And again, it would be at least far-fetched to say that the whole array of therapeutic treatments are to be reduced to such counterfactual descriptions of double prevention, in the face of the innumerable processes of production that modern medicine has discovered in relation to its treatments. Hall’s insistence on the ‘previously neglected’ cases of double prevention that would set apart counterfactuals and production is I suspect slightly misguided, and one should not dismiss the possibility of finding production alongside the double prevention cases as amounting merely to ‘twists and turns’.²⁹

²⁹ “Of course, the debate can take various twists and turns from here: there are further stratagems one might resort to in an effort to interpolate a sequence of omissions between Billy and the bombing; alternatively, one might deny that causation without a connecting sequence of causal intermediates really is sufficient for action at a distance. It won’t profit us to pursue these twists and turns (but see my 2000b); suffice it to say that the stratagems fail, and the prospects for a replacement for the sufficient condition seem hopeless.” (Hall 2004 §4.1). A word of warning is in place here—talk of absences as causal factors is in itself very problematic, because one needs to show what the terms designating absences refer to, and this is an age-old problem, that starts from Plato’s *Sophist* and Aristotle’s *Physics* 1.9, going up to Russell’s semantic struggles with assertions of negative existentials or Wittgenstein’s ‘derivative’ explanation of negation; see Pippin 1979; Swindler 1980. Again, starting from Plato’s ‘in a certain sense not-being is’ as uttered by the Stranger of the *Sophist*, all the solutions I know of for this problem work by interpreting (the reference to) non-existence or non-being as oblique references to existence or being, which of course parallels the above interpretation of absence-based counterfactual relations as standing for existing processes of production.

Michael Strevens has provided in fact a framework to accommodate these different levels of production and (counterfactually defined) causation by absence. Again, in brief, Strevens argues that we should distinguish between low-level fundamental events at reality's bedrock, on the one hand, and the high-level, non-fundamental events involved in our normal counterfactual assertions, on the other. The fundamental events are *concrete*, and they involve *causal influence*, where this causal influence could be covered by a process approach, a counterfactual approach, or a nomological dependence approach, and should provide us with 'picture of the physical world as a densely reticulated web of causal influence relations among reality's fundamental constituents'.³⁰ On the other hand, the high-level events are *abstract*; in particular, the events making up our normal counterfactual claims are obtained by isolating a chunk of the causal web, producing a model that deductively entails the effect from the cause and represents the causal influences in the chunk, and further abstracting this model as much as possible without invalidating the causal entailment of the effect, until we obtain the difference-makers.³¹

Strevens' example is that of the breaking of a vase by a heavy rock. The model corresponding to this concrete event will describe in its (sets of) premises "in exacting, excruciating detail"³² the properties of every molecule and fundamental particle in the rock, the particles of the vase and the surrounding air. To obtain the high level events and features, these premises will be replaced by more abstract statements, which are entailed by the premises and have the same subject matter or a part of it, without invalidating at the same time the causal entailment of the vase's breaking. One arrives thereby at the difference-makers from the high-level – the approximate size, consistency, mass, and velocity of the rock. Depending how large the chunk of the causal web is selected, one can start from Suzy's act of throwing the rock, isolating other high-level factors,³³ or, even if Strevens does not say it here, from any preventative factor of Suzy's throwing or preventative of the preventative (double preventive process). What Strevens does say in *Causality Reunited* is that omissions and absences qualify as causal difference-makers on his account, insofar as, by the process of abstraction, parts of the causal web can be

³⁰ "To complete the story, a metaphysics of causal influence is required, [...] Suppose we have the correct account of causal influence in hand; we can then use it to see what the influences are according to modern physics. Presumably the exchange of bosons will play a central role, as will the interaction between mass and space-time, and ultimately the nodes of influence may be values of fields at space-time points rather than particles. I will not dwell on these matters here; that I leave to philosophers such as Dowe. However physics turns out, I expect it to give us a picture of the physical world as a densely reticulated web of causal influence relations among reality's fundamental constituents." (Strevens 2013, pp. 307-308)

³¹ In fact, Strevens discusses two types of difference making—those that result directly from the abstraction of the model representing causal influences in the chunk of the causal web (resulted from the application of what he calls the 'kairitic' criterion) and one that is also derived from and relative to a normative framework, within which values also find their place (Strevens 2013, pp. 306, 315–319). The distinction between two types of difference making does not matter though to the point I draw here from Strevens' work.

³² Strevens 2013, p. 311.

³³ 'More generally, a high-level event is a cause of the breaking if there exists some causal model, corresponding to some section of the relevant causal web, in which that event appears when abstraction is done.' (Strevens 2013, p. 311)

ascribed negative properties and be considered causes.³⁴ In a different publication, as a response to Hall's critical review of his recent volume on scientific explanation,³⁵ Strevens has provided a response to the problem of pre-emption, which goes exactly on the lines described above.³⁶

Thus, we end up with an entirely coherent unificatory ontic alternative to Hall's ontic pluralism, where this unificatory ontic picture is consistent with conceptual pluralism, i.e., with ways of describing causation, be it pragmatically, epistemically, methodologically, etc. This is how Strevens puts it:

Let me pause now to respond to Hall's argument for causal pluralism. That argument, you will recall, proceeds by showing that certain intuitively compelling theses about causation cannot be satisfied simultaneously by a single relation; rather, the theses fall into two groups, corresponding to conceptually and (presumably) metaphysically distinct relations. I believe that the argument is correct in every important way. There are two kinds of causal relation. One has the properties of locality, intrinsicness, and transitivity, while the other is a difference-making relation that allows omissions and other absences to play a causal role, and for the existence of which counterfactual dependence is typically sufficient. The first is causal influence; the second is causal difference-making. *But this causal duality does not amount to causal pluralism.* Why not? On my picture, the world has a single, objective, homogeneous causal structure: the web of influence. (thee homogeneity consists in each strand of the web having the same essential nature, to be specified by the correct metaphysics of influence: conserved quantity transmission, nomological dependence, or whatever.) Where does causal difference-making fit in, then? That one high-level event makes a difference to another high-level event is not a fact that exists over and above the web of influence, constituting a distinct kind of causal being; it is, rather, an abstract fact about the web of influence itself. Facts about difference-making stand to facts about influence in the same about a center of mass summarizes certain physically pertinent information about the underlying distribution of mass, so a fact about difference-making summarizes pertinent information about the fundamental study of causation, the network of relations of influence.[...] My two causal relations, then, do not imply as Hall's do distinct, coexisting causal structures. There is, at base, a single kind of causal stuff—*influence*—with a single kind of fundamental structure—the web. There are, however, two ways of describing the web. The first is by stating facts about influence; these may be used to provide a complete inventory of causal reality, without remainder. The second is by stating facts about difference-making. Such facts do not add anything to what would be contained in a complete enumeration of the connections in the web, any more than facts about center of mass add anything over and above what is contained in a complete description of the universe's distribution of mass (Strevens 2013, pp. 313–314).

³⁴ 'Can omissions or absences be causal difference-makers, and therefore causes? Yes: after the kairetic account's abstraction procedure has been applied, a piece of the causal web may be ascribed negative properties, and these may properly be cited as causes. In the case of the rock and the vase, for example, the abstract model will specify that there were no significant forces on the rock other than the earth's gravity. The absence of such forces—forces sufficient to deflect the rock from its vase-ward trajectory—is therefore one of the causes of the vase's breaking.' (Strevens 2013, p. 313).

³⁵ Hall (2012)

³⁶ Strevens 2012, pp. 496–501.

Again, these do represent knock-out arguments against ontic pluralism, and one might disagree with parts of Strevens' scheme,³⁷ but they do undermine the inevitability with which certain philosophers have viewed this ontic pluralism,³⁸ i.e., they cut out the seemingly inevitable relationship between conceptual pluralism and ontic pluralism about causation. Perhaps, given the correlations of the Einstein-Podolsky-Rosen/Bohm experiment, one should entertain the thought of ontic causal pluralism in physics, where one would seem to need a purely counterfactual relation to account for causation at distance, although, as it well know, this is a matter of huge controversy.³⁹ However, in medicine at least, given the largely accepted monistic premises of medical theory and practice, including the practice of aggregating evidence, the ontic pluralism is an extremely implausible theory, and the burden of proof is on the side of those propounding it. If production and dependency are two distinct types of ontic causation, *how* could they (or evidence of them) ever be aggregated in medical studies? If they are two distinct types of ontic causation *why* would ever want one to aggregate them, given that each should count as sufficient for causal relationships? Whether or not one agrees with RWT, according to which one *should* have evidence of both difference-making and production, it is commonplace that de facto the two have been often aggregated in the history of medical discoveries. But why? Again, we should come back to the arguments provided by Russo and Williamson, the authors of RWT:

“But there is a second problem that besets pluralism, namely, that it inherits the difficulties of monistic accounts. This problem affects the pluralist who notes that there are two types of evidence for causal claims—mechanistic and probabilistic—and concludes that there are two types of causal claim, mechanistic and probabilistic. This is clearly a fallacious inference, and, worse, opens the pluralist up to the objections of section 5. Suppose that the pluralist advocates two notions of cause, a mechanistic, cause1, and a probabilistic, cause2. Take any particular causal claim, e.g., ‘smoking causes cancer’, that the pluralist cashes out in terms of one or other of these notions but not both (there must be some such claim, for otherwise she is not a pluralist but rather takes causality to be one thing that has two aspects or components). Now the evidence for this claim is multi-faceted, consisting of observed dependencies and mechanistic/theoretical considerations. But the pluralist’s analysis of this claim will be single-faceted, say ‘smoking is a cause1 of cancer’. But then the pluralist opens herself up to the epistemological problems of monism. If this particular use of ‘cause’ is mechanistic, cause1, then how can it be that, even when the mechanism is established and uncontroversial, further probabilistic evidence is cited in support of the causal claim? However, if the use is probabilistic, cause2, why are mechanisms invoked as evidence, even when there is ample probabilistic evidence? The pluralist can’t explain the variety of evidence for the claim: if pluralism is right, it should be possible that the evidence

³⁷ For instance, in special sciences like biology, as far as the *concrete* events are concerned, one might want to descend just down to the level of fundamental biology.

³⁸ ‘The two important, independent positive arguments for disunity to be considered here are Hall’s argument for two concepts of causation, and a collection of arguments purporting to show that causal facts are responsive in part to facts that concern what is in the broadest sense normal, and so depend on the complex and heterogeneous domain of normative facts. This paper will sketch a theory of causation that accepts the premises of both arguments, yet that is simple, principled, and unified. Whether you think that the theory will ultimately succeed or not, then, you should conclude at least that *the premises do not lead to disunity with the inevitability that some philosophers have supposed*’ (Strevens 2013, p. 301, italics added).

³⁹ Berkovitz 2007.

just be mechanistic, or just be probabilistic. So, while the pluralist may say that different uses of the word ‘cause’ can refer to different relations, some particular use must refer to a single relation” (Russo and Williamson 2007, pp. 166–167, italics added).

Russo and Williamson go on to argue that the epistemic account of causation can accommodate both evidential pluralism and ontic monism, which is a perfectly coherent solution, and which provides not only an ontically monistic reunification, but also a conceptual one, since the concepts of difference-making and production are further subsumed under the concept of causation derived from the inferential practices of an omniscient being, while evidential pluralism remains in place. But clearly, since both difference making and production are viewed as part of the same causal relation, then, on their reading as well, difference-making should be integrated into mechanisms; otherwise we would have again two different causal relations that somehow coincide when evidence is aggregated, and the set of questions they raise—as for instance the question as to the need of evidence of difference-making when a mechanism seems apparently establish—would surface again.

As I mentioned before, whether difference-making (and production) should be given a Humean reading or not is beyond the scope of this paper. It would certainly be advantageous for our practices and arguments if ontic monism was attended also by conceptual monism, but that would require a definite metaphysical story about what causation *is*, and a discussion about medical causation and mechanisms need not go so far. Strevens unifies ontically the causal relations, while remaining agnostic as to what the “web of causal influence” really amounts to and while agreeing with Hall that we do have two concepts of causation (presumably reading them as methodological and pragmatic); Russo and Williamson unify ontically the causal relations and provide also conceptual reunification, siding with the Humean *ontological* injunction that causation is not a relation in nature. But we need not enter into that discussion in order to say something significant about medical mechanisms, tempting as it might certainly be to follow the Humean path.

Think about the sort of interrogation that happens so often in the Platonic dialogues, where Socrates asks his interlocutors what the very definition of a certain kind is, as in the Theaetetus for instance, where the question is how knowledge should be defined, and where the young Theaetetus responds innocently that it is to be found in geometry and would include the art of the cobbler and other craftsmen - Socrates commenting ironically that they are being offered too much. Yet, it has often been noted that the sort of definition Socrates asks might just be too much to ask or impossible to respond.⁴⁰ Theaetetus certainly has a point that the art of the cobbler is knowledge, a simple one—the art or science of making shoes, even if he cannot raise to the high expectations set up by the general enquiry of the dialogue. One finds oneself in an analogous situation when affirming that both production and difference-making should be parts of medical or biological mechanisms. ‘But what *is* causation?’ it will immediately be asked. ‘What is our *concept* of causation?’ ‘How about causation at a distance in physics?’ These are questions are legitimately asked, and should be investigated,⁴¹ but one needs not

⁴⁰ See for instance Ryle (1966)

⁴¹ My tendency would be to say that this very concept answering to the essence of causation would correspond to the manifestation of a capacity or a law, where difference-making and production would be the *propria* for such manifestations. But to paraphrase again Socrates from *Gorgias*, I aim to be among those who takes pleasure in being shown to be wrong.

respond to them definitively in order to make a point about difference-making and medical mechanisms, which is just that, again, on an ontic monistic reading, medical mechanisms should not be separated from difference-making.

Coming back to the down-to-earth level, it might be that in certain cases, their difference-making is *counter-factual*, in which case evidence of difference-making from population level is evidently needed.⁴² In most cases, although *factual*, their difference-making needs to be weighed against the difference-making obtained at population level, given the problems of complexity and masking. In certain other cases, it is the difference-making from the population level that is weighed against the mechanistic difference-making, as in the difficult situations where population studies show contradictory results, have indefinite results, are not sufficiently targeted against the precise control and test groups, etc.⁴³ Finally, in extremely rare cases, knowledge of mechanisms will suffice to draw causal conclusions, in the absence of population studies, because the difference-making grasped at the mechanistic level is sufficient, and I will discuss these cases in the next section, in the context of Howick's criticism of RWT.

4 Exposing the Vanities Indeed

In espousing the answer that can be given to Howick's criticism, it should be said that on the view of mechanisms as linked to difference-making, the epistemic claim of RWT, namely that we need evidence of *both* production and difference-making (since taken on their own, neither production nor difference-making alone can rule out spurious causation, accidental conjunction, wrongheaded difference-making) is also grounded in the ontic elements of the causal relations studied by medicine—among others, we need them *both* epistemically, because ontically they are both part of the subject-study in question. In other words, RWT and ontic monism about causation go hand in hand. They do not entail each other but ontic monism is the natural interpretation of RWT's epistemic claim, and it has been like this from the beginning, albeit in a specific Humean form. Note that when mechanisms as such are invoked, or when 'evidence of mechanisms' goes into the formulation of RWT, then, given that epistemically the difference-making and production of mechanisms needs to be complemented by the difference making of population studies (just like, to put this

⁴² Incidentally, it should be said that the criticism advanced against the MDC account that in certain mechanisms, e.g. the homeostatic ones, 'nothing happens' (e.g. Skipper and Millstein 2005) does not affect the account of mechanisms presented above. The MDC account requires activities; the above account requires difference-making, which should be in place at least counterfactually.

⁴³ It might be that we get at population level will not be the real difference making for the causal factor C or a mechanism M (which is what might interest us precisely in medical treatment) but a difference making that is the summation, the net effect in the context of the whole organism where other mechanisms show up, a difference making that can be further mitigated by the very fact of considering it at the level of populations, etc. We will get a global effect E which is statistically relevant (on the level say, of generic causation, whatever that might be) but incorrect for the level of what C does, within a particular organism, where the singular causal relations, the real ones, take place. That is to say, if complexity and masking constitute problems for mechanisms, they can also be problem for difference making assessments at population levels as well, and one should add to this the sort of conundrum for RCTs espoused in Worrall (2007). When such problems show up for population assessments, it is mechanisms which can alleviate them, as for instance the history of atherosclerosis clearly shows.

conversely, the difference making of population studies needs to be complemented by the production and the difference making of mechanisms), RWT states that both mechanisms and population level dependencies are necessary, and I will be using the two formulations alternatively, as the context will require.

Now, Howick is not an explicit advocate of the ontic pluralistic view of causation. Rather, in reading him, one would say *prima facie* that he tacitly adopts a monistic view.⁴⁴ But then things get complicated, because in some arguments he also adopts tacitly a pluralistic view. So my argument in this section has the form of a disjunction, which I will express more generally, since I take Howick to be an (open-minded) representative of a way of thinking current in the EBM area. The disjunction goes as follows: in the framework of EBM, either one adopts an (ontically) pluralistic view of causation, which faces unpalatable consequences from a philosophical point of view and goes against the entire medical practice of treating causation monistically, or, one adopts a monistic view of causation, in which case one should adopt the Russo-Williamson thesis and take mechanisms as (at least) necessary for establishing causal claims.

Recall from §1 that, granted that mechanisms are necessary for establishing causal claims, how much we need to know about them may differ. In cases of diagnosis and explanation, it might just be sufficient to know that the mechanism exists, i.e. that the ‘black box’ is in place, because this will offer us evidence of production. In difficult cases of prediction, in order to face the problem of extrapolation, and/or when the results of population studies are ambiguous, inconclusive, or show a failure of expected results, one will want to open the black box and know as much as possible the details of mechanisms, because in this way, evidence of difference-making from the microstructural level will usefully complement the difference-making of the population studies. In both cases, given the necessity of both production and difference-making, and the (extremely probable) insufficiency of difference-making from the microstructural level, mechanisms are necessary but not sufficient for causation claims; they (almost) always need to be attended by population level assessments. And a consequence of the ontic causal monism I have advocated above is that, in both cases, we gain evidence of mechanisms from evidence of difference-making, even if mechanisms are not *just* difference-making, of course, and *even if* evidence of difference-making is obtained from population level studies.

Notice that this move is not allowed to the pluralist. On a pluralist view, strong evidence of difference-making from population level could not be evidence for the existence of a mechanism, if mechanistic causation is a different type of causal relation from the difference-making one. Neither would it be warranted, from a pluralistic

⁴⁴ A monistic account that would seem to follow from his general adherence to the actual practice of medicine and also from his definition of mechanistic reasoning, which is provided after surveying the main mechanistic accounts in the philosophical literature “Mechanistic reasoning is an inferential chain (or web) linking the *intervention* (such as HRT) with a patient-relevant outcome, via relevant mechanisms” (p. 929, italics added). That is to say, Howick does not define mechanisms in terms of difference-making, but asks that mechanistic evidence of causation be understood in terms of interventions—which is a special type of counterfactual approach. This bit on mechanistic reasoning is in fact ambiguous between the definition of mechanism in term of production alone and that in terms of production and difference making, and at any rate, as we shall see, Howick is also ambiguous between the monistic and pluralistic views on causation. Note parenthetically that Howick’s approach to mechanistic reasoning asks too much, since it implies that in order to count as evidence, mechanisms should solve on their own the problem of extrapolation. More on this will be discussed towards the end of the present section.

standpoint, to say that microstructural evidence of a mechanism eliminates claims of spurious causation made from population level, if mechanisms and difference-making population results do not point to the same causal relations.

This is the root of the problem. Howick is indeed an open-minded EBM proponent, in that he grants some role to mechanisms,⁴⁵ and one has much to learn from his discussion of the difficulties of extrapolation and how mechanisms are involved with it. But his arguments against RWT rest on an ambiguity between the pluralist and monist views on causation, in the sense in which he blocks the inference from population studies to a mechanism, which is only possible on a pluralistic stance. Thus, for instance, Howick writes that even if cases such as gastric or duodenal ulcer seem to point to the need of mechanisms, there are other “anecdotes” that might point into the different direction, in the sense in which, correlations or difference making alone seem to have been sufficient for causal claims.

Yet these same anecdotes could be used to point out the dangers with requiring mechanistic reasoning alongside comparative clinical studies to establish causation. Countless mothers and babies would have been saved had Semmelweis’s intervention been adopted after the results from his comparative clinical study became available, and life expectancy in many countries would have risen decades earlier if government strictures on smoking had been introduced before the mechanism linking smoking and lung cancer was established. Hence, Gillies (whom Russo and Williamson ironically cite as a source) uses the same cases to argue that it is unwise to require mechanistic reasoning when there is strong evidence from comparative clinical studies (Howick 2011, p. 930).

What Howick says about correlations separated from mechanistic reasoning being sufficient for causal conclusion is very interesting and we should be grateful to him for bringing this aspect into discussion. However, one can dispute his conclusion on two counts. The first is that, on a pragmatic scenario,⁴⁶ we could act according to incomplete knowledge (i.e., in the absence of detailed knowledge of mechanisms or not knowing if mechanisms exist, which entails not being sure that spurious correlations or wrongheaded difference-making have been ruled out) since no harm could ensue from prevention against potential factors of morbidity; protecting against what reasonably appear to be risk factors could not constitute the basis for a theoretic argument against RWT, which is concerned with *establishing* causal claims. The second is that disregarding any inference to (or evidential relevance for) mechanisms starting from rigorously established difference-making at population level can be blocked only on a pluralistic view of causation, as argued above.

The same holds for Howick’s argument that the mechanisms of some drugs are not completely understood, and yet they are being used based on clinical studies alone, and have been used at a time when one had no idea about their mechanism at all. This deserves a bit of clarification. Again, for cases of diagnosis and explanation, one just needs to know that a mechanism exists, and evidence for the existence of a mechanism can from rigorous population level assessments; this evidence of difference-making, *together* with some minimal insight that the microstructural entities at stake at really

⁴⁵ Howick (2011, p. 927).

⁴⁶ A pragmatic scenario that Howick himself seems to admit, distinguishing between ‘practice’ and ‘theory’; Howick (2011, p. 931).

there, can establish the existence of a mechanism.⁴⁷ For cases of prediction, where there is the problem of extrapolation, one *should* know more about mechanisms, and Howick's examples, drawn especially from psychiatry and neurology, point precisely to an area in medicine where we would want to know more but unfortunately at this stage we do not.

"Moreover, there are many counterexamples where medical interventions have been accepted on the basis of evidence from comparative clinical studies alone. To name a few (see Howick 2011 for a more complete list), aspirin was used for a century before its analgesic mechanism was identified, the mechanism for general anesthesia is still not well understood, and deep brain stimulation (DBS) is currently used to suppress tremors in patients with advanced Parkinson's and to cure other motor function disorders such as Tourette's Syndrome, yet researchers have not been able to identify the mechanism of DBS with any certainty." Howick 2011, p. 930

The scenario in which the inference to mechanisms from (rigorously established) difference-making from a non-microstructural level is blocked comes with a vengeance from the comment of a personal communication on RWT.

'Jon Williamson (in personal correspondence) worries that I have misunderstood his position. He claims that [" "] the most salient mechanistic evidence is *evidence that there is an appropriate mechanism linking the putative cause to the putative effect*. This isn't evidence from a mechanism to intervention claims, but evidence that supports the claim of the existence of a mechanism. Such evidence could be got from an RCT [randomized trial], for example. RWT [the Russo Williamson Thesis] says that one needs evidence that there is a linking mechanism, as well as evidence that C makes a difference to E. So my main worry is that we are talking at cross-purposes![" "] This understanding of mechanistic evidence seems to overlap substantially with my mechanistic reasoning (a difference being that mine acknowledges that several mechanisms are often involved). Williamson's emphasis, however, is *on evidence for a mechanism, whereas I emphasize the importance of the mechanism providing evidence*. Unless the evidence for a mechanism is evidence "from a mechanism to intervention claims," I do not see how it helps us establish a hypothesis about the effects of an intervention. In addition, I do not understand how evidence for a mechanism can be gained from an RCT (noncircularly, i.e., without claiming that if an RCT reveals an effect that there must be a mechanism, which is a separate claim). Most importantly, when he claims that one needs a mechanism "as well as evidence that C makes a difference to E" he seems to concede that "difference-making" evidence is sufficient to establish causation' (Howick 2011, pp. 933–934, italics and underline added).

Again, that is the root of the problem—mechanisms are not taken as difference-making. If they are not taken as difference-making, then a good RCT (which establishes difference-

⁴⁷ For DBS for instance we have of course some minimal information on the mechanism; we know at least that for Parkinson syndrome it is targeted at thalamus, subthalamic nucleus, and a portion of the globus pallidus, and we know that these regions are involved in the tremor, rigidity and bradykinesia of the respective syndrome. As for the aspirin example, Howick is of course right, it has been used apparently ever since Hippocrates, under the form of extract from the willow tree. One interesting detail though is that, while the mechanism of its action was not known, its extreme popularity at the turn of the XXth century after being synthesized might not have been always fortunate, and it has been suggested that their use might have increased the mortality of the Influenza pandemic between 1918 and 1919 (Starko 2009). Of course, one would also have to think at how many cases of gastro-duodenal ulcer this use of aspirin must have produced.

making) could not offer evidence for the existence of a mechanism. Howick claims not to “understand how evidence for a mechanism can be gained from an RCT (non-circularly, i.e., without claiming that if an RCT reveals an effect that there must be a mechanism, which is a separate claim). (p. 934)”. But this is not a separate claim at all. Or, differently put, it needs to be a separate claim on an ontically pluralistic view of causation, which, due to its implausibility, carries the burden of proof on the shoulder of its proponent, and which, at any rate, would need to be assumed as such and differentiated from the monistic view.

Prima facie, Howick appears to be right in saying that, if a good RCT is taken to be sufficient to establish that a mechanism exists and to draw causal inferences, then it seems that difference-making is sufficient for causation.⁴⁸ However, Williamson’s claim, if I understand it correctly, is not that a good RCT is sufficient to establish that a mechanism exists. His claim is that an RCT provides “evidence that supports the claim of the existence of a mechanism”. In this vein, evidence of difference-making is also evidence for mechanism; *if* the difference-making grasped from population level were sufficient for causation, *then* evidence of difference making *would* also *establish* the existence of a mechanism.

The trouble is that the difference-making of population studies alone could not establish causation, and that has been one of the main points of RWT. On the one hand, perfect RCTs are almost impossible to realize in practice, because of the infinite randomization they would require in order to eliminate spurious causation, and that is why evidence of production coming from mechanisms is necessary.⁴⁹ On the other hand, and relatedly, as I have mentioned in §1, such evidence of production is *de facto* involved, more or less tacitly, in the current practice of population studies. That is to say, the good RCTs which Howick and the other EBM proponents cite as examples of valid scientific practice almost invariably presuppose extensive knowledge of mechanisms in their design, implementation and interpretation, and one only needs to think here of the extensive laboratory research of pharmaceutical companies before proposing treatments for population testing, biased by financial interest as that research might sometimes be.

⁴⁸ Howick also appears to be right that talk of ‘evidence for a mechanism’ in this context is slightly odd, since mechanisms should provide evidence for causation in conjunction with the population studies.

⁴⁹ Howick admits as true that comparative clinical studies sometimes support spurious relationships, but responds that mechanistic reasoning does not prevent adoption of spurious hypotheses. That is a weak argument, it seems to me, in the sense that, while RWT requires that mechanisms and population level evidence of difference making are both necessary, it is hardly a counter-argument to say that they both might not be jointly sufficient. If mechanisms can rule at least some spurious causation, then surely mechanistic evidence is valuable and should be taken as necessary. “Russo and Williamson (2007) could insist that, for cases in which the mechanisms had not been identified, the medical community should have rejected claims that the treatments were effective, which brings us to their theoretical argument. Following Glennan, Russo and Williamson claim that “if there is no plausible mechanism from C to E, then any correlation is likely to be spurious” (2007, 159). It is true that comparative clinical studies sometimes support spurious relationships.[...] *But mechanistic reasoning does not prevent adoption of spurious hypotheses*. Bloodletting was adopted on the basis of both mechanistic reasoning (derived from the humoral theory) and comparative clinical studies (observations of patients recovering after being bled), yet it is safe to say that bloodletting was useless or harmful in most cases. The example of bloodletting may be unfair because we now know that the evidence supporting its efficacy was weak. But this reinforces a point of this article: quality (tightly controlled, unbiased, etc.) evidence rather than quantity of evidence helps reduce the likelihood of spuriousness.” (Howick 2011, p. 931, italics added)

Moving on, the same equivocation between the monistic and pluralistic view is present also in Howick's discussion of extrapolation, which was also referred to in §1. The problem of extrapolation is viewed by Howick as a problem for mechanisms, and not as a problem of causal relations in medicine per se, which could be approached with the help of mechanisms and population level assessments, as RWT would entail, and as proponents of EBM+ have argued.⁵⁰

EBM proponents are justified in being cautious about mechanistic reasoning for at least two reasons. First, biochemical mechanisms are difficult to identify. Bloodletting, placing babies to sleep on their stomachs, antiarrhythmic drugs, and many other arguably useless or harmful therapies have been adopted on the basis of reasoning from what we now believe to be wrongly identified mechanisms. Unfortunately, the more recent examples that involved more plausible mechanistic reasoning may have led to more acute harm. The apparent knowledge of what happens to some of the mechanisms under intervention lends an aura of acceptability, which, in turn, leads to more prolific use of a harmful treatment. For example, some estimate that antiarrhythmic drugs (adopted on the basis of reasoning from some but not all relevant mechanisms) killed more people every year than were killed in action during the whole of the Vietnam War (Howick 2011, p. 934).

As mentioned previously, I am worried that Howick is arguing here against a straw man. No one has said that mechanisms alone could solve the problem of extrapolation. They can help solving it, in conjunction with population studies. Moreover, they are crucial for (making steps towards) solving it, insofar as the problem of extrapolation appears due to erratic behavior of certain mechanisms, to several mechanisms being involved in producing the same effect, etc., as otherwise Howick very rightly notes. Antiarrhythmic drugs killed more people than the Vietnam War because the microstructural assessments were not backed up by population level studies. But one can also cite numerous other examples of failures of extrapolation due to reasoning based on population studies alone, as for instance the use of beta-blockers in hypertension, or the various treatments and diets involved in the long history of atherosclerosis.⁵¹ Such failures of extrapolation due to reliance on population studies alone were followed by microstructural research which looked into the difference-making within mechanisms, which in turn triggered more specific population studies.⁵²

One last argument advanced by Howick against RWT is that, in exceptional cases, mechanisms could be sufficient for establishing causation.

"Accordingly, to be accepted, mechanistic reasoning must satisfy the following desiderata: 1. Knowledge of mechanisms upon which the mechanistic reasoning is based is not incomplete; that is, the mechanisms linking the intervention with the outcome have been identified and their behavior under intervention established. A not

⁵⁰ Clarke et al. (2014a, pp. 346-348).

⁵¹ For the case of beta blockers see Messerli et al. 1998; Khan and McAlister 2006. For the case of atherosclerosis, see the series of articles by David Steinberg (2004), Steinberg (2005a, b).

⁵² Howick presents convincingly Cartwright-like arguments against mechanisms along being able to solve the problem of extrapolation due to their complexity, paradoxical behavior under intervention, etc. But Cartwright has an equally convincing battery of arguments as to the limitations of RCTs in predicting effectiveness; see for instance Cartwright and Munro (2010). To belabor the point, I could add here that Cartwright's arguments both against RCTs and mechanisms tend to look similar and have a common root, which is not surprising, because, as I said, extrapolation is a problem of causal relations as such, not of mechanisms specifically or of population studies specifically.

incomplete understanding of the mechanistic chain linking the intervention with the clinically relevant outcome involves correct identification of relevant mechanisms. But identifying one inferential pathway is not enough. Mechanistic reasoning must factor in the complex and stochastic nature of most biochemical mechanisms, so the other desiderata is: 2. The probabilistic and complex nature of the mechanisms is explicitly considered when inferring from mechanisms to any claims that a particular intervention has a patient-relevant benefit. When the two desiderata have been met, the mechanistic reasoning in question can be judged to be of sufficiently high quality to support a hypothesis about an intervention's effects. The following real example suggests that mechanistic reasoning might sometimes suffice to support causal hypotheses without comparative clinical studies. (Somewhat ironically, this example of "high-quality" mechanistic reasoning also counts against Russo and Williamson's claim that both types of evidence are required.)" (p. 938).

It is I believe correct that in exceptional cases, mechanisms could be sufficient for establishing causation. At least one could say is that cases like that of the radiotherapy treatment for obstruction due to large nodular goiters (the example cited by Howick) or the famous case of extracorporeal oxygenation for infant pulmonary hypertension (for which the introduction of a clinical trial has raised great ethical questions due to its apparent effectiveness)⁵³ call for a serious consideration of this possibility. Does this possibility contradict RWT? It does not, because if ontically mechanisms are made up of both production and difference-making, then in such extremely rare cases, evidence of mechanism will amount to just that—evidence of production and evidence of difference-making, in line with the letter and the spirit of RWT. The question that should be asked in return is rather—if mechanisms are *not* defined in terms of difference making, how would such cases be ever *justified*?

Notice that, in describing the conditions under which evidence of mechanisms would be sufficient, Howick appears precisely to descriptions of difference-making "Knowledge of mechanisms upon which the mechanistic reasoning is based is not incomplete; that is, the mechanisms linking the intervention with the outcome have been identified and their behavior under intervention established" (p. 938). And this is my point. Only if difference-making is linked to mechanisms that is it conceptually possible for mechanisms to be sufficient for diagnosis and prediction, without RCTs or any other population studies—in those (extremely rare) cases in which we know that grasping difference-making is not made difficult by the problems of masking and complexity.

There is no irony in the fact that the examples of "high-quality" mechanistic reasoning count in favor of Russo and Williamson's claim that both types of evidence are required; such examples speak against ontic causal pluralism, and imply that difference-making should be associated to mechanisms.

5 Conclusion

I have argued that difference making should take central stage when it comes to general role played by mechanisms in establishing causal claims in medicine. The difference-

⁵³ See Worrall 2008 for discussion.

making of mechanisms provides additional compelling reasons to accept the Russo-Williamson thesis, and include mechanisms in the protocols for Evidence-Based Medicine (EBM), as the EBM+ research group has been advocating. Via difference-making, the criterion of robustness in the protocol for evidence of mechanisms provided by Clarke et al. becomes specified. More work needs to be done though with respect to the criteria of evidence of mechanisms overall. In particular, one would welcome an epistemological theory that would underpin in a unitary way the use of these criteria to classify the quality of evidence. In light of the central importance of difference-making, such an epistemological theory might well be the Inference to the Best Explanation, for which, among others, Michael Strevens' work might turn out very important. But that is the subject of a different paper.

Acknowledgments This paper was written as part of the project “Grading Evidence of Mechanisms in physics and biology”, funded by the Leverhulme Trust.

<http://blogs.kent.ac.uk/jonw/projects/grading-evidence-of-mechanisms-in-physics-and-biology/>

I am grateful for useful comments to Kristoffer Ahlstrom-Vij, Veli-Pekka Parkkinen, Jon Williamson and an anonymous referee for this journal.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Bechtel, W., & Abrahamsen, A. (2005). Explanation: a mechanist alternative. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 36(2), 421–441.
- Berkovitz, J. (2007). Action at a distance in quantum mechanics. *Stanford Encyclopedia of Philosophy*, accessed October 2014. <http://plato.stanford.edu/entries/qm-action-distance/>.
- Bird, A. (2007). *Nature's metaphysics: laws and properties*. Oxford: Oxford University Press.
- Cartwright, N., & Munro, E. (2010). The limitations of rcts in predicting effectiveness. *Journal of Experimental Child Psychology*, 16(2), 260–266.
- Casini, L., McKay Illari, P., Russo, F., & Williamson, J. (2011). Recursive Bayesian nets for prediction, explanation and control in cancer science. *Theoria*, 26(1), 495–4548. <https://russofederica.files.wordpress.com/2010/02/rbnpec.pdf>.
- Cassini, L. (2015). Can interventions rescue Glennan's mechanistic account of causality? *British Journal for the Philosophy of Science*. Advanced access, published online March 2015.
- Clarke, B., Gillies, D., Illari, P., Russo, F., & Williamson, J. (2014a). Mechanisms and the evidence hierarchy. *Topoi*, 33(2), 339–360.
- Clarke, B., Leuridan, B., & Williamson, J. (2014b). Modeling mechanisms with causal cycles. *Synthese*, 191, 1651–1681.
- Darby, G., & Williamson, J. (2011). Imaging technology and the philosophy of causality. *Philosophy & Technology*, 24(2), 115–136.
- de Vreese, L. (2006). Causal pluralism and scientific knowledge: an underexposed problem. *Philosophica*, 77, 125–150.
- Dragulinescu, S. (2012). On ‘Stabilising’ medical mechanisms, truth-makers and epistemic causality: a critique to Williamson and Russo's approach. *Synthese* 187(2), 785–800.
- Glennan, S. (1996). Mechanisms and The Nature of Causation. *Erkenntnis*, 44, 49–71.
- Glennan, S. (2002). Rethinking mechanistic explanation. *Philosophy of Science*, 69(S3), 342–353.
- Godfrey-Smith, P. (2008). Causal pluralism. In H. Beebe, C. Hitchcock, & P. Menzies (Eds.), *The Oxford handbook of causation*. Oxford: Oxford University Press. <http://petergodfreysmith.com/CausalPluralismPGS-07-Final.pdf>.

- Hall, N. (2004). Two concepts of causation. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 181–204). Massachusetts: The M. I. T. Press. http://www.power-structuralism.ox.ac.uk/_data/assets/pdf_file/0018/27504/Hall_TCOC.pdf.
- Hall, N. (2012). Comments on Strevens' depth. *Philosophy and Phenomenological Research*, 84(2), 474–482.
- Howick, J. (2011). Exposing the vanities—and a qualified defense—of mechanistic reasoning in health care decision making. *Philosophy of Science*, 78(5), 926–940.
- Howick, J., Glasziou, P., & Aronson, K. (2013). Problems with using mechanisms to solve the problem of extrapolation. *Theoretical Medicine and Bioethics*, 34(4), 275–291.
- Illari, P. (2011). Mechanistic evidence: disambiguating the Russo–Williamson thesis. *International Studies in the Philosophy of Science*, 25(2), 139–157.
- Illari, P., & Williamson, J. (2012). What is a mechanism? Thinking about mechanisms across the sciences. *European Journal for Philosophy of Science*, 2(1), 119–135.
- Khan, N., & McAlister, F. A. (2006). Re-examining the efficacy of β -blockers for the treatment of hypertension: a meta-analysis. *CMAJ*, 174, 1737–1742.
- Lewis, D. (1973). Causation. *Journal of Philosophy*, 70, 556–567.
- Lewis, D. (1986). *Postscript to 'Causation'*. *Philosophical Papers* (Vol. 2, pp. 172–213). New York: Oxford University Press.
- Longworth, F. (2006). Causation, pluralism and responsibility. *Philosophica*, 77, 45–68.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67(1), 1–25.
- Messerli, F., Grossman, E., & Goldbourt, U. (1998). Are beta-blockers efficacious as first-line therapy for hypertension in the elderly? A systematic review. *JAMA*, 279, 1903–1907.
- Pippin, R. B. (1979). Negation and not-being in Wittgenstein's tractatus and Plato's sophist. *Kant-Studien*, 70(1–4), 179–196. <http://philosophy.uchicago.edu/faculty/files/pippin/negation.pdf>.
- Psillos, S. (2004). A glimpse of the secret connexion: harmonizing mechanisms with counterfactuals. *Perspectives on Science*, 12(3), 288–319.
- Russo, F., & Williamson, J. (2007). Interpreting causality in the health sciences. *International Studies in the Philosophy of Science*, 21(2), 157–170.
- Russo, F., & Williamson, J. (2011). Epistemic causality and evidence-based medicine. *History and Philosophy of the Life Sciences*, 33(4), 563–582.
- Ryle, G. (1966). *Plato's progress*. Cambridge: Cambridge University Press.
- Skipper, R. A., & Millstein, R. L. (2005). Thinking about evolutionary mechanisms: natural selection. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36, 327–347.
- Starko, K. M. (2009). Salicylates and pandemic influenza mortality, 1918–1919 pharmacology, pathology, and historic evidence. *clinical Infectious Diseases*, 49(9), 1405.
- Steinberg, D. (2004). An interpretive history of the cholesterol controversy: part I. *Journal of Lipid Research*, 45, 1583–1593.
- Steinberg, D. (2005a). An interpretive history of the cholesterol controversy, part II: the early evidence linking hypercholesterolemia to coronary disease in humans. *Journal of Lipid Research*, 46, 179–190.
- Steinberg, D. (2005b). An interpretive history of the cholesterol controversy, part III: mechanistically defining the role of hyperlipidemia. *Journal of Lipid Research*, 46, 2037–2051.
- Strevens, M. (2012). Replies to Weatherston, Hall, and Lange. *Philosophy and Phenomenological Research*, 84(2), 492–505.
- Strevens, M. (2013). Causality reunified. *Erkenntnis*, 78(2), 299–320. <http://www.strevens.org/research/expln/Unicause.pdf>.
- Swindler, J. K. (1980). Parmenides' paradox: negative reference and negative existentials. *Review of Metaphysics*, 33(4), 727–744.
- Vandenbroucke, J. (1996). Evidence-based medicine and “Medecine d'Observation”. *Journal of Clinical Epidemiology*, 49(12), 1335–1338.
- Williamson, J. (2011). Mechanistic theories of causality. *Philosophy Compass*, 6(6), 421–432.
- Worrall, J. (2007). Why there's no cause to randomize. *British Journal for the Philosophy of Science*, 58, 451–488.
- Worrall, J. (2008). Evidence and ethics and medicine. *Perspectives in Biology and Medicine*, 51(418–431), 2008.