

Information and Preferences Over Redistributive Policy: A Field Experiment

Lucy Barnes*

School of Politics & International Relations, University of Kent
l.c.barnes@kent.ac.uk | <https://kent.academia.edu/LucyBarnes/>

Avi Feller

Department of Statistics, Harvard University
avifeller@fas.harvard.edu
<http://statistics.fas.harvard.edu/people/avi-isaac-feller>

Jake Haselswerdt

Department of Political Science, University of Missouri
haselswerdtj@missouri.edu

Ethan Porter

Department of Political Science, University of Chicago
evporter@uchicago.edu | <http://www.ethanporter.com/>

October 27, 2014[†]

Abstract

Does better information about policy change voter attitudes? We ask and answer this question with specific reference to redistributive policy, a core concern in the comparative literature about advanced industrial democracies. Most scholars of political economy have highlighted the potential for asymmetries in information to skew redistributive policy towards the interests of the rich. We evaluate the causal impact of providing better information through a field experiment by randomising the timing of receipt of a new policy providing budget information to taxpayers.

*Corresponding author

[†]'Pre-analysis plan' prepared for submission to CPS Transparency Special Issue. This project is funded by Omidyar Network. Thanks to Paul Phillips and Rohan Grove at HMRC, Michael Hallsworth from the Behavioural Insights Team, Ian Neale at YouGov and Laura Bacon at Omidyar Network. Don Green, Brendan Nyhan and Jason Reifler, John Sides and Richard Thaler provided invaluable advice at various stages of the experimental and survey design. The usual disclaimer applies.

This pre-analysis plan submitted details our motivation and expectations, both grounded in theoretical approaches from comparative political economy; and a detailed plan for the implementation and analyses of the experiment. It also highlights some results from a pilot (online survey) experiment which indicate that the findings may be somewhat surprising. Given the nature of this project– the collaboration HMRC, with a major external partner, the design also provides unique opportunity to document how the research process is both enhanced and constrained by practical requirements, an angle not often emphasized in scholarly treatments of research design, but critical for transparency of the scientific process.

Information and Preferences Over Redistributive Policy: A Field Experiment

1 Introduction

The importance of informed voters is a long-standing and disputed issue which draws attention from all subfields in political science. From a normative fear that voters may be unable (or unwilling, given their finite time) to shoulder the demands made by democratic theory (**Landemore2013**), to classic empirical treatments asking how informed voters really are (**Converse; Caplan**), the idea that information matters is nearly unchallenged. This proposal outlines an experimental empirical approach to answer the question: does better information about policy change voter attitudes? We ask and answer this question with specific reference to redistributive policy, a core concern in the comparative literature about advanced industrial democracies.

A lack of voter information, particularly if skewed with income, may contribute to what has become a common concern in both media and scholarly treatments of economic inequality. the lack of voice and representation of low income voters' true interests may skew policy in favour of the rich. While the 'matter with Kansas' may not be that poorer voters are distracted by religion (**GelmanEtAl2008; Bartels2008a**), an extensive scholarly canon in both American and comparative politics suggests that lack of voter information penalises redistributive egalitarian policies (**Bartels2005; Bartels2008a; IversenSoskiceWSWP6; Arnold2012; FowlerMargolis2014; FrankoTolbertWitko2013; Campbell2010**).

But these claims are made typically without reference to a broader lack of consensus on the claim that information matters overall, let alone with a specific direction in terms of policy bias (**Gilens2001; Bartels1996a; Druckman2005**). Further, if voters do lack information, it is unclear how easy it is to remedy the situation, as political learning appears to be very difficult (**NyhanReifler2010**). In this research project we bring these two literatures together— along with more serious attention to the informational content of the status quo, as highlighted in

recent work by **Mettler2011** and **Gingrich2014**

The wide range of interesting and plausible macro- political economy accounts of the links between socioeconomic status, information, preferences and vote choice hangs on the individual-level mechanism that greater information allows voters to better vote their economic interests; the study we propose will ascertain whether or not even the most basic implications of this assumption are true: does information alter preferences? Does it do so so as to bring them more in line with (‘enlightened’) self-interest?

The means for plausibly answering these questions are an innovative field experiment in collaboration with the British tax authority, Her Majesty’s Revenue and Customs (HMRC). HMRC are providing every income taxpayer in the country a tax statement summarising their annual tax contribution and its distribution across various categories of government spending. By randomising the receipt of this information for individuals matched within YouGov’s panel of survey respondents, we can directly and experimentally determine whether this information on the impact of government policy on individuals’ pocketbooks affects their views. The delivery of the information by the government’s own tax bureau also means that its credibility, as well as the likelihood of compliance with the treatment, is greatly improved compared to other assessments of the effect of information.

This study is also of more than academic interest. It was motivated by the apparent contradiction in the expected effect of information on attitudes implied by two different sets of policy makers. In the United States, Democratic and liberal politicians have pushed for a federal tax receipt to provide information about all the benefits that government provides (**PorterKendall2011**). The implicit reasoning is that once voters realise the value provided by government programs, they will be more supportive of them. In contrast, in the United Kingdom, Conservative politicians have been the driving force behind the implementation of the tax receipts policy (**MontgomerieGummerProposal**), and their expected impact illustrated by the endorsement of the policy by anti-tax campaign group the Taxpayers’ Alliance (**TaxPayersAlliance**). That is, both supporters and opponents of greater redistribution seem to think that once people have

all the facts, attitudes will shift in their favour.¹

In the next section we outline the theoretical framework according to which information about the existing policy situation should be expected to change attitudes towards redistribution. This involves a more careful elaboration of the role of the status quo in considering the determination of preferences than has hitherto been typical. However, the major contribution of this project and the characteristics which make it an ideal example with which to illustrate various issues surrounding research transparency lie in its empirical research design.

As intimated above, taxpayer receipts detailing taxpayers' individual contribution to the government budget, as well as a breakdown of how that money is allocated across spending areas, have been rolled out (online) in the United States. In the United Kingdom, a paper version of this information will be sent to 16 million taxpayers between October and December of this year (2014). In collaboration with HMRC, the government bureau implementing this policy, we will randomise receipt of the information within a representative subset of this population. This experimental leverage, paired with the solicitation of attitudes about redistributive government policy, will allow us to ascertain the causal impact of government-provided information about taxing and spending on attitudes towards redistribution.

This design represents a major advance in its causal credibility relative to existing treatments of redistribution preferences in the public at large, which tend to rely on observational data (**Margalit2013; LupuPontusson2011**). Yet, as a field experiment, it also improves upon studies on the impact of information, which tend to rely on experiments either in the lab, or within a survey, for their leverage (**Mettler2011; Gilens2001**). This situates the study in a uniquely advantageous point in the trade-off between internal and external validity. Discerning the impact of information in the 'real world' is more relevant for policy-makers seeking to change information levels in the population.

Moreover, with a clear theoretical motivation, and deductively derived hypotheses, we are in a position to directly implement a transparent research design. Specifying our hypotheses in ad-

¹Technically, they could both be right, if the US and UK status quo is located in a different place relative to voters' ideal positions. But if policy is skewed in different directions in the two countries, why is this so? At least prima facie, there is a puzzle here.

vance, on the basis of the theoretical framework and refined by a pilot study already conducted, we are able to demonstrate insulation from ‘ p -value hacking’ under multiple comparisons. In terms of transparency, the additional complications of implementing the study collaboration with an external partner will provide a unique angle on adapting to ‘exogenous shocks’ to the empirical research design.

In the next section we elaborate explicit hypotheses, and outline their theoretical basis. ?? outlines the design of the study and the analyses we plan in order to test our hypotheses. ?? outlines our expectations about feasibility, and outlines contingency plans in case the first-best implementation is impossible. We then describe our pilot study and present its results and implications for the power of the main study in ???. In lieu of established conclusions, ?? highlights the contributions that will be made by this study, particularly in the context of its potential publication in the CPS special issue on Transparency.

2 Theory: Information and Redistributive Policy Preferences

Our expectations that information about the policy status quo should affect preferences seems unobjectionable on its face. That is, at least at the extreme, if voters have no information whatsoever as to what government does or is able to do, then they will be unlikely to have any ability to form preferences over policy. But in fact, the role of information about the existing policy landscape has received only indirect attention in the theoretical literature on preferences towards redistribution. Thus in this section we lay out explicitly the theoretical basis for this importance of information about the status quo. There are four mechanisms by which understanding existing policy may affect preferences. The first fits easily into a stylized account of preference formation which should be quite familiar, once we acknowledge the potential importance of policy *feedback effects*. Information about the status quo may also allow voters to more directly map their *material* interests in to the policy space. Third, existing policy information may affect preferences because of important ways in which a ‘full’ spatial model of politics is overly demanding: since most political reforms and proposals take the form of ‘increases’ or ‘de-

creases', the space in which political preferences are relevant is strongly anchored by the status quo. Thus policy directions represent an important *heuristic* mechanism whereby the location of the starting point becomes extremely important. We discuss each of these mechanisms in turn before turning to the direct evidence that exists in the redistributive context. Finally, there may be a mechanism operating at the level of voters' *expressed preferences*, particularly in survey questions, which also often focus on 'increase' or 'decrease' formulations.

First, a large literature on policy feedback effects notes the impact that existing policies can have on individual attitudes. While much of this literature focuses on the effects of existing policy in motivating citizens participation in politics, and mobilising particular constituencies, the more relevant effects of policy here are on the ideas that citizens form about their personal identities (**Pierson1993**), and in shaping perceptions about the appropriate role of government. In this context, changes to policy may compel changes in attitude, "but this dynamism depends on information reaching citizens" (**Campbell2012a**). We would argue that this necessity for information is not only at issue when policy changes, but that even when policy is stable, the mechanisms by which citizens' ideas are shaped by policy requires that they know about the policy in the first place.

A second status quo operates through a slightly different channel, by changing material interests of particular groups. The seminal example here is the creation of old-age Social Security programs creating a committed constituency in favour of government action by virtue of the material benefits that it generates (**Campbell2003a**). On the other hand, such feedback can be negative: interaction with the welfare system in the United States has been shown to generate more mistrust of government than support (**Soss**). Similarly, opinion may respond to existing policy in a 'thermostatic' manner (**Wlezien1995**). In cases of benefit claimants, the relevant 'information' is their experience in dealing with program administrators. However, we argue that informational effects of the status quo should also operate at less immersive levels. That is, information that indicates to citizens whether or not they benefit from particular government programs should change the way in which they match their material position to policy preferences (assuming there is something 'learned' in the process of informing citizens). This superior

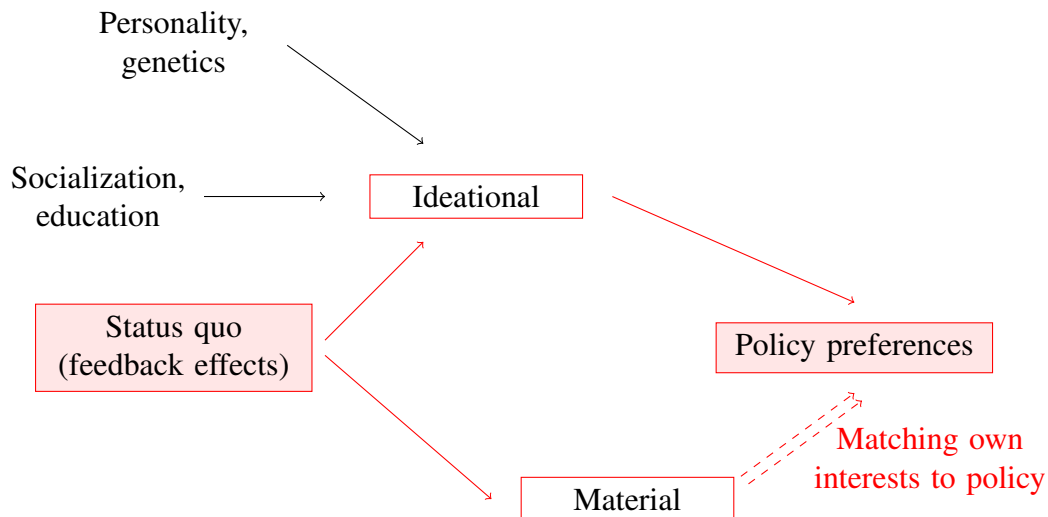


Figure 1: A stylized model of preference formation. The diagram is not meant to be ‘complete’: presumably personality and socialisation may also affect material position and thus material interests, and so on. However, it captures the main ideas, to enable us to illustrate the location of status quo effects via the ‘feedback effect’ and ‘material’ mechanisms. In particular, the traditional ‘policy feedback effects’ of Pierson and Campbell would fit in the diagram in driving the arrow between ‘ideational’ and ‘policy preferences’, as well as directly along the ‘material’ line in the figure. Information allowing voters to better situate their material interests relative to the status quo (the ‘matching’ element in the diagram) is the logic associated with Bartels, or Iversen and Soskice.

matching of interests to policies under higher information has been identified in survey experiments at the level of specific policies (Mettler2011); and where more visible welfare states provide more information about the delivery of benefits at the macro-level (Gingrich2014).

?? shows a stylized model of preference formation that summarizes these first two mechanisms.

The other two mechanisms by which the status quo ought to affect policy preferences are closely linked to one another, but subtly. Under classic spatial theories of policy competition, voters need quite demanding estimates of positions– of parties and their own interests– in the policy space. If we believe that voters are (at least somewhat) limited in their ability to go through the complex calculations associated with the spatial model, then using the direction from the status quo as a *heuristic* shortcut to deciding their own policy positions to the necessary level of precision (RabinowitzMacDonald1989). Knowing the position of the status quo thus becomes extremely important. Furthermore, even if voters do evaluate policy positions more

comprehensively in the policy space, the location of the status quo provides a useful benchmark or anchor.

This is related to the fourth and final mechanism by which we may measure changes in attitudes in consequence of providing extra information to voters. Even without the anchoring impact of the status quo, many survey questions used to gauge policy preferences take the form akin to “Should government spending (on policy area X) increase, stay about the same, or decrease?” In these contexts it is clear that the accurate perception– or at least a *consistent* perception– of the status quo across individuals is required to interpret responses as meaningful with regard to abstract positioning in the policy space (**BeramendiRehm**). Although we use questions of both this and the non-anchored type (see below), the distinction between the two kinds of preference is quite subtle.

In previous studies of public attitudes, the status quo has been investigated usually with reference to public resistance to policy change (**Hansen1998**). However, our focus on the role that the status quo may play turns more closely on the ability of individuals to find diagnostic information in their current environment. **KuklinskiEtAl2001** find that the combination of general information about the need for trade-offs, combined with material motivation for good decision making, improves policy choices in the sense of making respondents’ preferences more consistent in the context of tradeoffs between existing goals. In our study, therefore, the question is whether information about the status quo can provide such information and motivation (**McCallChin2013; Trump2013**). In principle, since the tax amounts in particular are specific to the individual, this should provide material motivation. The spending information is illustrative of the necessary tradeoffs required to allow tax reductions.

2.1 Budget Knowledge and Redistribution Preferences

How would we expect these general theoretical ideas to play out with concrete application to redistributive policies? In particular, how does information on the budgetary status quo affect these preferences?

The general literature on information and preference formation has not typically attended to the implications of policy information on individuals with different types of material interest. The main exceptions to this have come from the political economy literature highlighted in the introduction. Of particular interest in the context of our experiment, and the theoretical model articulated above, is the role of better information in allowing voters to match their material interests to policy choices.

In this vein, **IversenSoskiceWSWP6** derive a formal argument whereby the provision of information polarises public opinion. Under conditions of uncertainty, they argue,

“those with low information have a strong tendency to vote centrist (whatever their true interest). We show analytically that this follows from a generalization of the standard spatial voting model to cover limited information”.

Intuitively, under uncertainty a centrist position is closer than an extreme one, in expectation, to an individual's true preference. With the provision of better information, then, voters ‘learn’ that in fact their interests are better represented by policy positions that lie more to the extremes. The kind of information that should have this effect includes better information about what taxes actually pay for, when the preference being elicited concerns the level and distribution of taxation. Under ignorance about how much he pays in taxes and receives from benefits, an individual may want to maintain the status quo for many reasons (**Hansen1998**). Depending on the net gains or losses revealed by the details of where money is spent, with the provision of information this individual may shift his position substantially. Essentially, the provision of information allows voters to shift from the risk-minimising position to their ‘true’ interest; and given that the former is the same (centrist) point for all, while the latter includes a wide heterogeneity, better information should lead to polarisation in the sense of a greater number of individuals being able to identify extreme positions as to their benefit, and doing so.

Secondly, it follows from this argument as well as others, that better information on the budget should allow voters to align their expressed preferences with their material interests (**IversenSoskiceWSWP6; Bartels2008a**). For example, a low-income pensioner who thinks

that their tax contribution goes largely to finance national defence may express preferences more opposed to tax increases than her material position should suggest, given the reality of the pension share in the total budget. Being apprised of this reality, she should be able to express a preference over taxation that better reflects these material interests. Equally, voters may be mistaken about how much tax they pay; by revealing this amount at the individual level, matching the financial contribution to the perceived benefits on the spending side will better enable preferences to reflect material reality.

To the extent that these kinds of mistakes and lack of information are unevenly distributed across material positions, the aggregation of these heterogeneous effects may have an overall bias. Given that low income voters are typically less well informed than higher income voters, this may have a pro-redistribution impact on aggregate.

Finally, better information may also have an effect on preferences regardless of personal position. Without accurate knowledge of status quo spending, voters may be unable to know which policies help them materially, and which do not. For example, a heavy consumer of healthcare and pensions (see ??) might not realise the implications of tax cuts on their material position without knowing the extent to which these items dominate the government budget. This kind of misperception— overlooking the degree to which they benefit from government spending under the status quo— has been found to drive a lack of support for government intervention and public policy in the United States, with better policy information increasing support for policies that benefit the middle class, at the expense of those which disproportionately benefit the rich **Mettler2011**

This is also where the ideational effects of policy are most likely to be relevant. If the existence of public welfare programmes instills the belief that support for those who need them is a proper role for government to play, then correcting over- or underestimation of the effort that is actually dedicated to that goal will lead to changes in attitudes.

In the UK context, misperceptions of the extent of the government budget that goes on redistributive programs has been argued to underpin support for radical reforms of unemployment and disability benefits in the name of deficit reduction, despite the relatively small share of

spending that these programs constitute (**TaylorGooby2013**).² The first material mechanism whereby policy information may shift attitudes, then, is if the material benefits of government tax and spending policies are consistently underestimated, and better information rectifies this misperception. This would be consistent with prior effects of information on British voters, leading to greater support for higher taxes on high income voters (**Sturgis2003**). Recent experimental research from the United States also indicates a slight pro-government bias with the provision of better information (**PorterAPSA2013**).

2.2 Hypotheses

We are now in the position to articulate the hypotheses to take to the data.

Following **IversenSoskiceWSWP6** the impact of greater political information should be polarising (since uncertain voters rationally vote centrist). Note that this is contra **AlvarezBrehm2002** for whom additional information reduces variance in preferences due to the reduction in uncertainty. The difference between the two expectations is the difference between within-individual variance (which is the key conceptual element for the latter authors, and which should be reduced by reducing uncertainty) and between-individual heterogeneity, which is the more obvious focus of distributive politics, and of (**IversenSoskiceWSWP6**).

Hypothesis 1 *The treatment group will have more polarised attitudes on redistributive policy than the control.*

The status quo ante of ‘unenlightened self-interest’ implies that this polarisation should be predictable in terms of individual interests **Gingrich2014; IversenSoskiceWSWP6**. With more information, higher income individuals should espouse greater resistance to redistribution; lower income individuals should prefer more. The same sorting should apply to ideology, with rightist respondents better able to align their general preference for less government intervention

²Unfortunately, the decision not to separate the more explicitly redistributive welfare policies from pension expenditures in the HMRC figure was a decision that was taken at a political level far beyond our ability to influence. In fact, spending on unemployment insurance, social exclusion and social protection comprise 30% of this ‘welfare and pensions’ category (**OBRTables**).

with the more redistribution-specific reductions in tax levels in the presence of information about the status quo.

Hypothesis 2 *Information on the policy status quo should increase the influence of interests on attitudes towards redistribution. Specifically, (high) income should be a better predictor of (disfavourable) views on redistribution in the high information (treatment) group.*

Finally, if errors about the status quo are systematic, then the provision of information (following **Gilens2001**), the distribution of tax statements – which provide policy-specific information about government’s redistributive efforts – may change the overall cast of preferences. In light of the conflicting theories and we might not anticipate a particular directionality in favour of more or less redistribution. We do not necessarily have strong theoretical priors about the effect of information on the central tendency. If anything, the lower level of information among the natural beneficiaries of expanding government programs implies that the aggregate impact of these opposing shifts should be a shift in favour of redistribution. However, this is in conflict with **Gilens2001** finding that the impact of policy-specific information of attitudes is greater among those with higher levels of general knowledge– assuming again that these high-information individuals are relatively concentrated among the rich. Moreover, at least in the British context, support for the provision of greater information via the tax statements has come from anti-tax, anti-redistribution political actors. Thus policy-makers’ ‘hypothesis’ about the direction of this effect certainly operates in the opposite direction to that which we derive from the literature as:

Hypothesis 3 *The provision of budget information (the treatment) may change attitudes towards redistributive tax policy, shifting preferences in favour of more redistribution.*

Thus, though the simplest of our hypotheses, this is also the one in which we have the weakest theoretical support. However, in policy terms it is the most important open question, and speaks directly to the puzzling contrast between pro-tax, liberal support for tax and spending information in the American context, and anti-tax, Conservative support in the UK. Although we have limited justification for anticipating an effect in either direction, the operating assumptions of policy-makers do seem strongly premised on the idea that such information matters.

3 Research Design: A Government Policy Field Experiment

As noted in the introduction, in January of 2012 the Conservative MP for Ipswich (in the UK) made the case in parliament for a bill to be introduced legislating the provision of ‘tax statements’, (also known in the US as ‘tax receipts’) to all taxpayers. This policy is being implemented this autumn by the revenue collection agency, HMRC. Sixteen million Pay As You Earn (PAYE) taxpayers will receive a paper statement detailing their personal contribution to the Exchequer (via income tax and social security contribution, but not indirect nor other taxes).³ The statements will also contain a breakdown of how that money was used by the government. An example version of a tax statement, of the type to be used can be seen in ??.⁴ The government-issued information will differ from this statement in its provenance (from the tax collection arm of the government), increasing the credibility of its information. It will also be personalised to the level of each individual’s total taxable income, and tax contribution.

The key intuition behind our design is to leverage the staggered roll-out of this information. By randomly assigning a group of taxpayers to receive the statements (at the time we will elicit their preferences) and a group to not (yet) receive them, we are able to manipulate exposure to the information the statements contain. In other words, subjects in control will not yet have received the treatment at the time they are surveyed; subjects in the treatment group will have received the statements at this time.

HMRC are implementing the roll-out of the tax statements across the population over almost two months in the autumn of 2014. A pilot round of 240, 000 statements will be delivered on October 22nd, to help HMRC gauge the resources they may need to cope with any response from recipients.⁵ Once this has been assessed, the full roll out will begin in November, and just over 2.5 million statements per week will be distributed over the subsequent six week timeframe.

³PAYE taxpayers comprise the bulk of employees in the UK. They typically do not have to file a tax return, but rather taxes are withheld and any necessary adjustments based on final (employment) income are made by HMRC at the end of the tax year.

⁴This is an example as used from the pilot study.

⁵Of particular concern to HMRC is that the summaries will be construed as demands for payment, and elicit large volumes of enquiries for their customer relations centres.

The information on this page shows you how your Tax and National Insurance Contributions were spent in 2013-4. This does not include indirect taxes such as VAT and other duties. All information comes from HRMC.

Use	Your contribution
Welfare and Pensions	£3,400.35
Health	£1,673.04
Education	£1,175.65
National Debt Interest	£642.09
Business and Industry	£226.09
Defence	£488.35
Criminal Justice	£425.04
UK Contribution to EU Budget	£54.26
Overseas Aid	£81.39
Environment	£144.70
Culture	£162.78
Housing and Utilities	£135.65
Government Administration	£189.91
Transport	£253.22

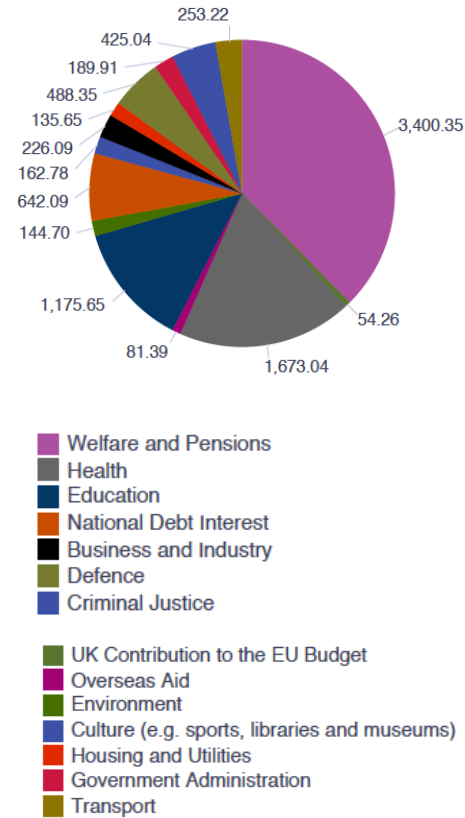


Figure 2: Example tax statement for an individual with an income of £37500 per year.

By randomly selecting a group of individuals to be included in the pilot wave, we will ‘piggy-back’ on the pilot. The elicitation of attitudes is being implemented by YouGov, so the treatment (and control) groups will be samples from their panel of respondents. By providing HMRC with a list of YouGov respondents who must be included in the pilot wave, we are able to know about the receipt of information without compromising taxpayer confidentiality.

In the context of information treatments that can be delivered in the field, this is a strong intervention. The tax statement arrives in the form of a letter from HMRC, with two important implications for compliance. First, the treatment group need undertake no active effort to get the information. Second, letters from the national revenue collection agency are difficult to ignore: often they contain important information about how much tax an individual should pay; whether any errors have been made; or if auditing is necessary. Thus the legal authority of the tax collectors stands behind the delivery of the information to the treatment group. In addition, the official nature of the information (rather than information provided directly by researchers, for example) makes it more credible and more important as an account of government activity.

3.1 Randomization and Data Collection

Our implementation strategy as of September 5th is outlined in ???. The only two major obstacles to a clear implementation of our preferred design hinge on the matching capacity of the HMRC information technology team, who are implementing the print run for the tax statements, and ensuring appropriately informed consent from the respondents, given the information flows involved. These points are highlighted in the diagram and will be resolved in the week of September 15th.

The project will then proceed as follows (as per the chart):

1. we select a random subset of the YouGov panel, of about 13000 people. 8000 are randomly assigned to the treatment condition, and their consent to the experiment is elicited.
2. YouGov release to HMRC the names and identifying details of these respondents. Name,

first line of address, and six-digit postcode in the UK allows for a precise identification of individuals, thus respondents must agree to have these details released by YouGov to HMRC.

3. Based on these identifying characteristics, HMRC's information systems team ('Aspire', as in the flow-chart) will match the records of these respondents. Based on prior experience with matching to name, address, postcode characteristics, HMRC anticipate around a 60 to 75 percent match rate, hence our over-recruitment for the target treatment group size of 5000.
4. HMRC will indicate to YouGov which respondents they successfully matched, and include them in the pilot wave of the tax statement roll-out on October 22.
5. Between October 22 and November 3 (when the full roll-out begins), YouGov will survey both our treatment and control groups.

We can straightforwardly check the balance of the treatment and control groups on the demographic characteristics they have already contributed to the YouGov panel. The slightly more problematic element is in ensuring compliance in the control group. Although they will not have been deliberately assigned to it, some of the control group may in fact receive their tax statements in the pilot wave. To match both groups' records between HMRC and YouGov is not something that HMRC believe is feasible, given the size of the overall policy for them and their IT team. However, since the pilot group is a relatively small number of taxpayers, the probability of any individual being included in the first wave is only 1.5%. We are also going to ask whether respondents recall receiving the statement, at the end of our questionnaire eliciting attitudes, and thus will be able to identify at least some of the 'always-takers' in the control group. We should also know how many always takers receive a statement but do not report recalling it from the comparable share in the treatment group. Thus, while this is problematic, it is a small issue and one whose impact on our results can be estimated.

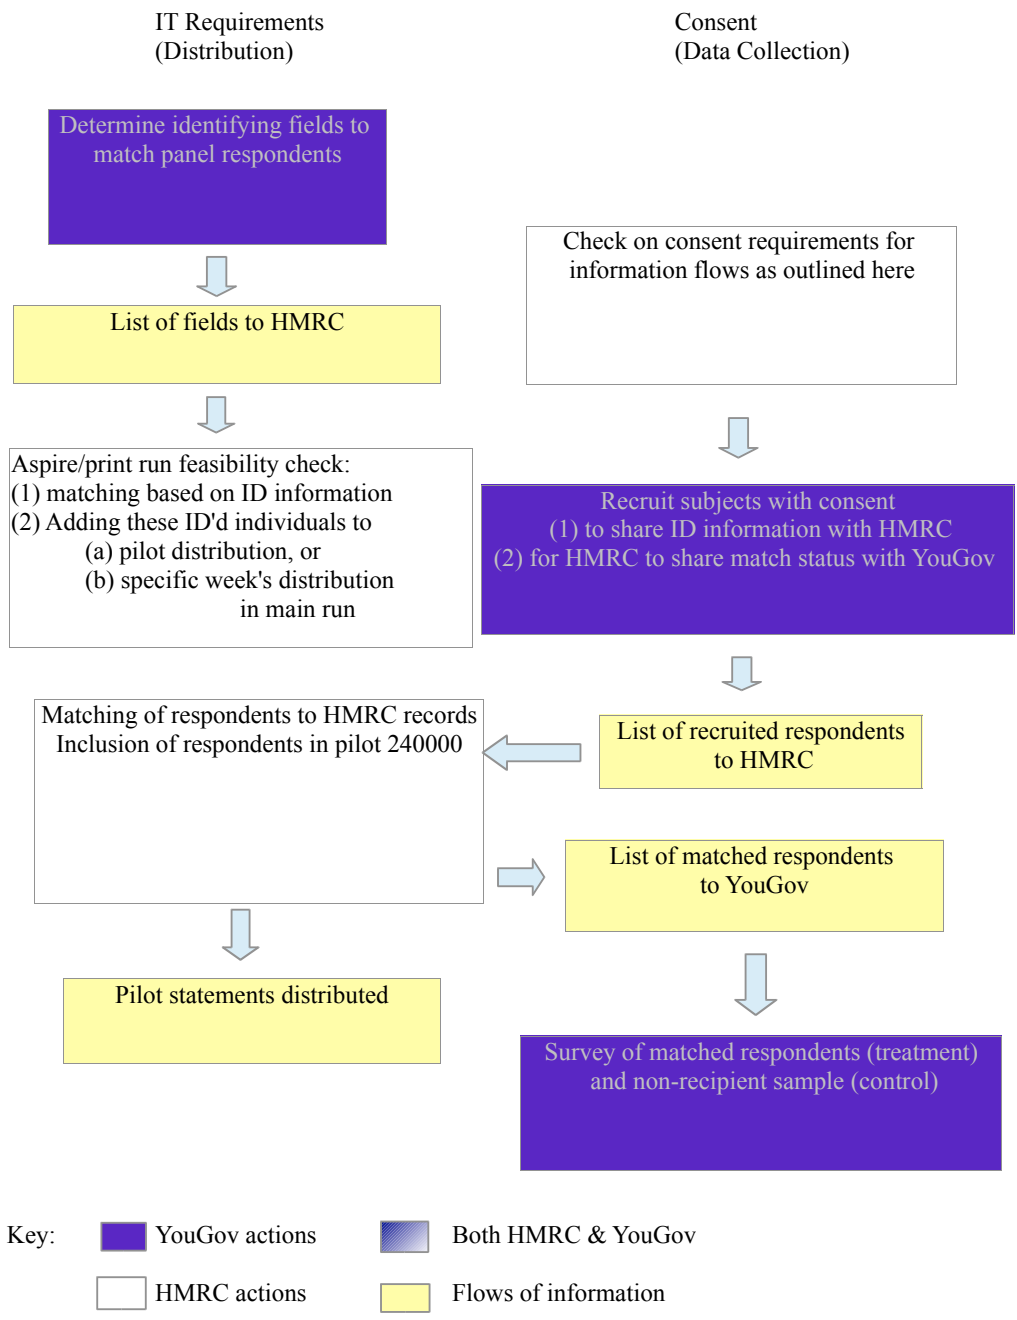


Figure 3: September 5th Implementation Chart.

3.2 Ethics, Confidentiality and Experimental Registration

This project has been approved by the Internal Review Board at the University of Chicago for the ethics implications. We are also bound to high standards regarding confidentiality as we must conform to HMRC’s best practice, and YouGov’s standards for confidentiality.

Once the implementation has been finalised, we also plan to register the study with EGAP.

3.3 Operationalising Preferences

To measure preferences over redistribution we must attend to both attitudes about the overall level of government activity, and views on the way that it is structured (**Barnes2014**). We focus on attitudes towards taxation both because the statements are most directly framed around the impact of revenue collection on the taxpayers, and because, for questions about the structure of government redistributive efforts it is much more transparent how attitudes relate to redistribution, compared to spending which are more often targeted based on ‘horizontal’ as well as vertical characteristics.

We will thus focus on three questions directly tied to the redistributive impact of taxation. First, we will ask respondents:

Generally, how would you describe the amount of taxes you pay overall? We mean all taxes together, including income tax, National Insurance, VAT, Council Tax etc.

. Responses are solicited over a five point scale from ‘much too low’ to ‘much too high’. This operationalisation mirrors that used by the International Social Survey Program (ISSP), so we have an additional way to link to other studies to situate our descriptive findings.

Second, we will ask two questions to gauge attitudes towards progressivity between the bottom and middle of the income distribution, and between the middle and the top. Having provided an estimate of the average tax rate paid by the middle quintile (which is around 35%, taking all taxes together including indirect taxation), we ask respondents about what the top and bottom income quintiles should pay,

How much do you think a household in the top fifth of income (an average household income of £83,950 a year) should pay?

with the options:

1. Less than 10% (or, much less than the middle fifth)

2. Between 10 and 30% (or, a little less)
3. Between 30 and 40% (or, about the same)
4. Between 40 and 60% (or, a little more)
5. Over 60% (or, much more)

We ask the question about the bottom quintile analogously, also providing the average income of a household within that quintile to respondents. Note that in principle this gives us a cardinality to the scale; though in practice we believe it will be most realistic to treat these as ordered rather than numeric outcomes. We also use average incomes rather than focusing on marginal rates, since these are both more intuitive (particularly in thinking about progressivity), and more easily applied to the idea of total taxes across all tax types.

3.4 Analyses

The hypotheses that emerge from the theory involve different levels of empirical complexity to test. Perhaps unsurprisingly, since the strongest theoretical expectations are about the most precise empirical relationships, the strongest theoretical implications (hypothesis 1) is the most complicated to test, and the weakest prediction (hypothesis 3) the simplest.

However, the clear theoretical basis makes the empirical comparisons very obvious, even where their implementation is relatively complex. This allows us to specify the intended comparisons in advance, to consider the required power of the study, and the necessary corrections for multiple comparisons.

3.4.1 Hypothesis 1: Polarization

The expectation that providing information about the nature of government spending (and thus redistribution) should increase the variance in preferences over redistribution turns out to be the most complicated hypothesis in empirical terms.

Reducing Outlier Responses

We are able to apply two techniques to testing the idea that information has a polarising effect on preferences (including for the pilot study)– the first involves outlier responses, while the second uses quantile regression. In case of the impossibility of fitting appropriate variance components models (see below), we thus retain leverage over the first empirical hypothesis.

First, we follow **Ansola** **here Meredith Snowberg 2014** in using outlying responses as a measure of the variation in attitudes. Higher polarization implies an larger share of ‘outlier’ responses on the preference questions. Thus there should be fewer such outliers in the control group if, as per hypothesis 1, information increases polarization. Given that our outcomes are measured on a five-point scale, these outliers do not necessarily have direct substantive implications but are an indication of diversity of attitude. Since we can code outlying responses into two categories (outliers, and not), this yields a relatively simple empirical specification.

Second, we can use quantile regression to ascertain if the treatment has different effects on individuals located at different points in the (potential) distribution of outcomes. This follows a large econometric literature on the estimation of heterogeneous treatment effects. In our case, the interest in polarisation implies that the treatment should have a ‘bigger’ effect at higher quantiles. Note that the interpretation of quantile treatment effects not as indicating that the same individuals who were ‘high’ before are now even higher, but that the high group as a whole is shifted out, is in fact more consistent with our theoretical motivation than the other (erroneous) interpretation.

3.4.2 Hypothesis 2: Treatment \times Income

The second of our theoretical hypotheses is, essentially, that the treatment effect should vary with income. We expect that the provision of information that indicates to respondents whether they benefit or lose from spending should allow them to update their preferences over taxation to reflect that reality. Given the overall redistributive slant of spending, those with high incomes who receive the treatment should be less supportive of high tax rates in general, and on them-

selves; for those with low incomes the treatment should, conversely, have the opposite pattern. For taxes on those with low incomes the predictions are more ambiguous (since higher taxes on the poor paying for extra spending is ambiguous in distributive terms). Thus although we can analyse this outcome, for completeness, we have no strong expectations deriving from theory.

To ascertain whether these differential effects by income are observed in the data, we can both (a) consider any effects of the interaction between treatment and income in classical regression models; (b) model the treatment as a random effect, with the slope varying according to income. The classical regression models will (again) be straightforward ordinal logistic models, of the form:

$$\text{logit}(\Pr(y_i \leq j)) = \alpha_j - \tau z_i - \gamma(z_i \times \text{income}_i) - \beta x_i - \varepsilon_i \quad (1)$$

Random Effects Models

In order to consider heterogeneous treatment effects by income group, we can divide respondents into income quintile based on the information in their profile. We use quintiles to divide respondents since the questions that we are interested in also focus on the relative tax burden on different income quintiles, and because respondents are more likely to be able to accurately place themselves in income quintiles than in smaller groups. Under these analyses we will estimate models of the form:

$$\text{logit}(\Pr(y_i \leq j)) = \alpha_j - \tau z_i - \beta x_i - \gamma_{q[i]} - \psi_{q[i]} z_i - \varepsilon_i \quad (2)$$

where q indexes income quintiles, and thus $\psi_{q[i]}$ represents the treatment effect on a particular income group. We use a standard varying intercept-varying slope hierarchical model to estimate treatment effect heterogeneity across income quintile.

3.4.3 Hypothesis 3: Average Intent-to-Treat Effect on Preferences over Redistribution

The notion that providing information about the spending status quo should lead to support for higher levels of taxes is the weakest of our hypotheses, predicated on the idea that without the provision of information, voters tend to overestimate the level of taxation, or to underappreciate the usefulness of the spending programs it pays for. Equally, and most important for our questions about the structure of taxation (the burden on the rich and poor, relative to the middle), voters may overestimate the degree to which spending is focused on the poor, and so when given better information, think that the tax structure should reflect that, reducing the amount that the poor are perceived to owe, and increasing that of the rich.

This idea lends it to very straightforward empirical expectations based on the average treatment effect for all of our three outcome variables: with the provision of information in the treatment condition, respondents should assess the overall levels of taxation as lower, relative to their ideal— a positive treatment effect under our coding of the tax levels variable. This same average effect should be seen for the level of taxes on the rich. For the level of taxation borne by the first income quintile, relative to the middle, hypothesis 3 entails a negative average treatment effect.

Assuming that our randomisation strategy is effective, the analyses required to identify the effect of being assigned to the treatment are fairly straightforward. We consider the difference between the preferences of those who received the tax statement to those of individuals who received no such information. We can compare outcomes in the raw data, and with regression analyses using the covariates available in respondents' YouGov profiles to improve the efficiency of the estimates. Since our outcome variables are ordered responses, simple comparisons of means is improved upon by modelling the ordered (but not cardinal) nature of the outcomes explicitly; thus we use ordered logistic specifications to model the data.

4 Contingency Plans

Given that the implementation of our intended design is somewhat complicated, we have two backup strategies to allow us to learn something from the distribution of the tax statements. The main uncertainty remaining centres on our ability to match the individuals in the YouGov survey with HMRC's tax records, either due to limitations in the matching process, or because of confidentiality and consent problems. Thus we have two ideas that do not rely on this linkage. First, we may be able to exploit as-if random of assignment to the pilot wave; secondly, without any information from HMRC, we can use an encouragement design.

4.1 Plan B: As-if-random Distribution Characteristics

In the absence of our intervention, selection to the pilot wave of 240,000 will be based on the digits of individuals' National Insurance (NI) number (akin to the Social Security number in the United States). Individuals who meet the criteria for inclusion in the population to receive the statement will be allocated to the scheduling of the roll-out based on the last two digits of their NI number. These last two digits are as-if random.

However, the last two digits of the NI number is not a piece of information that YouGov has about their panellists. Thus to pursue this design, we would need to recruit within the YouGov sample based on individuals' willingness and ability to provide this detail. Unlike in the United States, the NI number is not a widely used nor required piece of information, so individuals who can provide their digits may be quite unrepresentative of the population as a whole—likely people who have moved jobs recently, or do so frequently. Most damagingly for our purposes, otherwise similar people who have had greater contact with government benefits and services may be more likely to recall their NI number. In addition, asking for this potentially sensitive detail may make recruitment to the study much more difficult, and thus reduce our sample size considerably.

4.2 Plan C: Encouragement Design

If HRMC changes course and resists participating in the experiment as described, we have a two-pronged back up plan in place. The first part of the backup plan is a standard encouragement design. Sometimes used by researchers interested in teasing out the causal effects of treatments that would be unethical to deprive of subjects of (McDonald et al 1992), or when the treatment in question cannot be effectively randomized (AlbertsonLawrence2009), encouragement designs proceed by randomly assigning subjects to receive (or not) an encouragement to comply with the treatment. In our case, when statements begin to roll out in November, all subjects will be randomly assigned to one of three encouragement conditions. In condition 1, the strong encouragement conditions, subjects will receive a strongly worded message that tells them to be on the look out for the statement that will soon be arriving in their mailboxes. In condition 2, the weak encouragement condition, subjects will receive a message previewing a number of major national events due to occur over the next month, including the distribution of the statements. In condition 3, the no-encouragement conditions, subjects will not receive any encouragement. (Subjects in conditions 1 and 2 will receive their encouragement at the same time, when the statements are beginning to be mailed.) Then, after the statements have all been mailed, subjects in all three conditions will be surveyed about their redistributive preferences. In this case, we will be able to observe the causal effects of being randomly encouraged to read one's statement. If there are differences between the conditions, we will be able to draw the credible inference that they are due to differences in compliance caused by the encouragement.

The second part of our backup strategy relies on particular knowledge about the distribution of the statements. To set the mailing schedule for the statements, HMRC has created an algorithm that effectively randomizes the population. Every week for the six weeks that HMRC is distributing the statements, we will conduct a survey of different subjects. At the end of each survey, all subjects will be asked if they recall receiving their receipt, and then divided into treatment and control accordingly. As suggested by SoveyGreen2011 to increase the validity of these self-reports we will ask subjects who report receiving the statements some elementary recall questions about their content. Subjects enrolled in the encouragement design will also an-

swer questions meant to test their recall. Though suboptimal, this back-up strategy will provide substantive insight into the causal effects of the statements.

5 Results from a Pilot Survey Experiment

To ascertain the feasibility of the field experiment, and to road-test the survey questions, we administered a pilot survey experiment between the 14th and 18th of August, 2014. The survey was administered by YouGov. In it, 533 respondents in the treatment condition were shown an approximation of the tax statement that they would receive from HMRC (as illustrated in the example ?? above). An equally sized control group (497 respondents) did not see any information. We then asked the same set of questions as outlined above, and as planned for the field experiment.

Obviously there are a number of critical ways in which this pilot differs from the final experiment. First, the treatment was administered online, rather than by letter. While this obviates some concern about compliance (we do not need to worry about statements lost in the mail, for example), it primarily introduces differences in the way that individuals may have responded to the information (AckermanGoldsmith2011). Second, our pilot study was conducted independently from HMRC, so the statements neither bore the imprimatur of the government agency, nor were they personalised by income beyond showing individuals the tax payments for an individual at the mid-point of the household income category reported by the respondents to YouGov. For both of these reasons the credibility and relevance of the information in the treatment was lower in the pilot than it will be in the main experiment.

Despite these limitations, the results are instructive, and discussed below.

5.1 Outcome Variance

As the analytical techniques necessary to ascertain effects on the variance are more demanding both in terms of the number of observations required, and in terms of the statistical techniques

involved, so far we can present only the first set of results– the indirect results– from the pilot study. We classified the extreme responses (‘Much too high’ and ‘Much too low’ for the tax amount question, and ‘Much less’ and ‘Much more’ in the relative tax burden questions for the top and bottom income quintiles) as outliers, and modelled these responses in logistic regressions along with controls for income and its square, gender, and age, as in the other specifications.




Figure 4: Coefficient estimates of the treatment effect on the probability of giving an outlier response to the three tax preference questions. Estimates derived from logistic regression models with controls for income, income², gender, age. Dark grey bars represent 95% confidence intervals.

The essential results, as indicated in ??, is that the pilot study provides no evidence that information has a polarising effect on attitudes. The 95% confidence intervals overlap zero effect for both outlying responses on the overall level of taxation and for the top quintile; and in the latter case the point estimate itself is extremely close to no effect. Only for taxes on the bottom quintile does the treatment appear to affect the likelihood of outlier responses, but here it is in the opposite direction from the theoretical prediction: those in the treatment condition

are less likely to espouse extreme attitudes. However, given the initial distribution of outliers in this last category, it is impossible to ascertain whether this is a true ‘variance’ effect rather than a ‘level’ effect, as 542 of the 545 extreme responses are respondents claiming that the bottom quintile should pay ‘Much less’, and only three espousing ‘Much more’.⁶ On the other hand, this pattern is consistent with the idea that better information can reduce uncertainty for voters (AlvarezBrehm2002).

5.2 Treatment × Income

Plotting our outcomes according to income and treatment status gives a first insight into the data. For both questions about the overall level of taxation, and of bottom incomes versus the middle, income has some effect but no difference between treatment and control groups is apparent (a finding repeated in more involved analyses below). However, for attitudes towards the relative tax rates at the top, we see (just) a discernible gap between the treated and control groups at low levels of income, using lowess smoothed predictions and their confidence intervals.

At low income levels we see not only that the treated group is more favourable towards higher taxes on the rich, but also that the slope of the two lines differs, such that the relative positions switch. This indicates that (at these low levels of income), in the control group increasing income increases support for high taxes on the rich, in opposition to the material interests of these individuals. Further, although these individuals are concentrated in the far left of the figure, the income levels involved here encompass the majority of the population. In the treatment group, there is less of the ‘wrong’ reaction of preferences to increasing incomes, providing partial support to the idea that the policy information allows lower income voters better to align their attitudes with their material position.

Beyond the simple visualisation of the raw data, we can investigate hypothesis 2 with reference to models interacting the treatment with income levels. The most straightforward way

⁶In light of this, if we take advocating higher taxes on the bottom quintile as outliers, categorising ‘About the same’, ‘More’, and ‘Much more’ responses as outlying on the basis that they empirically rarely observed, rather than theoretically at the extreme of the possibilities, the estimated treatment effect is $\beta = 0.10$, but with a standard error of 0.24, so again indistinguishable from zero.




Figure 5: Support for higher levels of taxation on the top income quintile, by income and treatment status. Lines indicate smoothed lowess predictions; shaded areas represent 95% confidence intervals.

to do this is to model the attitude outcomes as dependent on the treatment, income, and their interaction in classical ordered logistic models. The results from such models are presented in ??.

This strategy reveals no evidence of a significant interaction in the determination of preferences over the overall level of taxation (columns 1 to 3), nor the preferred rates of taxation on the bottom quintile. However, it does appear that the provision of information has a differential effect on attitudes towards taxes on the richest quintile, depending on the respondent's income level. Specifically, the interaction between income and the treatment is negative, indicating that at higher incomes, the positive baseline impact of the treatment is muted. The precise interpretation of the magnitude of these effects is not transparent given the ordinal modelling strategy

DV:	Overall			Bottom fifth			Top fifth		
Model:	1	2	3	4	5	6	7	8	9
Treatment	0.17	0.17	0.12	0.41	0.36	0.41	0.70	0.71	0.71
<i>t value.</i> ⁷	<i>0.59</i>	<i>0.56</i>	<i>0.37</i>	<i>1.22</i>	<i>1.06</i>	<i>1.13</i>	<i>2.27</i>	<i>2.29</i>	<i>2.16</i>
Treatment × Income	-1.23	-1.19	-0.73	-0.06	0.24	-0.18	-4.81	-4.85	-5.19
	<i>-0.59</i>	<i>-0.57</i>	<i>-0.33</i>	<i>-0.03</i>	<i>0.10</i>	<i>-0.07</i>	<i>-2.27</i>	<i>-2.29</i>	<i>-2.34</i>
Treatment × Income ²	2.17	2.07	1.53	-0.67	-1.26	-0.85	3.23	3.33	4.44
	<i>0.80</i>	<i>0.77</i>	<i>0.54</i>	<i>-0.23</i>	<i>-0.42</i>	<i>-0.28</i>	<i>1.18</i>	<i>1.22</i>	<i>1.55</i>
Income	1.84	1.66	1.75	3.46	2.74	3.14	3.87	3.98	4.40
	<i>1.24</i>	<i>1.11</i>	<i>1.09</i>	<i>2.05</i>	<i>1.60</i>	<i>1.77</i>	<i>2.51</i>	<i>2.56</i>	<i>2.70</i>
Income ²	-1.81	-1.66	-1.98	-3.00	-2.57	-2.65	-4.10	-4.25	-5.16
	<i>-0.93</i>	<i>-0.85</i>	<i>-0.96</i>	<i>-1.41</i>	<i>-1.20</i>	<i>-1.20</i>	<i>-2.07</i>	<i>-2.13</i>	<i>-2.43</i>
Male		0.09	-0.01		0.11	0.15		-0.11	-0.14
		<i>0.72</i>	<i>-0.06</i>		<i>0.81</i>	<i>1.03</i>		<i>-0.91</i>	<i>-1.05</i>
Age		-0.40	-0.42		-2.72	-3.13		0.05	0.44
		<i>-1.03</i>	<i>-1.01</i>		<i>-6.49</i>	<i>-6.96</i>		<i>0.13</i>	<i>1.05</i>
Information			0.70			-0.55			0.56
			<i>3.18</i>			<i>-2.31</i>			<i>2.59</i>
<i>N</i>	967	967	837	923	923	820	920	920	811
<i>AIC</i>	2175	2177	1893	1752	1713	1504	2313	2316	2036

Table 1: Coefficients from ordinal logistics regression results of the interaction effects between income and the treatment. Italic values are *t*-values.

and the presence of the higher-order income terms, but ?? illustrates the patterns in the expected values.

Here we can see the impact of information is completely different at the high and low income groups. For the bottom income quintile, the treatment group sees many more individuals advocating higher taxes on the rich than on the middle quintile (the fourth diamond is much higher than the fourth circle). In contrast, in the top income group this pattern is reversed. In both cases the ‘action’ seems to be coming from the same place, that is, with fewer ‘about the same’ and ‘lower’ responses under treatment (control) for the bottom (top) quintile.

The second possibility is to consider particular income levels as groups by which the effect of the treatment might vary in a random effects model. Adopting this approach (not reported here for reasons of space) reinforces the conclusions implied by the interaction models: that lower income voters support higher taxes on the rich (though there is no effect in the other dimensions of redistribution).

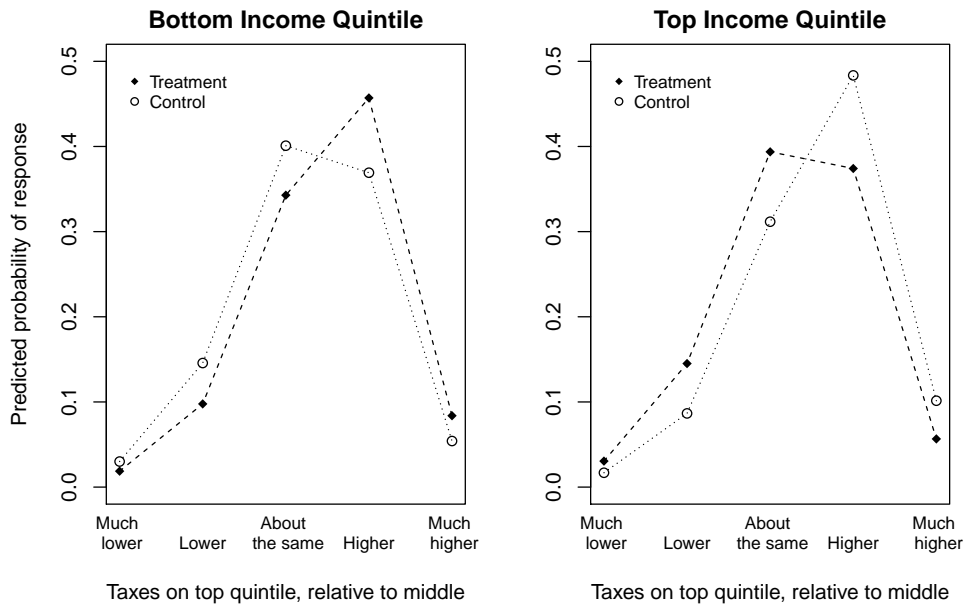


Figure 6: Predicted (fitted value) probabilities of each view on the relative burden the top quintile should bear, for the bottom and top income quintiles under treatment and control. Solid diamonds and dashed lines indicate treatment and open circles with dotted lines are control group averages.

5.3 Average Treatment Effect on Preferences

Finally, we consider the average treatment effect on preferences, again for the three outcomes of interest: attitudes towards the overall level of taxation; and those towards high and low incomes relative to middle incomes. For reasons of space, we do not present the full results, but only that where we see an effect in the pilot data. Here, there is no evidence of an overall treatment effect on attitudes towards overall levels of taxation, nor on attitudes towards taxing the rich more aggressively than the middle. In contrast, we do see support for taxes on the first income quintile differing, as illustrated in ??.

The pilot data thus indicates that, contrary to the literature, but in keeping with plausible Conservative party strategy in rolling out the tax statements immediately before the election, that the direction of the effect of information on preferences overall is in the direction of decreasing support for progressivity— but in the guise of leading to calls for higher (relative) taxes on lower incomes rather than lower taxes at the top. In contrast, the effects on support for higher (lower) overall tax levels seem very small. At least in this pilot study, then, the provision of de-




Figure 7: Share of respondents supporting different relative rates of taxation on the bottom income quintile, by treatment status. Vertical lines (in blue and red) represent 95% confidence intervals.

tailed budget information, which gives detailed information both on the level of spending (in terms of an individual’s tax contribution) and its distribution (across spending areas) seems to have more obvious effects on attitudes towards the shape of government policy than its overall size.

5.4 Power Implications from the Pilot Study

We plan to conduct a formal power analysis for the final pre-analysis plan. Even based on back-of-the-envelope calculations, however, this experiment is very well powered. In particular, we can assume 5,000 individuals in each treatment condition and a set of predictive baseline

covariates that explains around 15 percent of the residual variance, in line with the data from the pilot study. For a binary outcome, such as our outlier status measure, a conservative calculation for the Minimal Detectable Effect is around two percentage points. Therefore, even if the main effects remained the same as in the pilot study, we would have the power to detect changes in outlier share that are of the same order as the pilot estimates for total tax amount, and for taxes on the bottom quintile (in percentage point terms, these are approximately a three-point change for amount, and a nine-point change for bottom quintile taxes).

6 Contribution

First, we make explicit the extension of voter ignorance and political information to the inclusion of the policy status quo, as well as the position of parties (**Downs1957**) or voters' own preferences (**AlvarezBrehm2002**). Although the idea of voter uncertainty about the status quo has received some attention in theoretical treatments (**Banks1990**), it has been less explicitly attended to in empirical analyses of public opinion.

Second, and more significantly, the results of our field experiment will provide important evidence on the effect of information on preferences over redistribution. The most basic assumption— not easy to test— is that information in general will change (at least some) voters' attitudes towards policy. But even this has not been conclusively demonstrated, as the imperviousness of voters to new (especially disconfirming) evidence shows (**NyhanReifler2010**).

Thus in empirical terms we provide an extension of **Gilens2001** and most of the existing work on information in the provision of the relevant information in the field rather than in a purely experimental setting. If the finding that information matters to policy preferences has real-world implications, we need to know whether information levels are in fact amenable to intervention. Although a single-country study, our findings will highlight some of the assumptions, and assess some of the implications, of recent work situated centrally in comparative politics (for example, **Gingrich2014**).

The major strength of our study is in its field experimental design. This allows us to go beyond the impact of informational survey or laboratory treatments in order to consider the effects of information in the real world, where the context is much different. Substantively, competition for scarce attention, and the difficulty of translating information seen in one context to another are two important differences between the provision of information and the creation of informed voters.

Methodologically, not only is this important in helping to gauge the external validity of previous experiments on information effects, it also allows us to draw conclusions about real policy initiatives designed to inform voters about the government budget. Such initiatives are currently underway in both the UK and US. This contribution is reinforced by the nesting of a survey experiment—our pilot study—within the full design. Thus our design sits in a ‘sweet spot’ in terms of a rigorous experimental approach in a substantively important, real-world context. The tension in political science research between answering questions of real theoretical importance, and answering them with a credible identification strategy, motivates the ambitious nature of the proposal, and justifies the additional efforts involved in collaboration with government actors on a national-level policy.

In terms of transparency, this ambitious and original research design will provide an additional insight into the process of empirical political research. The implementation implications of collaboration with external partners on such a scale does represent a potential source of modifications to the design that are beyond the researchers’ control— but assessing policy in the field speaks with much more directness to the theoretical questions of interest, for reasons of external validity. The methodological challenges involved in the trade-off between internal validity and the implementation of empirical research that tracks important, theoretically motivated questions, will be directly documented through this project.