

Do attitudes towards policy changes depend on beliefs about policy levels?*

Itzhak Rasooly[†]

February 17, 2025

ABSTRACT. In theory, voter attitudes towards policy *changes* (e.g., whether to increase the minimum wage) ought to depend on their beliefs about the current *level* of the relevant policy variable. In this paper, I test this hypothesis using a large-scale ($n = 5,000$) and pre-registered survey experiment that spans four different policy areas. The experiment yields four main results. First, voters have both inaccurate and biased beliefs about the levels of the policy variables. Second, voters' attitudes are remarkably unresponsive to changes in their beliefs about levels: for example, exogenously increasing average beliefs about the top tax rate by ~ 8.5 percentage points does not increase the share who want to cut the top tax rate. Third, the observed unresponsiveness cannot be rationalised by a model in which voters form attitudes towards policy changes by comparing actual and preferred policy levels. Fourth, although attitudes are unresponsive to the quantitative information presented, they can be swayed by qualitative arguments.

KEYWORDS: political attitudes, survey experiment

JEL CODES: C91, D72

*For comments on the experimental design, I am especially grateful to Jesper Akesson, Dylan Balla-Elliott, Adam Brzezinski, Jeanne Hagenbach, Chris Roth and Jasmine Theilgaard. For useful comments and discussions, I would also like to thank Julia Cagé, Emeric Henry, Théo Marquis, Eduardo Perez-Richet, Stefan Pollinger, Nina Roussille, Roberto Rozzi, Toby Shevlane and Eugenio Verrina. This project received IRB approval from the Paris School of Economics (approval number 2023-033) and received funding from the Investissements d'Avenir program of the Agence Nationale de la Recherche (reference ANR-17-EURE-001).
Links: [[Pre-registration](#)] [[Experiment](#)] [[Replication package](#)]

[†]Sciences Po & Paris School of Economics.

1. INTRODUCTION

In theory, voter attitudes towards policy *changes* (e.g., whether to increase the minimum wage) ought to depend on their beliefs about the current *level* of the relevant policy variable. To illustrate, consider [Meltzer and Richard's \(1981\)](#) influential model of redistribution under political competition. In that model, each voter has a preferred rate at which taxes ought to be set; thus, whether a voter wants taxes to be increased depends on whether their preferred rate exceeds the actual rate. For this reason, *increasing* voter beliefs about the actual rate will generally *decrease* the share who want taxes to rise and *increase* the share who want tax cuts. Of course, this logic does not just apply to taxes: similar remarks apply to any model in which voters have stable, well-defined preferences over policy levels (see, e.g., the models surveyed in [Persson and Tabellini, 2002](#)).

In practice, however, one may doubt whether voter attitudes towards changes are as responsive to policy levels as such models may suggest. First, determining optimal levels is cognitively taxing: even a trained economist may struggle to decide the rate at which marginal income should be taxed, or the level at which the minimum wage ought to be set. Second, although voters often do have strong beliefs about how policy ought to be changed, one can think of reasons why these might be rather disconnected from beliefs about the underlying policy reality. For example, if an individual is eager to signal their nationalistic credentials, they may demand that the number of immigrants be reduced almost regardless of what the current level of migration happens to be.

This all raises a simple question: do attitudes towards policy changes depend on beliefs about policy levels? For example, does support for tax cuts depend on beliefs about current tax levels; does support for minimum wage increases depend on beliefs about the current minimum wage level; and does support for cutting unemployment benefits depend on beliefs about the generosity of the current benefits system? Answering these questions does not just inform our theoretical frameworks, but also helps us understand the likely results (and possibility of ‘backlash’) following changes to quantitative policies.

In this paper, I study these questions using a large-scale ($n = 5,000$) and fully pre-registered survey experiment that spans four different policy areas: the ‘Living Wage’ (the minimum wage in the UK), the top marginal rate of income tax, refugee inflows, and unemployment benefits. To generate exogenous variation in voter beliefs about current policy levels, I introduce and use the method of *random intervals*. Specifically, individuals are either informed that the policy variable lies in a high interval, or that it lies in a low interval; while these intervals are always truthful, they systematically shift beliefs in the expected directions. Af-

ter exogenously shifting beliefs, I then measure whether voters want to change the relevant policy variables (e.g. by changing the top marginal tax rate).

In addition to shocking individual beliefs about levels, I include two additional treatments that help one to interpret and contextualise the results. In a *qualitative treatment*, individuals are presented with a series of arguments that aim to ‘push’ their attitudes in a particular direction; this treatment allows one to benchmark the effects of the quantitative interventions. Meanwhile, in the *control group*, I explicitly measure individuals’ ‘ideal points’ (e.g. what they believe that the top tax rate ought to be) along with baseline beliefs about the policy variables. This allows one to obtain descriptive evidence on the accuracy of voter beliefs as well as to compute the theoretically predicted impact of the exogenous shifts in beliefs about policy levels.

The experiment yields four main findings.¹ First, the (incentivised) individual beliefs about the policy variables are rather inaccurate. Very few respondents can state the exact values of the policy variables even though most of the policies studied in the experiment (e.g. the top tax rate) are widely debated and discussed. Somewhat more surprisingly, the majority of respondents’ estimates are not even within 10% of the true value for any of the policy variables with the exception of the Living Wage. I also show that respondent beliefs are typically biased, and that these biases can operate in subtle ways. For example, whilst the median respondent greatly underestimates current refugee inflows to the UK, the mean respondent greatly overestimates them.

Second, I find that individual attitudes towards policy changes are remarkably unresponsive to exogenous changes in their beliefs about policy levels. For example, exogenously increasing average beliefs about the top tax rate by ~ 8.5 percentage points does not increase the share who want to cut the top tax rate. Likewise, exogenously increasing average beliefs about refugee inflows by $\sim 260,000$ (a factor ~ 11 relative to the baseline) has no discernible effect on the share who want fewer refugees. More generally, the exogenous shocks do not have significant effects on voter attitudes in any of the four policy areas studied.

Third, although voter attitudes are unresponsive to the quantitative information provided, they can be strongly swayed by more qualitative arguments. For example, the *deadweight loss argument* against minimum wages dramatically decreases the share who want minimum wages to be increased. Similarly, the *diminishing marginal utility argument* for wealth redistribution substantially increases the share who want a higher top tax rate. This means that the unresponsiveness documented earlier cannot be attributed to a general impossibility of

¹ All findings can be reproduced using the data and code in the [replication package](#).

influencing attitudes in my particular context (e.g. due to general inattentiveness to the information presented in the survey).

Fourth, the observed unresponsiveness to quantitative information cannot be reconciled with a model in which individuals derive their attitudes towards policy changes by comparing their belief about the current policy level with a stable ideal point. To show this, I measure the distribution of these ideal points and thus compute the share who should (in theory) change their policy attitudes following the exogenous shock to beliefs. I show that these predicted effects are generally substantially larger than the observed effects. For example, increasing individual beliefs about refugee inflows by about $\sim 260,000$ is predicted to increase the share who want fewer refugees by ~ 30 percentage points. In contrast, the observed effect is just 0.006 percentage points; it is thus easy to reject the null hypothesis that the predicted and observed effects are equal.

This paper contributes to several literatures across economics and related disciplines. First, it contributes to the literature on information provision; see [Haaland et al. \(2023\)](#) for an overview. In particular, this paper is the first to use the tools of information provision to study the general (and rather fundamental) question of whether voter attitudes towards changes depend on beliefs about levels. Reassuringly, one can detect traces of the general results reported here in two studies of immigration attitudes (namely, [Sides and Citrin, 2007](#) and [Hopkins et al., 2019](#); although see also [Grigorieff et al., 2020](#)). For example, [Hopkins et al., 2019](#) find that downwardly shocking beliefs about the share of the US population that is foreign-born—which might be viewed as a rough proxy for the current level of immigration—fails to increase the share of respondents who want more immigration. Despite this link, however, neither of these papers study the broader and more general question of whether attitudes towards policy changes depend on beliefs about policy levels. For this reason, they only examine one policy area (immigration), shock beliefs using a method that is vulnerable to salience confounds, and do not measure the distribution of ideal points (which is crucial for the analysis, as discussed below).²

² It is perhaps also worth noting the large number of papers in the information provision literature whose dependent variable overlap with one of the policy areas considered here. For example, [Cruces et al. \(2013\)](#); [Karadja et al. \(2017\)](#); [Alesina et al. \(2018\)](#); [Hoy and Mager \(2018\)](#); [Fehr et al. \(2022\)](#); [Roth et al. \(2022a\)](#); [Fehr et al. \(2024\)](#) study attitudes towards taxation (or government spending); while [Lergetporer et al. \(2021\)](#); [Facchini et al. \(2022\)](#); [Alesina et al. \(2023\)](#) study attitudes towards immigration. These studies are less relevant since they do not examine whether attitudes towards changing a policy depend on beliefs about the level of that same policy. For example, [Alesina et al. \(2018\)](#) study how beliefs about one topic (immigration) influence attitudes towards another (redistribution); [Cruces et al. \(2013\)](#) study beliefs about a personal characteristic (individual's standing in the income distribution) as opposed to a policy level; and [Kuziemko et al. \(2015\)](#) study how beliefs about the *share* who pay the estate tax (which is distinct from the *level* of the estate tax) influence attitudes towards changing the estate tax.

Second, this paper also contributes methodologically to the literature of information provision by introducing a new method—random intervals—for obtaining exogenous variation in beliefs. As I argue later on, this method avoids the salience confounds that afflict classical information provision. In addition, the method is much more flexible than classical information provision: it can be used to shock beliefs upwards, downwards, and by greater or lesser amounts. Finally, the method is portable and can be used to shock beliefs about essentially any (quantitative) variable, whether macroeconomic indicators (e.g. inflation), the prevalence of ‘pro-social’ behaviours, or the estimated returns to a particular investment.

Third, this paper contributes to the literature in political science on the ‘thermostatic’ model of public opinion (Wlezien, 1995; Soroka and Wlezien, 2010). According to this model, increases in the level of a ‘policy variable’ (e.g. defence spending) tend to reduce the share who want ‘more’ of the policy and increase the share who want ‘less’ — this is termed a ‘thermostatic’ response to the policy change.³ This hypothesis is precisely what I test with my experiment, with the important caveat that my experiment studies changes in *beliefs* about policy levels (not changes in the policy levels themselves).⁴ This caveat notwithstanding, this paper provides the first experimental test (and falsification) of the logic underlying the thermostatic model: contrary to what the model predicts, I do *not* observe thermostatic responses in any of the four policy areas that I consider.

Fourth, this paper contributes to the literature on beliefs about the economy and economic policy (see, e.g. Stantcheva, 2020). The descriptive findings on the inaccuracy of voter beliefs follow a long tradition of work in economics and political science. For example, Lewis (1978), Gideon (2017), Rees-Jones and Taubinsky (2020) and Stantcheva (2020) find, as I do, that voters typically underestimate the top marginal tax rate, a result that might be explained by the use of linear heuristics to approximate the (non-linear) tax schedule (Rees-Jones and Taubinsky, 2020). This paper extends these findings to topics that are less studied in the literature (namely, beliefs about unemployment benefits and the minimum wage). It also presents a finding that is in tension with existing results: namely, that the median respondent actually underestimates refugee inflows. As I argue later on, this discrepancy may be due to differences in the response scales used when eliciting voter beliefs.

³ In addition to proposing that changes in policy affect voter attitudes, Soroka and Wlezien (2010) also propose that changes in voter attitudes affect policy. Although this second hypothesis creates difficulties when testing the thermostatic model using observational data (Brezna, 2017), it has no bearing on the experiment whose results I report here.

⁴ This distinction matters: individuals can only respond ‘thermostatically’ to changes in policy if they are aware of these changes, which typically requires that these changes are reported in the mass media (Soroka and Wlezien, 2010). Thus, the thermostatic model requires both that individuals respond thermostatically to changes in their beliefs about policy levels (the assumption tested here) *and* that beliefs about policy levels are revised in light of policy changes (an additional assumption that I do not test).

Finally, the finding that the qualitative arguments presented have a much greater effect than the quantitative information presented fits with recent work on stories and narratives within economics; see, e.g., [Graeber et al. \(2022\)](#). Since the qualitative arguments are often drawn from basic economic theory, they also build on a literature demonstrating the ability of studying economics to alter political attitudes, whether by making individuals more supportive of free markets ([Fischer et al., 2017](#)), more likely to view market prices as fair ([Whaples, 1995](#)) and more like to view voluntary transactions as mutually beneficial ([Goossens and Méon, 2015](#)). The effects of the qualitative arguments reveal that individuals do not need to receive months of sustained economic instruction for political attitudes to be shaped by economic theory.

The remainder of this article is structured as follows. Section 2 presents the experimental design and analysis plan. Section 3 contains the results. Finally, Section 4 concludes with a discussion of why individuals fail to react ‘thermostatically’ in the realm of politics when they (presumably) do so in more routine economic environments.

2. EXPERIMENTAL DESIGN

In this section, I outline the design of the experiment.⁵

Sample. The experiment was conducted online in February 2024 using the Prolific platform. The experiment was large-scale: in total, $n = 5,000$ participants completed the survey (this was based on power calculations that are reported in [Haaland et al., 2023](#)). Since the experiment concerned UK government policy, I required all respondents to be adult UK residents. In order to minimise the chance of careless responses, I also required participants to pass an attention check (at the start of the survey) and to have a Prolific approval rate of at least 99%. Finally, I constructed the sample so that it roughly matched the age and gender distribution of the UK adult population.⁶

Overview. Before describing the experimental details, I first provide a broad overview of the experiment’s general structure. Subjects were randomly allocated into one of three treatment groups: an ‘information manipulation treatment’ (~50% of participants), a ‘qualitative

⁵ The exact wording of the survey can be viewed using [this link](#); see also Appendix B.

⁶ Specifically, the sample was constructed so that the fraction of the sample in each of a series of age ‘bins’ matched the corresponding fraction in the UK adult population (as estimated by [ONS, 2024](#)). The bins were: 18-24, 25-29, 30-34, 35-39, 40-44, 45-49, 50-54, 55+. Participants were also selected so that roughly half of those in each age ‘bin’ were female. Given that the experimental sample comprised a reasonably high fraction of all UK residents on the Prolific platform, it was difficult to make the sample representative along additional dimensions. As a result, I conduct robustness checks to verify that none of the descriptive results are driven by unrepresentative aspects of the sample.

treatment’ (~25% of participants) and a control group (the remainder). The information manipulation treatment studied the main question of the paper, namely whether changing beliefs about policy levels affects attitudes towards policy changes. The qualitative treatment benchmarked these results by studying the extent to which policy attitudes can be influenced by qualitative arguments. Finally, the control group directly elicited participants’ ‘ideal points’ (i.e. preferences over policy levels) to see if these were consistent with the results obtained from the information provision treatment.

Topics. In all treatment groups, subjects were asked about the following topics:

1. The National Living Wage, i.e. the minimum hourly wage that (nearly all) employees aged 23 or older are entitled to by law.
2. The top marginal tax rate on income.
3. Refugee inflows (more specifically, the number of refugees who were granted refugee permission following an asylum claim in the year ending September 2023).
4. Unemployment benefits (measured as the weekly payments of New Style Job Seeker’s Allowance that an eligible 25 year old would receive).

I selected these policies because they are naturally quantitative (in the sense that each is parameterised by a particular number). They are also salient within British political discourse and continue to provoke debate. In particular, the third topic (along with immigration more generally) was generally rated by the British public as one of the ‘most important issues facing the country’ in the months leading up to the survey ([YouGov, 2024](#)).

Information manipulation. As mentioned above, roughly half of participants were allocated to the information manipulation treatment (the main experiment). Participants in this treatment were asked about all four policy areas; and these policy areas were presented in a random order. Given each policy, the experiment proceeded in three steps. In the first step, participant beliefs about levels were randomly shocked. In the second step, participant beliefs about levels were measured. In the third step, participant attitudes towards policy change were elicited. I elaborate upon each of these steps below.

Shocking beliefs about levels. To generate exogenous variation in participants’ beliefs about levels, I provided participants with a ‘hint’ that the true policy level lies within a certain interval. However, different participants received different intervals: specifically, each participant was equally likely either receive a ‘low’ or ‘high’ interval (with independent randomisation across policy areas). [Table 1](#) displays the intervals that were displayed to the participants. Note that, although the intervals were chosen to systematically alter subject beliefs, they

were also truthful: as a result, the true answer always lies in the intervals’ intersection.

Table 1: Random intervals

Topic	Low interval	High interval	True value
Living Wage	£5/hour to £11/hour	£10/hour to £20/hour	£10.40/hour
Taxes	10% to 50%	40% to 90%	45%
Refugees	400 to 40,000	30,000 to 3,000,000	36,003
Benefits	£10/week to £90/week	£80/week to £800/week	£84.80/week

Notes. This table shows the intervals that subjects were shown in the information manipulation treatment. A subject was equally likely to be shown the ‘low’ or ‘high’ interval (for a given topic).

Before proceeding, I will briefly comment on some advantages of the ‘random interval’ approach (which, quite surprisingly, does not seem to have been used to study the impact of beliefs in previous work.) First, it is very general and can be used to generate exogenous variation in any (quantitative) belief. Second, it naturally controls for salience effects: although providing the interval may not just revise beliefs but also make them more salient, this is true in both the ‘low’ and ‘high’ interval groups. In this respect, it compares favourably to the more common practice of comparing information provision to a pure control (see [Haaland et al., 2023](#)), which can in theory generate salience confounds ([Akesson et al., 2022b](#)). Third, it is easier to implement and somewhat more flexible than existing methods that generate exogenous variation in beliefs while controlling for salience confounds. For example, [Akesson et al. \(2022a\)](#) shock beliefs using different expert estimates, which requires gathering these estimates and ties one to the particular estimates made by these experts. In contrast, there is tremendous flexibility in how one constructs the random intervals, as well as in the number of intervals that one chooses to display to different subjects.

Measuring beliefs. After randomly shocking subjects’ beliefs about levels, these beliefs were then measured. In general, this was done using free responses, i.e. subjects could enter any number they wished (provided the numbers were non-negative and were not so high that they violated pre-set ceilings).⁷ The one exception was the question about the top marginal tax rate, which was measured on a sliding scale from 0% to 100%. Immediately after measuring subjects’ beliefs, I asked subjects how confident they were in these beliefs (measured on a five-point scale from ‘very confident’ to ‘very unconfident’).

In the literature on information provision, belief measures are often incentivised through payments for correct guesses ([Haaland et al., 2023](#)). However, strong incentives can create

⁷ The ceilings were: £100/hour (minimum wage), 10,000,000 (refugees), £2,000/week (benefits). Since estimates of the top marginal tax rate were measured on a sliding scale from 0% to 100%, they were bounded by a ceiling of 100%. The ceilings were only binding for 0.02% of estimates.

problems if they encourage subjects to search for the answers online; for this reason, [Roth et al. \(2022b\)](#) and [Akesson et al. \(2022a\)](#) avoid incentives entirely. I take a compromise approach between these two extremes. Although answers are incentivised (subjects were told that they were more likely to win a £200 prize if they were more accurate), subjects were only told about the incentivisation once (at the start of the survey). In addition, subjects were told that they could lose their experimental earnings if they left their browser tab at any point during the experiment (one can track this using the method developed by [Permut et al., 2019](#)). It should perhaps be noted that, even if some subjects did search for answers online despite these measures (which was not easy to do in some of the policy areas), this does not invalidate the instrumental variables approach described below.⁸

Measuring attitudes towards change. After measuring subjects’ beliefs about levels, I then measured subjects’ attitudes towards policy changes. In the literature on information provision, such attitudes are often captured using ‘incentivised’ measures, e.g. by asking subjects whether they would like to donate money to an organisation that lobbies to shift policy in a particular direction ([Haaland et al., 2023](#)). While this approach has its merits, it also has some drawbacks. First, such measures are rather noisy since they also capture whether a subject wants to donate money that they have just received, along with whether they believe that donating to the relevant activist organisation is likely to change policy. Second, it is unclear whether such measures are fully ‘incentivised’: after all, subjects are extremely unlikely to change national policy by making a small donation to an activist organisation. Given these points, I instead measure policy preferences by simply asking subjects about these preferences, as is standard in political science (see, e.g., [Bartels, 2016](#); [Campbell, 2009](#); [Newman et al., 2015](#); [Horn et al., 2023](#)). More specifically, subjects were asked whether they would like to increase the relevant policy variable, decrease the relevant variable, or leave it unchanged; and were then asked for their degree of confidence in this judgement.

Demographics. At the end of the survey, I asked participants for some demographic-type information; this was done in all treatments. Specifically, I asked for their age, gender, country of residence (England, Scotland, Wales or Northern Ireland), highest level of education, employment status, average monthly post-tax income, views towards economic policy (left-wing, centrist, right-wing), whether they voted in the 2019 election, and (if so) which party they supported. I also asked respondents if both of their parents were born in the UK. These variables allow for various heterogeneity analyses and robustness checks (see Section 3).

⁸ In any event, the data presented in Section 3.1 suggest that online searches were rare: not one subject in the control group was able to give an exactly correct answer to the question about the Living Wage, the question about refugee inflows, or the question about unemployment benefits.

Qualitative treatment. As mentioned above, around 25% of subjects were allocated to a qualitative treatment that attempted to benchmark the results of the main experiment. As usual, subjects in this treatment were asked about the four topics above, again presented in random order. For each topic, around half of subjects were presented with a qualitative argument designed to sway their attitudes in a particular direction; the remaining subjects were not shown an argument. Subject attitudes towards policy changes, as well as confidence in these attitudes, were then measured (in the same way as before). The survey then concluded with the same set of demographic-type questions.

The arguments presented were ‘qualitative’ in the sense that they attempted to push subjects in a particular direction (e.g. towards wanting to raise taxes) without attempting to resolve the more difficult and quantitative question of what the optimal policy level should be. Often, these arguments were drawn from basic economic theory. The arguments were: (i) the *dead-weight loss argument* against minimum wages (which observes that minimum wages can prevent mutually beneficial trades if set above the marginal productivity of some workers); (ii) The *diminishing marginal utility argument* in favour of wealth redistribution and thus higher taxes (see [Pigou, 1920](#) and [Lerner, 1944](#) for relevant historical background); (iii) a *‘legal realities argument’* for allowing more refugees that stressed the current lack of safe and legal routes for those wishing to claim asylum in the UK; (iv) An *incentives argument* against unemployment benefits. More detail on the arguments can be found in [Appendix B](#).

Control group. The final ~25% of respondents were allocated to a control group. As usual, participants in this group were asked about the four policy areas, presented in random order. For each policy area, subject beliefs about levels (and confidence in these beliefs) were first elicited; this was done in the same way as in the main experiment. I then elicited subjects’ preferences over policy levels: for example, a subject who had just been asked to estimate the top marginal rate of income tax was then asked what the top marginal rate ought to be. As usual, I elicited subjects’ confidence in these attitudes and concluded the survey with the same set of demographic-type questions.

One use of the control group was to provide information on the prior beliefs and their accuracy. Note that, in contrast to those in the main experiment, these prior beliefs were unadulterated by the information contained within the ‘intervals’. More importantly, however, the control group also allows one to estimate what the results of information provision should be under the assumption that individuals have well-defined, stable preferences over policy levels (see [Section 3.4](#)).

Pre-registration. The experimental protocols were pre-registered on the AEA RCT Reg-

istry. In addition to pre-registering these protocols, I also pre-registered a detailed analysis plan specifying which analyses would be performed, which participants would be excluded, and so forth. With one relatively minor exception⁹, all of the analyses were conducted exactly as specified in the pre-registration. However, I occasionally also conduct additional analyses; in such cases, these analyses are always explicitly flagged as not pre-registered.

Two features of the analysis plan should perhaps be emphasised. First, the plan committed to dropping participants whose answers suggested inattentiveness (based on pre-set criteria). Ultimately, only 8 subjects were dropped according to these criteria, which might suggest that the criteria were quite conservative (or that the subjects were unusually attentive). Second, given that the analysis plan contained several hundred hypothesis tests, I pre-registered that p -values would be corrected using [Westfall and Young \(1993\)](#)’s correction on a ‘table by table’ basis. This is a bootstrap-based procedure that allows one to control the family-wise error rate while making minimal assumptions and allowing for correlation between tests (I use the implementation developed by [Jones et al., 2019](#)). For completeness, however, I also report conventional (i.e. uncorrected) standard errors throughout.

3. RESULTS

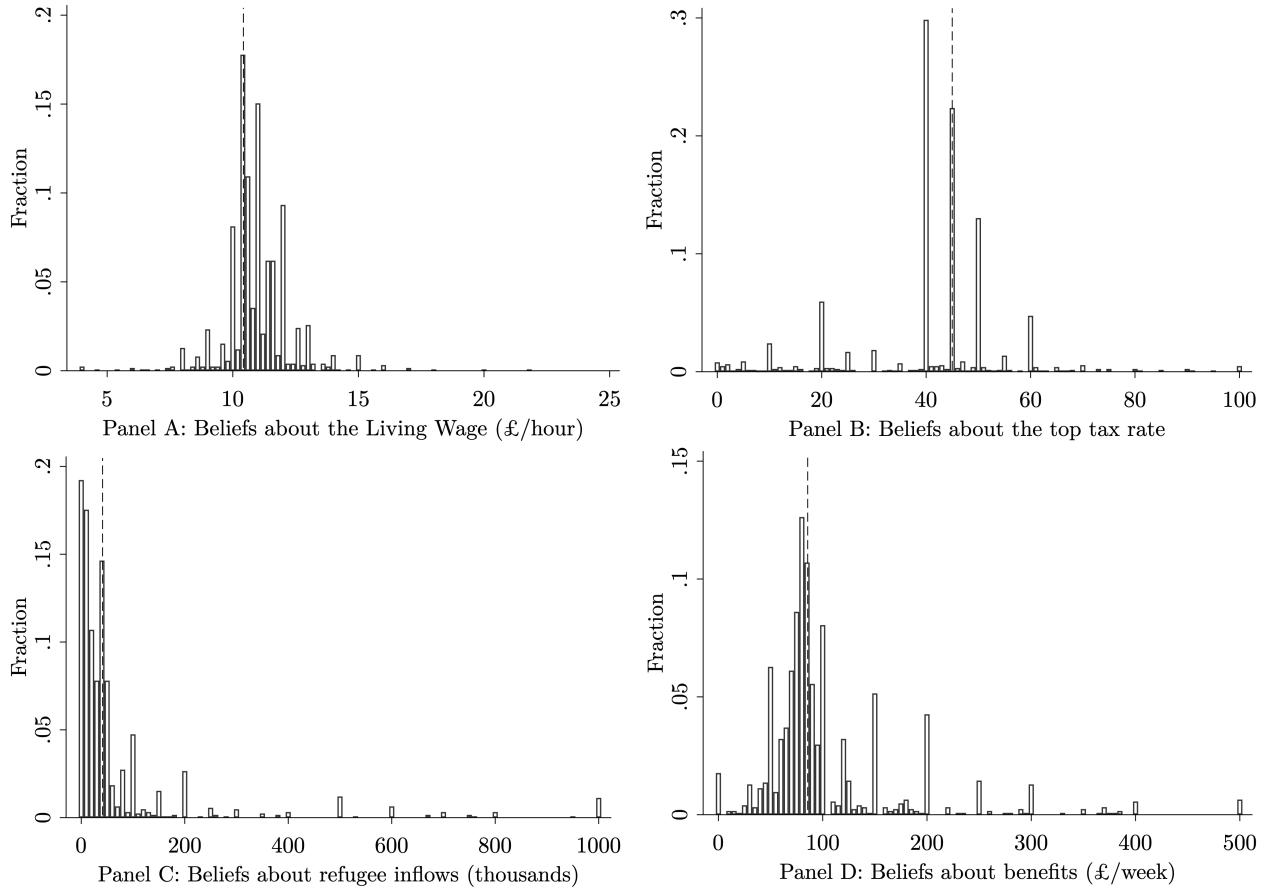
3.1 ACCURACY OF BELIEFS

I begin by outlining some descriptive results about the accuracy of voter beliefs. To do this, I examine data from the control group; recall that subjects in the control group were not given any ‘hints’ so provide the best source of evidence on beliefs in the wider population. [Figure 1](#) below plots the distribution of beliefs across the four policy areas: the Living Wage (Panel A), the top tax rate (Panel B), refugee inflows (Panel C), and benefits (Panel D). In each case, the vertical dotted line represents the true number; thus, for example, Panel A has a dotted line at £10.40/hour (the Living Wage at the time of the survey). Note that beliefs are binned into intervals and top censored for ease of visual exposition (this is not done in the formal analyses); details are provided in the notes below the figures.

First, the figures reveal that voter beliefs are generally quite inaccurate. No subject is able to give an exactly correct answer to the question about the Living Wage, the question about refugee inflows, or the question about unemployment benefits (22% can correctly identify

⁹ The one exception relates to the (pre-registered) computation of Westfall-Young corrected p -values. While it was possible to compute this correction for all regression-based estimates, it was not possible in a handful of cases — notably, those in which standard errors had themselves been computed via the bootstrap (see, e.g., the standard errors of medians reported in [Table A1](#)). In such cases, corrected p -values were computed using the method of [Holm \(1979\)](#).

Figure 1: Beliefs about the policy variables



Notes. This figure shows the beliefs of respondents in the control group about the policy variables. Beliefs are binned into intervals: specifically, they are rounded to the nearest £0.20/hour (living wage), 10,000 (refugees), and £5/week (benefits). Beliefs about refugee inflows that exceed 1 million are plotted as 1 million; an analogous comment applies to beliefs about benefits that exceed £500/week.

the top tax rate). Moreover, the large majority of beliefs about the top tax rate, refugee inflows and unemployment benefits differ from the true value by at least 10%: see Table 2. An important exception are beliefs about the Living Wage, which are within 10% of the true value in 67% of cases. This contrast is perhaps not surprising in light of models of rational inattention (Maćkowiak et al., 2023): while respondents have little incentive to hold accurate beliefs about the current level of refugee inflows, they might benefit from knowing the current level of the National Living Wage (e.g. if they are a low earner and want to ensure that they are paid the amount to which they are legally entitled).

Second, we see that, even when beliefs are roughly centred on the true value, they remain statistically biased (see Table A1). Despite the relative accuracy of beliefs about the Living Wage, one can easily reject the hypothesis that both the mean and median of the belief distribution are equal to the true value: instead, average beliefs are too high. Meanwhile,

Table 2: Overestimates and underestimates

Topic	Living Wage	Taxes	Refugees	Benefits
Underestimate	6.4	50.3	55.2	35.2
Within 10%	67.2	25.8	8.1	29.4
Overestimate	26.4	23.9	36.8	35.4
<i>n</i>	1,243	1,243	1,243	1,243

Notes. This table shows the fraction of the control group who overestimate or underestimate each policy variable by at least 10%.

both mean and median beliefs about the top marginal tax rate are too low; again, it is easy to reject the hypothesis of unbiasedness. A slightly more complex pattern appears in the case of unemployment benefits: although the mean belief (£104/week) is far too high, the median belief is very accurate and statistically indistinguishable from the true value (£80.40/week). Finally, a rather unexpected pattern arises in the case of refugee inflows: the median belief (25,000) is much too low, whereas the mean belief (79,043) is much too high.

It is useful to contextualise these results in light of recent evidence on the descriptive accuracy of voter beliefs. The general idea that voters have inaccurate beliefs about important policy variables is well established: for example, [Lewis \(1978\)](#), [Gideon \(2017\)](#), [Rees-Jones and Taubinsky \(2020\)](#) and [Stantcheva \(2020\)](#) all find that voters have inaccurate beliefs about the top tax rate. However, the exact patterns of inaccuracy documented here are rather new. In particular, the finding that most respondents underestimate refugee inflows somewhat contrasts with the general finding that individuals overestimate the number of migrants ([Citrin and Sides, 2008](#); [Grigorieff et al., 2020](#); [Hopkins et al., 2019](#)) and refugees ([Thorson and Abdelaaty, 2023](#)) in developed countries. While several explanations for this discrepancy are possible, one explanation is that this is largely driven by differences in response scales. Typically, studies ask individuals to estimate the number of migrants or refugees as a share of the total population; here, respondents are asked about the absolute number. Given that respondents are generally unaware of the true number migrants and refugees, it is reasonable to suppose that their answers can be heavily swayed by this difference in the question framing.

Before proceeding, I investigate whether the results are robust to re-weighting the sample to make it better reflect the voting behaviour of the general UK population.¹⁰ Such re-weighting is useful since the sample was only constructed to be representative based on age and gender, and turns out to heavily under-represent Conservative voters. The general patterns documented above (high levels of inaccuracy along with formally biased beliefs)

¹⁰ This exercise was not pre-registered in the original analysis plan.

persist even after this re-weighting is conducted; if anything, rates of inaccuracy become slightly worse (see Table A2 for the details).

Finally, I consider heterogeneity in accuracy rates. To do this, I define the ‘accuracy’ of a subject’s answer as (the negative of) the absolute difference between the answer they give and the true value. I then regress accuracy on the socio-demographic variables that are listed in the previous section. Somewhat surprisingly, the results do not strongly point to any important predictors of accuracy: see Table A3 for the details. This further suggests that the inaccuracy identified above is robust and persists across sub-groups.¹¹

3.2 DO INDIVIDUALS RESPOND ‘THERMOSTATICALLY’?

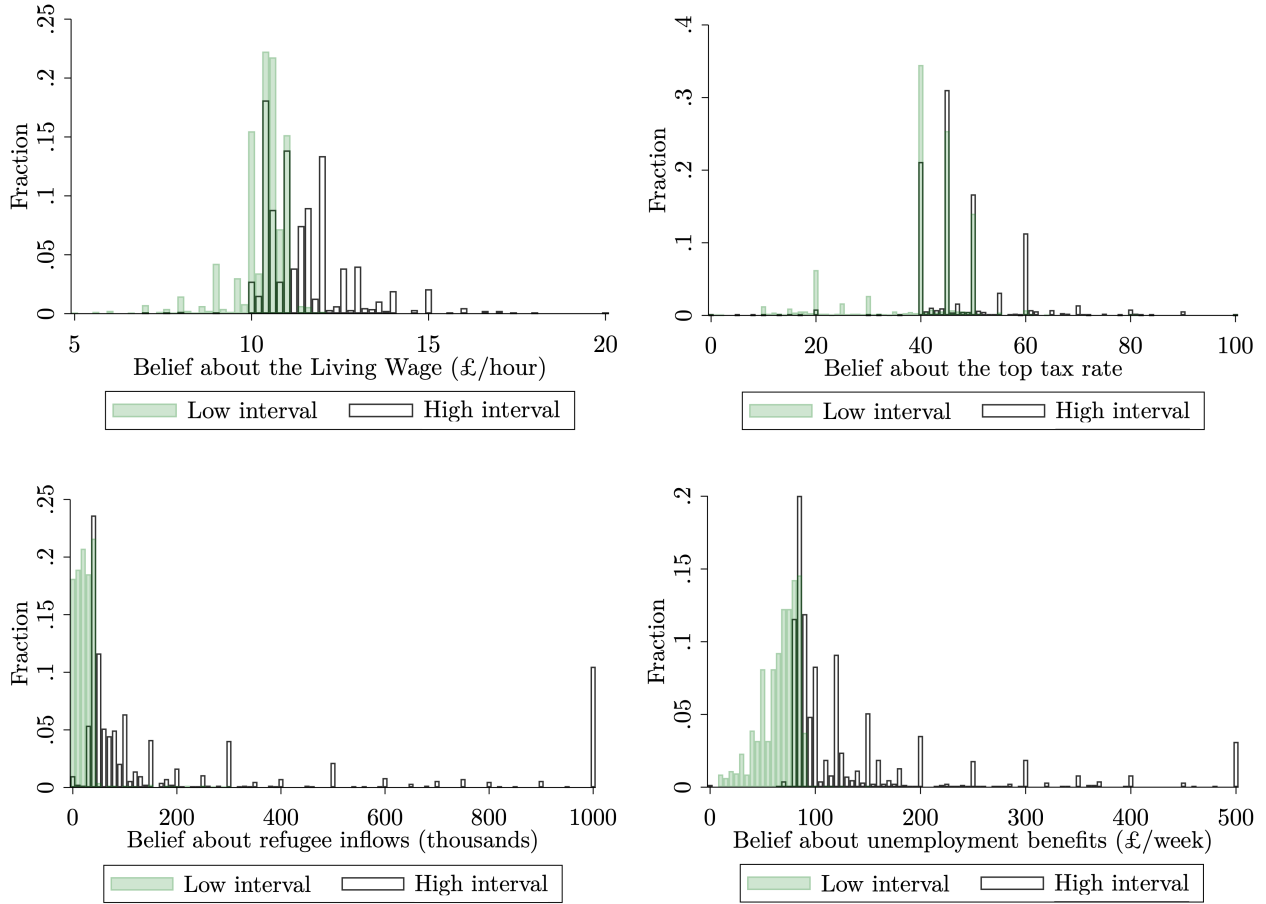
I now turn to the central question of the paper, i.e. whether individuals react ‘thermostatically’ to shocks to their beliefs about policy levels. To do this, I use an instrumental variables approach (similarly to, e.g., Akesson et al., 2022a). Specifically, I use subjects’ treatment assignment (i.e. whether they are exposed to a high or low interval) as an instrument for beliefs about policy levels in order to uncover the effect of these beliefs on attitudes towards policy changes. As usual, this approach requires both that the instruments are relevant and that the exclusion restriction holds; I discuss both these requirements below.

Instrument relevance. Figure 2 provides some visual evidence on whether the instruments are relevant; that is, on whether assigning subjects to a high, as opposed to low, interval substantially shifts their beliefs. As usual, the analysis is conducted separately for each policy variable; beliefs regarding each policy variable are shown in a separate panel. The green (translucent) distributions depict beliefs in the low interval treatments and the white (transparent) distributions depict beliefs in the high interval treatments; as before, the beliefs are binned into intervals and top censored for ease of visual exposition.

It is obvious from the figures that treatments reliably shift beliefs. Indeed, plotting the cumulative distributions reveals that, in every case, the distribution of beliefs amongst subjects in the high interval first-order stochastically dominates the distribution of beliefs amongst subjects in the low interval treatment. Table 3 confirms these results more formally by regressing beliefs on treatment assignment: the corresponding F -statistics are all above 200 and thus far in excess of the usual recommended thresholds (Andrews et al., 2019). It is perhaps worth noting that these effects are not mechanical: although subjects were always given hints, they were free to disregard these hints if they wished to do so.

¹¹ Instead of regressing accuracy on the demographic variables, one can instead regress beliefs on the demographic variables (see Table A4). None of the demographic variables appear to strongly predict beliefs.

Figure 2: Shifting beliefs



Notes. This figure shows the beliefs of respondents in the low and high interval groups. Beliefs are binned into intervals: specifically, they are rounded to the nearest £0.20/hour (living wage), 10,000 (refugees), and £5/week (benefits). Beliefs about refugee inflows that exceed 1 million are plotted as 1 million; an analogous comment applies to beliefs about benefits that exceed £500/week.

Exclusion restriction. Having shown that the instruments are relevant, I now turn to the question of whether the exclusion restriction holds; i.e., whether the instruments could influence the outcome (attitudes towards change) in a way that does not run through their effect on beliefs about policy levels. Typically, this restriction is not explicitly discussed in the information provision literature: in part, this presumably reflects an assumption that the restriction is so plausibly satisfied that it is hardly worth discussing.¹² This assumption is plausible: it is very hard to think of any relevant variable that information provision can change apart from the relevant beliefs. Nonetheless, it is still worth considering this issue in

¹² It can also reflect a focus on the effects of information provision *per se*: if one is only interested in the effect that providing information has (and uninterested using information as an instrument for beliefs), then the question of the exclusion restriction does not arise. However, it is also common to use information provision to study the impact of beliefs (Haaland et al., 2023).

Table 3: Instrument relevance

Topic	Living Wage	Taxes	Refugees	Benefits
High interval	1.188** [0.041]	8.500** [0.405]	257,448** [15,501]	93.6** [6.473]
Constant	10.272** [0.022]	39.740** [0.290]	22,550** [709]	66.8** [0.502]
<i>F</i> -statistic	827.44	440.9	275.82	209.12
<i>n</i>	2,475	2,475	2,475	2,475

Notes. This table shows the results of regressing beliefs about policy levels on a dummy variable indicating that the subject was shown a high interval. Adjusted *p*-values for treatment effects are computed using the Westfall-Young method (** $p < 0.01$, * $p < 0.05$).

some detail.

First, it is useful to emphasise two issues which do *not* introduce any bias into the estimates. The first issue is salience. In theory, information provision can not just update beliefs (measured as point estimates) but also affect the salience with which these beliefs are held (Akesson et al., 2022b). While this could conceivably be a problem for classical information provision, it is not a problem for the random interval approach used here: both treatment groups are presented with an interval, which means that the topic is made salient in both treatments. Second, as far as the exclusion restriction goes, one does not need to worry about so-called ‘cross learning’ (i.e. the possibility that changing beliefs about levels will in turn alter some other belief). The reason is that, although such learning could constitute a mechanism through which any observed effects occur, it operates via the independent variable of interest and thus it does not introduce bias into the instrumental variables estimates.

Indeed, the only violation that seems possible would seem to be an effect via confidence. That is, one could imagine that the instruments influence not just subjects’ beliefs (measured as point estimates) but also the confidence with which they hold these beliefs. However, this seems unlikely to bias our results. First of all, it is obvious that becoming more confident in one’s beliefs about a policy level (e.g., what the minimum wage happens to be) would have a substantial impact on the outcome of interest (e.g., whether one wants to cut the minimum wage). Thus, even if the treatments did influence confidence, it is not clear whether this would introduce any bias into the estimates. Second, as revealed by Table A5, the treatments do not influence measured confidence in three of the four policy areas. Thus, it seems extremely unlikely that confidence effects could explain the pattern of results documented in the subsequent sections.

Results. Having argued that the instruments are valid, I now turn to the main analysis. Table 4 reveals the results from the instrumental variable regressions; as usual, the analyses are conducted separately for each policy area. As stated in the pre-registration, the dependent variable is converted into two binary variables tracking whether a subject would like to increase or (respectively) decrease the policy variable. Each policy area is thus associated with two regressions (corresponding to these two measures of attitudes towards policy change).

Table 4: Instrumental variable regressions

Topic	Living Wage	Taxes	Refugees	Benefits
<i>Policy increases</i>				
Belief	-0.0104 [0.0141]	0.00279 [0.00231]	-2.21e-08 [7.02e-08]	-0.000265 [0.000210]
Constant	0.888** [0.153]	0.262** [0.102]	0.285** [0.014]	0.423** [0.026]
<i>Policy decreases</i>				
Belief	0.00192 [0.00310]	0.00165 [0.00173]	1.21e-07 [7.66e-08]	4.18e-05 [0.000153]
Constant	-0.0124 [0.0335]	0.0854 [0.0762]	0.390** [0.0150]	0.145** [0.0187]
<i>n</i>	2,475	2,475	2,475	2,475

Notes. This table shows the estimated effect of beliefs about policy levels on attitudes towards policy changes; beliefs are instrumented using treatment assignment. Adjusted p -values for treatment effects are computed using the Westfall-Young method (** $p < 0.01$, * $p < 0.05$).

Two results are apparent from the table. First, the estimated effects are very close to zero. For example, increasing individual beliefs about refugee inflows by $\sim 260,000$ (see Table 3) makes individuals 0.006 percentage points more likely to report wanting fewer refugees (a small but apparently ‘thermostatic’ reaction).¹³ Meanwhile, increasing beliefs about the top tax rate by 8.5 percentage points appears to *decrease* the share who want to cut taxes (a small but apparently ‘anti-thermostatic’ reaction). Second, and perhaps more importantly, none of the estimated effects are remotely significant by conventional criteria. It is perhaps worth noting that, although all p -values are computed with the multiple test correction (as stated in the pre-registration), this statement also applies to uncorrected p -values, none of which are significant at even the 10% level. Taken together, these points cast doubt on the

¹³ The 0.006 number is the ‘intention to treat’; the IV estimate from the table is simply this number divided by the average change in beliefs ($\sim 260,000$).

notion that individuals react ‘thermostatically’ in any of the policy areas considered in the experiment.

The lack of observed thermostatic responses is especially surprising once one recalls the exact structure of the experiment. In the experiment, subjects were first asked for their beliefs about policy levels and then asked about their attitudes towards policy changes. One might expect that this procedure makes beliefs about policy levels salient. The fact that beliefs about policy levels have such small effects, even after they have been made salient, provides especially strong evidence against the ‘thermostatic’ hypothesis. As we will later discuss, the lack of thermostatic responses is also hard to square with a model in which voters have stable, well defined preferences over policy levels.

Before proceeding, I conduct some robustness checks. As a first (pre-registered) robustness check, I repeat the instrumental variable analyses while controlling for socio-democratic characteristics. Such controls are useful since they can increase statistical power. As one would expect, controlling for these characteristics slightly reduces the standard errors associated with the estimates (see Table A6). However, it does not provide any more of a hint of thermostatic responses: if anything, adding controls shrinks the estimated coefficients slightly towards zero. As a second (pre-registered) robustness check, I repeat the instrumental variable regressions on the full sample, i.e. including those subjects who were dropped initially for providing apparently inattentive answers. As shown by Table A7, this yields near identical results as those obtained by the main regressions. This is to be expected given that just 8 of the 2,475 in the information manipulation provision treatment were dropped by the checks.¹⁴

3.3 ARE ATTITUDES JUST HARD TO CHANGE?

The previous section documents a lack of thermostatic reactions to changes in beliefs across our four policy areas. I now discuss how this unresponsiveness should be interpreted. A first interpretation is that the shocks fail to shift attitudes simply because political attitudes are generally difficult to change. Specifically, one might imagine that these attitudes are deeply ingrained (e.g., due to years of media consumption) and not the kind of thing that one might alter through an online experiment. Alternatively, but in a similar vein, one might suggest

¹⁴ One can also look for thermostatic responses by regressing attitudes towards change on beliefs about levels while controlling for the socio-demographic variables. Of course, this analysis should not be taken too seriously since it could reflect either omitted variable bias (if beliefs are correlated with uncontrolled variables that influence attitudes towards change) or reverse causality (e.g., if individuals construct beliefs about policy levels in a way that rationalises their attitudes towards policy changes). The correlational evidence does not point to thermostatic responses in the great majority of policy areas and measures of attitudes towards policy change (see Table A8).

that the subjects in the experiment do not pay enough attention to the information with which they are presented for the information to shape their political attitudes.

To test this interpretation, I turn to the qualitative treatment. Recall that subjects in this treatment were either exposed to a qualitative argument designed to sway their attitudes in a particular direction or asked to report their political attitudes without having been presented with an argument. To estimate the impact of these arguments, I regress attitudes on a dummy variable indicating treatment assignment. As usual, the analysis is conducted separately for every policy area; and political attitudes are measured using two different dummy variables which capture whether a subject wanted to increase or decrease the relevant policy variable.

The results are displayed in Table 5. As can be seen, the qualitative arguments have a remarkably strong effect on political attitudes. The ‘deadweight loss argument’ against minimum wages reduces the share who want to increase the living wage by over 20 percentage points (the share who want to decrease the living wage correspondingly increases, although more modestly). Meanwhile, the ‘diminishing marginal utility’ argument for increasing the top rate of income tax increases the share who want higher taxes on the rich by around 19 percentage points. One also observes substantial effects in the expected directions when examining the argument concerning refugees. The only argument that appears to be less persuasive is the ‘incentives argument’ against unemployment benefits. While this argument produces effects in the expected directions, the effect is only significant at the 5% level and for one of the two outcome variables once the multiple test correction has been applied.

Before proceeding, I check that the results are robust by re-estimating the effects after controlling for socio-demographic variables to increase power. In general, this produces similar estimates (see Table A9). The main difference is that the effect of the incentives argument against benefits on the share who want to cut benefits now appears to be rather stronger: as a result, the effect is now significant at the 1% level.

Taken together, the results here show that the political attitudes of the subjects in the sample can be changed through appropriate interventions. Although political attitudes are unresponsive to quantitative information presented, they can be swayed using qualitative arguments (see also Graeber et al., 2022). This benchmarking exercise aside, the results are also of independent interest insofar as they illustrate the ability of basic economic reasoning to change voters’ minds about important policy issues.

Table 5: Effects of the qualitative arguments

Topic	Living Wage	Taxes	Refugees	Benefits
<i>Policy increases</i>				
Argument	-0.205** [0.026]	0.193** [0.026]	0.132** [0.025]	-0.028 [0.026]
Constant	0.755** [0.017]	0.558** [0.019]	0.213** [0.016]	0.331** [0.018]
<i>Policy decreases</i>				
Argument	0.030** [0.009]	-0.073** [0.016]	-0.079* [0.028]	0.062* [0.025]
Constant	0.011** [0.004]	0.125** [0.013]	0.462** [0.019]	0.229** [0.016]
<i>n</i>	1,274	1,274	1,274	1,274

Notes. This table shows the results of regressing attitudes towards changes on a dummy variable indicating that a subject was shown a qualitative argument. Adjusted p -values are computed using the Westfall-Young method (** $p < 0.01$, * $p < 0.05$).

3.4 IS THERE A GAP IN THE DISTRIBUTION OF IDEAL POINTS?

Finally, I ask whether the results are consistent with a model in which voters form their attitudes towards policy changes by comparing their ‘ideal point’ (e.g., their preferred tax rate as in [Meltzer and Richard, 1981](#)) with their belief about the current policy level. As discussed earlier, any such model will generate thermostatic responses: for example, increasing beliefs about the policy level will weakly decrease the share who want the policy to be increased. In theory, however, the effect could be rather small. Indeed, if there were somehow a ‘gap’ in the distribution of ideal points over the relevant interval, the predicted thermostatic response would be zero! Given these observations, it is thus useful to compare the estimated effects with the effects that one would expect given the measured distributions of ideal points.

To do this, it is helpful to first introduce some notation. Fix a particular policy variable and let I denote a generic ideal point; that is, the policy level that a generic voter prefers. Let L and H respectively denote a generic belief about the policy level conditional on observing the low and high interval respectively. If voters form attitudes towards policy changes by comparing ideal points with beliefs about current policy levels, then they should want the policy to increase whenever their ideal point exceeds their belief about the current level. Thus, the share who want the policy variable to increase should be $\mathbb{P}(I > L)$ and $\mathbb{P}(I > H)$ in the low and high interval groups respectively. Thus, the difference across groups should

be $\mathbb{P}(I > L) - \mathbb{P}(I > H)$. One can obtain an unbiased estimate of this quantity using its sample counterpart, namely

$$\frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m \mathbb{1}(I_i > L_j) - \frac{1}{nm'} \sum_{i=1}^n \sum_{j=1}^{m'} \mathbb{1}(I_i > H_j),$$

where n is the sample size of the control group, m is the sample size in the low interval group, m' is the sample size in the high interval group, and $\mathbb{1}$ is the indicator function. Note that, intuitively, this just involves computing the fraction of pairs with the property that the ideal point exceeds the belief about levels in the low and high interval groups respectively. One can then compare this to the *actual* difference in attitudes observed between the high and low signal group, i.e.

$$\frac{1}{m} \sum_{i=1}^m \mathbb{1}(Y_i = 1) - \frac{1}{m'} \sum_{j=1}^{m'} \mathbb{1}(Y_j = 1),$$

where Y_i is a dummy variable that equals one if and only if individual i wants to increase the policy variable. To test the null hypothesis that the predicted thermostatic responses match the observed responses, one considers the difference the differences

$$\hat{d} \equiv \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m \mathbb{1}(I_i > L_j) - \frac{1}{nm'} \sum_{i=1}^n \sum_{j=1}^{m'} \mathbb{1}(I_i > H_j) - \left(\frac{1}{m} \sum_{i=1}^m \mathbb{1}(Y_i = 1) - \frac{1}{m'} \sum_{j=1}^{m'} \mathbb{1}(Y_j = 1) \right)$$

The hypothesis that the observed effect is the same as the predicted effect corresponds to the hypothesis that $d = 0$. To test this, it is necessary to compute the standard error of \hat{d} : I do this using the bootstrap with 10,000 repetitions.

Table 6: Observed and predicted thermostatic effects

Topic	Living Wage	Taxes	Refugees	Benefits
Actual Δ	0.012	-0.024	0.006	0.025
Predicted Δ	0.156	0.160	0.303	0.325
Δ in Δ	0.144**	0.184**	0.298**	0.300**
	[0.018]	[0.021]	[0.019]	[0.021]
n	3,718	3,718	3,718	3,718

Notes. This table compares the actual (first row) and predicted (second row) thermostatic effects. The standard error of the difference between these effects is computed using the bootstrap (with 10,000 repetitions). Adjusted p -values are computed using the Holm-Šidák method (** $p < 0.01$, * $p < 0.05$).

Table 6 reveals the result of this exercise. The first row reveals the difference in the share who want to increase the policy variable across the low and high interval groups (this corresponds to the ‘intention to treat’ from the instrumental variables estimates). In other words, the first row shows the *actual* effect of the belief shock. Meanwhile, the second row reveals the theoretically predicted effect (computed as above using the distributions of ideal points and beliefs about policy levels). The third row shows the estimated difference between the actual and predicted effects. The main message from the analysis is very clear: generally, predicted effects are at least an order of magnitude larger than the actual effect. For example, although increasing beliefs about refugee inflows reduces the share who want to increase it by an estimated 0.006 percentage points, it should in theory reduce the share by an estimated 30.3 percentage points. Given the large disparity between estimated and predicted effects, it is easy to statistically reject the hypothesis that voters derive their attitudes towards policy changes from stable, well-defined preferences over policy levels.¹⁵

3.5 ADDITIONAL ANALYSES

Before concluding, I report the results of two additional analyses. First, I compare how confident individuals are in their attitudes towards policy changes with how confident they are in their attitudes towards policy levels; recall that the latter is elicited in the control group. The results are displayed in Table 7. As can be seen, subjects are more confident about how policy should be changed than what policy ought to be in all four policy areas studied in the experiment. Needless to say, this difference can be explained in various ways. However, one explanation which is consistent with the previous analysis is that subjects do not derive their attitudes towards policy changes by first considering their attitudes about policy levels. This explains why subjects are able to have confident opinions about how policy changes while having very little idea about what policy ought to be.

Second, I examine the relationship between beliefs about policy levels and respondents’ ideal points; this comparison just uses respondents in the control group. Generally, the analysis suggests a positive relationship between the two variables, with estimated correlations ranging from 0.12 (in the case of benefits) to 0.59 (in the case of taxes) — see Table A10 for details. Again, this relationship admits multiple interpretations (unlike the causal estimates discussed in Section 3.2). However, a natural interpretation in light of the previous findings is that individuals simply do not have stable, well-defined preferences over policy levels. Thus, when forced to state the level at which a policy variable ought to be, their answers

¹⁵ While the pre-registration committed to conducting this analysis using attitudes towards policy increases as the dependent variable, it is also possible to conduct the analysis using attitudes towards policy decreases. As one might expect, these exercises yield very similar results.

Table 7: Confidence about levels and changes

Topic	Living Wage	Taxes	Refugees	Benefits
Changes	0.869	0.808	0.757	0.735
Levels	0.688	0.611	0.404	0.488
Difference	0.181** [0.012]	0.198** [0.014]	0.353** [0.015]	0.247** [0.015]
<i>n</i>	4,992	4,992	4,992	4,992

Notes. This table compares confidence about levels and changes using an unpaired *t*-test. Adjusted *p*-values are computed using the Westfall-Young method (** $p < 0.01$, * $p < 0.05$).

are heavily ‘anchored’ on the beliefs about policy levels that they have just provided.

4. CONCLUSION

In this paper, I study whether individuals react ‘thermostatically’ to large shocks in their beliefs about policy levels. The main result is that individuals’ attitudes are remarkably unresponsive to these shocks: for example, increasing average beliefs about refugee inflows by $\sim 260,000$ has no discernible effect on the share who want fewer refugees. This finding casts doubt on models in which individuals derive attitudes towards policy changes by comparing actual and preferred policy levels; and indeed such models predict effects that are generally much larger than those seen in the data.

This finding calls into question the body of evidence purporting to support the thermostatic model (see, for instance, [Soroka and Wlezien, 2010](#)). It is perhaps worth noting that, unlike the experimental results here, existing work is based on correlational evidence. The findings also have implications for how voter attitudes ought to be modelled. As discussed earlier, there is a long tradition of work in political economy that views voters as holding stable ideal points about policy variables (see, e.g., [Meltzer and Richard, 1981](#)). Instead of modelling voters in this way, it may be more productive to view voters as holding strong views about how policies should be changed without deriving these opinions from preferences over levels.

The findings also raise the question of why individuals fail to react thermostatically in the realm of politics when they (presumably) do so in more routine economic environments. For example, one would assume that whether a shopper wants to buy more apples depends on the number of apples that are currently in her shopping basket. One answer is based on cognitive complexity: although it is relatively easy to decide how many apples one wishes to buy, it is extremely difficult to decide (for example) the level at which the minimum wage

ought to be set. Another answer is based on model of rational inattention ([Caplan, 2000](#); [Maćkowiak et al., 2023](#)): if individuals sense that they are unlikely to determine political outcomes through their voting decisions, then they may be unwilling to spend the time and effort required to arrive at well defined preferences over policy levels.

REFERENCES

- Akesson, J., Ashworth-Hayes, S., Hahn, R., Metcalfe, R., and Rasooly, I. (2022a). Fatalism, beliefs, and behaviors during the COVID-19 pandemic. *Journal of Risk and Uncertainty*, 64(2):147–190.
- Akesson, J., Hahn, R. W., Metcalfe, R. D., and Rasooly, I. (2022b). Race and redistribution in the united states: An experimental analysis. Technical report, National Bureau of Economic Research.
- Alesina, A., Miano, A., and Stantcheva, S. (2023). Immigration and redistribution. *The Review of Economic Studies*, 90(1):1–39.
- Alesina, A., Stantcheva, S., and Teso, E. (2018). Intergenerational mobility and preferences for redistribution. *American Economic Review*, 108(2):521–554.
- Andrews, I., Stock, J. H., and Sun, L. (2019). Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics*, 11:727–753.
- Bartels, L. M. (2016). *Unequal democracy: The political economy of the new gilded age*. Princeton University Press.
- Breznau, N. (2017). Positive returns and equilibrium: Simultaneous feedback between public opinion and social policy. *Policy Studies Journal*, 45(4):583–612.
- Campbell, A. L. (2009). How Americans think about taxes: Lessons from the history of tax attitudes. In *Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association*, volume 102, pages 157–164. JSTOR.
- Caplan, B. (2000). Rational irrationality: a framework for the neoclassical-behavioral debate. *Eastern Economic Journal*, 26(2):191–211.
- Citrin, J. and Sides, J. (2008). Immigration and the imagined community in europe and the united states. *Political Studies*, 56(1):33–56.
- Cruces, G., Perez-Truglia, R., and Tetaz, M. (2013). Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *Journal of Public Economics*, 98:100–112.
- Facchini, G., Margalit, Y., and Nakata, H. (2022). Countering public opposition to immigration: The impact of information campaigns. *European Economic Review*, 141:103959.

- Fehr, D., Mollerstrom, J., and Perez-Truglia, R. (2022). Your place in the world: Relative income and global inequality. *American Economic Journal: Economic Policy*, 14(4):232–268.
- Fehr, D., Müller, D., and Preuss, M. (2024). Social mobility perceptions and inequality acceptance. *Journal of Economic Behavior & Organization*, 221:366–384.
- Fischer, M., Kauder, B., Potrafke, N., and Ursprung, H. W. (2017). Support for free-market policies and reforms: Does the field of study influence students’ political attitudes? *European Journal of Political Economy*, 48:180–197.
- Gideon, M. (2017). Do individuals perceive income tax rates correctly? *Public Finance Review*, 45(1):97–117.
- Goossens, A. and Méon, P.-G. (2015). The belief that market transactions are mutually beneficial: A comparison of the views of students in economics and other disciplines. *The Journal of Economic Education*, 46(2):121–134.
- Graeber, T., Zimmermann, F., and Roth, C. (2022). Stories, statistics, and memory. *CESifo Working Paper*.
- Grigorieff, A., Roth, C., and Ubfal, D. (2020). Does information change attitudes toward immigrants? *Demography*, 57(3):1117–1143.
- Haaland, I., Roth, C., and Wohlfart, J. (2023). Designing information provision experiments. *Journal of Economic Literature*, 61(1):3–40.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, pages 65–70.
- Hopkins, D. J., Sides, J., and Citrin, J. (2019). The muted consequences of correct information about immigration. *The Journal of Politics*, 81(1):315–320.
- Horn, A., Kevins, A., and Van Kersbergen, K. (2023). Workfare and attitudes toward the unemployed: New evidence on policy feedback from 1990 to 2018. *Comparative Political Studies*, page 00104140231178743.
- Hoy, C. and Mager, F. (2018). Can information about inequality and social mobility change preferences for redistribution? Evidence from randomized controlled trials in 11 high and middle-income countries. *Tax and Transfer Policy Institute Working Paper*.

- Jones, D., Molitor, D., and Reif, J. (2019). What do workplace wellness programs do? Evidence from the Illinois workplace wellness study. *The Quarterly Journal of Economics*, 134(4):1747–1791.
- Karadja, M., Mollerstrom, J., and Seim, D. (2017). Richer (and holier) than thou? the effect of relative income improvements on demand for redistribution. *Review of Economics and Statistics*, 99(2):201–212.
- Kuziemko, I., Norton, M. I., Saez, E., and Stantcheva, S. (2015). How elastic are preferences for redistribution? Evidence from randomized survey experiments. *American Economic Review*, 105(4):1478–1508.
- Lergetporer, P., Piopiunik, M., and Simon, L. (2021). Does the education level of refugees affect natives' attitudes? *European Economic Review*, 134:103710.
- Lerner, A. P. (1944). *The economics of control: Principles of welfare economics*. Macmillan and Company Limited, New York.
- Lewis, A. (1978). Perceptions of tax rates. *British Tax Review*, pages 358–366.
- Maćkowiak, B., Matějka, F., and Wiederholt, M. (2023). Rational inattention: A review. *Journal of Economic Literature*, 61(1):226–273.
- Meltzer, A. H. and Richard, S. F. (1981). A rational theory of the size of government. *Journal of Political Economy*, 89(5):914–927.
- Newman, B. J., Hartman, T. K., Lown, P. L., and Feldman, S. (2015). Easing the heavy hand: Humanitarian concern, empathy, and opinion on immigration. *British Journal of Political Science*, 45(3):583–607.
- ONS (2024). Principal projection – UK population in age groups. <https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationprojections/datasets/tablea21principalprojectionukpopulationinagegroups> [Accessed: March 2024].
- Permut, S., Fisher, M., and Oppenheimer, D. M. (2019). Taskmaster: A tool for determining when subjects are on task. *Advances in Methods and Practices in Psychological Science*, 2(2):188–196.
- Persson, T. and Tabellini, G. (2002). *Political economics: explaining economic policy*. MIT press.
- Pigou, A. C. (1920). *The Economics of Welfare*. Macmillan and Company Limited, London.

- Rees-Jones, A. and Taubinsky, D. (2020). Measuring “schmeduling”. *The Review of Economic Studies*, 87(5):2399–2438.
- Roth, C., Settele, S., and Wohlfart, J. (2022a). Beliefs about public debt and the demand for government spending. *Journal of Econometrics*, 231(1):165–187.
- Roth, C., Settele, S., and Wohlfart, J. (2022b). Beliefs about public debt and the demand for government spending. *Journal of Econometrics*, 231(1):165–187.
- Sides, J. and Citrin, J. (2007). How large the huddled masses? the causes and consequences of public misperceptions about immigrant populations. In *annual meeting of the Midwest Political Science Association, Chicago, IL*. Citeseer.
- Soroka, S. N. and Wlezien, C. (2010). *Degrees of democracy: Politics, public opinion, and policy*. Cambridge University Press.
- Stantcheva, S. (2020). Understanding economic policies: What do people know and learn. *Unpublished Manuscript, Harvard University*, pages 2309–2369.
- Thorson, E. and Abdelaaty, L. (2023). Misperceptions about refugee policy. *American Political Science Review*, 117(3):1123–1129.
- Westfall, P. H. and Young, S. S. (1993). *Resampling-based multiple testing: Examples and methods for p-value adjustment*. John Wiley & Sons.
- Whaples, R. (1995). Changes in attitudes among college economics students about the fairness of the market. *The Journal of Economic Education*, 26(4):308–313.
- Wlezien, C. (1995). The public as thermostat: Dynamics of preferences for spending. *American Journal of Political Science*, pages 981–1000.
- YouGov (2024). The most important issues facing the country. <https://yougov.co.uk/topics/society/trackers/the-most-important-issues-facing-the-country> [Accessed: March 2024].

APPENDIX A. TABLES AND FIGURES

Table A1: Average beliefs in the control group

Topic	Mean	Median	True value
Living Wage	10.9** [0.038]	10.9** [0.0978]	£10.42/hour
Taxes	40.7** [0.417]	40.0** [1.576]	45%
Refugees	79,043** [7,057]	25,000** [2,143]	36,003/year
Benefits	104.1** [2.104]	£84.4 [1.193]	£84.80/week
<i>n</i>	1,243	1,243	1,243

Notes. This table shows the mean and median beliefs in the control group about each policy variable; the standard errors of the medians are computed using the bootstrap. The asterisks correspond to tests of the hypothesis that the mean (or median) equals the true value. Adjusted *p*-values are computed using the Holm-Šidák method (** $p < 0.01$, * $p < 0.05$).

Table A2: Overestimates and underestimates (re-weighted)

Topic	Living Wage	Taxes	Refugees	Benefits
Underestimate	6.8	51.7	53.8	36.0
Within 10%	67.8	25.1	7.8	26.8
Overestimate	25.4	23.3	38.4	37.3
<i>n</i>	1243	1243	1243	1243

Notes. This table shows the fraction of the control group who overestimate or underestimate each policy variable by at least 10%. Responses have been re-weighted so the distribution of voting behaviour in the sample matches the distribution in the UK population.

Table A3: Heterogeneity in accuracy

Topic	Living Wage	Taxes	Refugees	Benefits
Age	-0.0105 [0.00250]	0.0270 [0.0294]	1,373 [758.9]	0.484 [0.194]
Male	0.217 [0.255]	-0.250 [3.480]	-28,296 [63,711]	18.76 [33.28]
Female	0.374 [0.251]	-3.509 [3.481]	-22,950 [60,235]	12.35 [33.26]
England	-0.229 [0.227]	-0.743 [2.398]	-75,608 [27,067]	6.701 [18.42]
Scotland	-0.142 [0.246]	2.509 [2.759]	-49,981 [27,279]	16.79 [19.97]
Wales	0.00439 [0.260]	1.362 [2.712]	-151,329 [56,953]	15.58 [19.12]
Further education	-0.0512 [0.0819]	3.001 [1.226]	-24,908 [23,631]	-3.721 [4.369]
Undergraduate	-0.364 [0.0888]	4.002 [1.195]	-11,740 [16,450]	-7.27 [4.355]
Postgraduate	-0.459 [0.120]	4.246 [1.346]	-19,534 [20,501]	-19.64 [8.082]
Both parents UK	0.0199 [0.0812]	-0.941 [0.978]	20,386 [33,899]	-1.042 [5.549]
Part time	0.121 [0.0916]	-0.184 [1.156]	-1,494 [14,269]	-0.116 [6.206]
Self-employed	0.00492 [0.0958]	1.408 [1.146]	-6,622 [15,046]	7.473 [6.611]
Unemployed	0.0717 [0.114]	0.754 [1.463]	-66,347 [54,063]	8.457 [6.616]
Student	-0.216 [0.164]	-0.0977 [2.008]	28,840 [44,085]	-17.33 [12.41]
Retired	-0.131 [0.115]	1.678 [1.055]	-44,017 [23,321]	-3.667 [6.722]

Table A3: Heterogeneity in accuracy (continued)

Topic	Living Wage	Taxes	Refugees	Benefits
£1,000-£1,500	0.140 [0.134]	-0.188 [1.676]	17,648 [50,258]	4.154 [5.566]
£1,500-£2,000	0.147 [0.130]	1.951 [1.590]	39,915 [45,495]	-3.138 [6.841]
£2,000-£3,000	-0.000421 [0.134]	2.409 [1.578]	25,909 [43,226]	-10.78 [6.618]
£3,000-£4,000	0.0134 [0.146]	4.023 [1.626]	12,983 [49,826]	-1.387 [7.029]
More than £4,000	0.123 [0.141]	4.177 [1.709]	20,894 [44,840]	-7.996 [7.216]
Centrist	0.0737 [0.0788]	0.586 [0.849]	8,295 [21,868]	-0.425 [5.521]
Right-wing	0.0891 [0.127]	1.230 [1.297]	-43,476 [41,450]	-4.749 [7.201]
Don't know	-0.0242 [0.124]	0.943 [1.480]	19,424 [33,318]	2.767 [7.529]
Did not vote	-0.251 [0.124]	-0.417 [1.540]	4,288 [36,612]	-4.161 [8.294]
Voted Conservative	-0.120 [0.126]	1.274 [1.535]	38,354 [33,628]	7.614 [7.365]
Voted Labour	-0.0854 [0.119]	0.620 [1.418]	29,984 [37,693]	6.682 [8.359]
Voted SNP	-0.239 [0.187]	-1.998 [2.425]	19,243 [40,084]	-28.12 [18.79]
Voted Lib Dem	-0.309 [0.164]	1.798 [1.636]	51,474 [33,414]	8.887 [8.300]
Constant	-0.246 [0.391]	-15.17** [4.795]	-69,996 [96,787]	-76.25* [38.61]
<i>n</i>	1,243	1,243	1,243	1,243

Notes. This table shows the results of regressing belief accuracy of subjects in the control group on demographic variables. Adjusted p -values for treatment effects are computed using the Westfall-Young method (** $p < 0.01$, * $p < 0.05$).

Table A4: Heterogeneity in beliefs

Topic	Living Wage	Taxes	Refugees	Benefits
Age	0.00747 [0.00313]	0.0225 [0.0364]	-1,490 [772.0]	-0.603 [0.216]
Male	-0.128 [0.438]	-3.892 [4.365]	51,970 [69,487]	-9.614 [37.05]
Female	-0.343 [0.435]	-7.470* [4.357]	39,037 [66,243]	-9.605 [37.03]
England	0.0807 [0.311]	3.116 [2.800]	86,458 [28,066]	-10.66 [20.29]
Scotland	0.108 [0.334]	8.054 [3.225]	57,481 [28,729]	-17.73 [22.15]
Wales	-0.246 [0.352]	6.313* [3.229]	163,825 [58,504]	-22.63 [21.18]
Further education	0.0655 [0.0992]	3.867 [1.501]	26,926 [24,244]	7.344 [5.230]
Undergraduate	0.135 [0.108]	5.058 [1.468]	14,344 [17,198]	6.023 [5.260]
Postgraduate	0.240 [0.148]	4.93 [1.646]	20,843 [21,400]	19.61 [9.013]
Both parents UK	0.104 [0.107]	-0.185 [1.181]	-20,761 [34,350]	0.0961 [6.269]
Part time	-0.110 [0.113]	-0.978 [1.385]	4,007 [15,105]	-5.673 [7.069]
Self-employed	-0.0487 [0.124]	2.88 [1.456]	8,304 [15,965]	-9.082 [7.342]
Unemployed	-0.126 [0.137]	0.645 [1.828]	69,192 [54,639]	-8.188 [7.617]
Student	-0.0109 [0.206]	2.527 [2.513]	-26,229 [45,101]	12.46 [14.41]
Retired	0.0663	1.388	45,258	4.101

Table A4: Heterogeneity in beliefs (continued)

Topic	Living Wage	Taxes	Refugees	Benefits
£1,000-£1,500	0.0264 [0.176]	0.341 [2.072]	-21,465 [50,906]	-10.63 [6.939]
£1,500-£2,000	0.0135 [0.169]	2.443 [1.964]	-41,217 [46,049]	-1.915 [7.900]
£2,000-£3,000	0.290* [0.173]	2.707 [1.933]	-22,482 [43,744]	9.257 [7.702]
£3,000-£4,000	0.268 [0.185]	6.214 [2.011]	-6,907 [50,388]	-3.886 [8.178]
More than £4,000	0.0661 [0.185]	4.744 [2.086]	-17,137 [45,451]	4.655 [8.440]
Centrist	-0.0189 [0.0990]	0.643 [1.053]	-4,861 [22,247]	1.495 [6.077]
Right-wing	-0.126 [0.158]	0.376 [1.575]	47,246 [42,166]	4.286 [8.068]
Don't know	-0.0734 [0.160]	-0.635 [1.782]	-12,864 [34,082]	-12.99 [9.089]
Did not vote	0.146 [0.161]	-4.403 [1.941]	-5,297 [37,422]	2.475 [9.730]
Voted Conservative	0.162 [0.156]	0.868 [1.923]	-36,006 [34,403]	-6.983 [8.635]
Voted Labour	0.107 [0.149]	-2.322 [1.793]	-33,026 [38,400]	-3.655 [9.443]
Voted SNP	0.281 [0.241]	-6.437 [3.058]	-23,813 [41,446]	27.71 [21.06]
Voted Lib Dem	0.191 [0.207]	-0.342 [2.045]	-54,757 [34,253]	-9.898 [9.584]
Constant	10.33** [0.582]	35.56** [5.919]	47,726 [101,578]	148.7** [42.75]
<i>n</i>	1,243	1,243	1,243	1,243

Notes. This table shows the results of regressing the beliefs of subjects in the control group on demographic variables. Adjusted *p*-values for treatment effects are computed using the Westfall-Young method (** $p < 0.01$, * $p < 0.05$).

Table A5: Effects on confidence

Topic	Living Wage	Taxes	Refugees	Benefits
High interval	-0.0459 [0.0191]	-0.101** [0.0196]	0.00661 [0.0145]	-0.0342 [0.0188]
Constant	0.678** [0.0133]	0.649** [0.0134]	0.151** [0.0101]	0.341** [0.0134]
<i>n</i>	2,475	2,475	2,475	2,475

Notes. This table shows the results of regressing confidence about policy levels on a dummy variable indicating that the subject was shown a high interval. Adjusted p -values for treatment effects are computed using the Westfall-Young method (** $p < 0.01$, * $p < 0.05$).

Table A6: Instrumental variable regressions (with controls)

Topic	Living Wage	Taxes	Refugees	Benefits
<i>Policy increases</i>				
Belief	-0.00378 [0.0133]	0.00274 [0.00222]	1.51e-08 [6.45e-08]	-0.000253 [0.000195]
<i>Policy decreases</i>				
Belief	0.00150 [0.00297]	0.00195 [0.00170]	6.47e-08 [6.58e-08]	4.29e-05 [0.000149]
<i>n</i>	2,475	2,475	2,475	2,475

Notes. This table repeats the analysis of Table 4 after controlling for socio-demographic characteristics. The estimated effects of the controls are suppressed but can be viewed using the replication package. Adjusted p -values for treatment effects are computed using the Westfall-Young method (** $p < 0.01$, * $p < 0.05$).

Table A7: Instrumental variable regressions (full sample)

Topic	Living Wage	Taxes	Refugees	Benefits
<i>Policy increases</i>				
Belief	-0.0100 [0.0142]	0.00290 [0.00232]	-1.55e-08 [7.02e-08]	-0.000267 [0.000210]
Constant	0.884** [0.154]	0.257* [0.102]	0.284** [0.0138]	0.423** [0.0257]
<i>Policy decreases</i>				
Belief	0.00125 [0.00318]	0.00166 [0.00174]	1.11e-07 [7.66e-08]	3.60e-05 [0.000154]
Constant	-0.00471 [0.0345]	0.0854 [0.0766]	0.392** [0.0150]	0.147** [0.0187]
<i>n</i>	2,483	2,483	2,483	2,483

Notes. This table repeats the analysis of Table 4 while including the full sample from the information manipulation treatment. Adjusted p -values for treatment effects are computed using the Westfall-Young method (** $p < 0.01$, * $p < 0.05$).

Table A8: Correlational evidence

Topic	Living Wage	Taxes	Refugees	Benefits
<i>Policy increases</i>				
Belief (high interval)	-0.0392* 0.00997	-0.00234 0.00139	-7.07E-08* 1.94E-08	-0.000115 0.0000566
Belief (low interval)	-0.000143 -0.00417	-0.00236 0.00140	-9.90E-07 6.36E-07	0.00112 0.000750
<i>Policy decreases</i>				
Belief (high interval)	0.00276 0.00296	0.000850 0.00122	4.55E-08 2.85E-08	0.0000735 0.0000555
Belief (low interval)	0.00373 0.00258	-0.00397 0.00117	8.04E-07 5.49E-07	-0.00141 0.000585

Notes. This table reports the results of regressing attitudes on changes on beliefs about levels (along with demographic controls, whose coefficients are suppressed). The regressions are conducted separately within the high interval and low interval groups. Adjusted p -values are computed using the Holm-Šídák method (** $p < 0.01$, * $p < 0.05$).

Table A9: Effects of the qualitative arguments (with controls)

Topic	Living Wage	Taxes	Refugees	Benefits
<i>Policy increases</i>				
Argument	-0.211** [0.0251]	0.183** [0.0254]	0.113** [0.0235]	-0.0382 [0.0251]
Constant	1.147** [0.197]	0.272 [0.232]	0.210 [0.248]	0.561* [0.313]
<i>Policy decreases</i>				
Argument	0.0323** [0.00907]	-0.0685** [0.0153]	-0.0496 [0.0254]	0.0753** [0.0233]
Constant	-0.0736** [0.0321]	0.114* [0.0682]	0.465* [0.252]	0.149 [0.107]
<i>n</i>	1,274	1,274	1,274	1,274

Notes. This table repeats the analysis of Table 5 after controlling for socio-demographic characteristics. The estimated effects of the controls are suppressed but can be viewed using the replication package. Adjusted p -values are computed using the Holm-Šídák method (** $p < 0.01$, * $p < 0.05$).

Table A10: Beliefs about actual and ideal policy levels

Topic	Living Wage	Taxes	Refugees	Benefits
Belief	0.924** [0.071]	0.716** [0.035]	0.636 [0.248]	1.161 [0.377]
Constant	3.265** [0.762]	12.710** [1.484]	48,394* [18,792]	24.059 [21.277]
R^2	0.245	0.354	0.08	0.014
<i>n</i>	1243	1243	1243	1243

Notes. This table regresses individual beliefs about actual policy levels on their preferred policy levels using responses from the control group. Adjusted p -values for treatment effects are computed using the Westfall-Young method (** $p < 0.01$, * $p < 0.05$).

APPENDIX B. SURVEY (FOR ONLINE PUBLICATION)

B.1 INFORMATION MANIPULATION TREATMENT

Consent block

Screen 1

Overview. The purpose of this study is to better understand the political views of voters in the UK. This study has received ethics approval from the Institutional Review Board of the Paris School of Economics.

Duration. The survey should take between 4 and 6 minutes to complete.

Payment. You will be paid around £9/hour on average, assuming that you complete the survey, pass the attention checks, and stay on this browser tab. In addition, you may be paid a bonus for answering questions correctly.

Anonymity. We will never ask for your name, and you can be sure that your responses will remain completely anonymous.

Contact. In case of any issues, please contact the Principal Investigator via the Prolific platform. If he fails to resolve your issue, please contact PSE's Institutional Review Board at: [redacted]

Do you consent to take part in this survey?

- Yes — I would like to take part in this study, and confirm that I am an adult living in the UK.
- No — I would not like to take part in this study.

Welcome block 1

Screen 2

What is your Prolific ID?

Note: your ID should automatically be generated in the text box below.

Screen 3

Welcome to the survey!

We are very grateful that you are participating since the answers that you give will form the basis of our academic research.

Screen 4

In surveys like these, there are occasionally participants who rush through the questions without reading them properly. Unfortunately, such participants can compromise the results of studies by providing essentially random answers. To show that you read questions carefully, please select ‘turquoise’ as the answer to the next question.

What is your favourite colour?

- Blue
- Red
- Purple
- Turquoise

Screen 5

You will be asked to guess the value of various quantities over the course of the survey. For each guess that you make that is within 10 percent of the correct answer, you will receive an additional “lottery ticket” for a £200 bonus. **Therefore, the more accurate your guesses, the more likely you are to win a large bonus payment!**

Screen 6

Please do not leave this browser tab at any point during the survey! For example, please make sure not to open a new browser, open a new browser tab, or change your browser tab. **If you leave this browser tab, you could lose your experimental earnings.**

Policy questions block

[*Note:* the order of policies was randomised within the block.]

Screen 7

We would now like to ask you some questions about the **National Living Wage**.

As you may know, the National Living Wage is the minimum pay per hour that most workers in the UK are entitled to by law. More specifically, nearly all employees aged 23 or older must be paid the National Living Wage for every hour of work they do.

Screen 8a (low interval treatment)

What do you think that the UK's National Living Wage is right now (in pounds per hour)?

Hint: it's between **£5/hour** and **£11/hour**.

Screen 8b (high interval treatment)

What do you think that the UK's National Living Wage is right now (in pounds per hour)?

Hint: it's between **£10/hour** and **£20/hour**.

Screen 9

How confident are you in your answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 10

Do you think that:

- The National Living Wage should be **increased**.
- The National Living Wage should be **decreased**.
- The National Living Wage should be kept just the same.

Screen 11

How confident are you in your answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 12

We would now like to ask you some questions about **income tax**. Income tax is a tax paid on most forms of income that an individual receives (including income from a job or self-employment).

Screen 13a (low interval treatment)

Consider a British tax resident who lives in England and currently earns £160,000 per year.

If that person were to earn an extra £100, what percentage of that £100 would they need to pay in income tax?

Hint: it's between **10%** and **50%**.

Screen 13b (high interval treatment)

Consider a British tax resident who lives in England and currently earns £160,000 per year.

If that person were to earn an extra £100, what percentage of that £100 would they need to pay in income tax?

Hint: it's between **40%** and **90%**.

Screen 14

How confident are you in this answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 15

We have been asking you about how much extra tax an English taxpayer who earns £160,000/year would pay if they receive extra income.

Do you think that:

- This tax rate should be **increased**.
- This tax rate should be **decreased**.
- The tax rate should be kept just the same.

Screen 16

How confident are you in this answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 17

We would now like to ask you some questions about **refugees**, i.e. migrants who have left their country because of fear of persecution.

Screen 18a (low signal)

In the year ending in September 2023, how many refugees do you think that the UK granted refugee permission to following an asylum claim?

Hint: it's between **400 (four hundred)** and **40,000 (forty thousand)**.

Screen 18b (high signal)

In the year ending in September 2023, how many refugees do you think that the UK granted refugee permission to following an asylum claim?

Hint: it's between **30,000 (thirty thousand)** and **3,000,000 (3 million)**.

Screen 19

How confident are you in this answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 20

Do you think that:

- The UK should accept **more** refugees.
- The UK should accept **fewer** refugees.
- The UK should keep accepting the same number of refugees.

Screen 21

How confident are you in this answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 22

We would now like to ask you about **government benefits** that are paid to people who are **unemployed**.

Screen 23a (low interval)

Specifically, please imagine a 25-year-old who is **out of work**. You can assume that they live with their parents, do not have children, do not have a disability or health condition, and do not have substantial savings.

How much do you think this person is eligible to receive every week in unemployment benefits?

Hint: it's between **£10/week** and **£90/week**.

Screen 23b (high interval)

Specifically, please imagine a 25-year-old who is **out of work**. You can assume that they live with their parents, do not have children, do not have a disability or health condition, and do not have substantial savings.

How much do you think this person is eligible to receive every week in unemployment benefits?

Hint: it's between **£80/week** and **£800/week**.

Screen 24

How confident are you in this answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 25

Do you think that:

- The UK should **increase** unemployment benefits.
- The UK should **decrease** unemployment benefits.
- The UK should keep the level of unemployment benefits just the same.

Screen 26

How confident are you in this answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

<i>Demographics block</i>

Screen 27

Finally, we would like to ask you a few questions about your personal circumstances.

What is your age?

What is your gender?

- Male
- Female

- Other

Where do you live?

- England
- Scotland
- Wales
- Northern Ireland

Screen 28

What is the highest level of education that you have completed?

- Secondary school up to 16 years
- Further education (A-levels, BTEC, etc.)
- Undergraduate degree
- Postgraduate degree (e.g. master's, PhD)

Were both of your parents born in the UK?

- Yes
- No

What is your current employment status?

- Full time employee
- Part time employee
- Self-employed
- Unemployed
- Student
- Retired

Screen 29

What was the monthly income of your household, after taxes, on average last year?

- Less than £1,000
- £1,000-£1,500

- £1,500-£2,000
- £2,000-£3,000
- £3,000-£4,000
- More than £4,000

Screen 30

When it comes to economic policy, are your views generally:

- Left-wing
- Centrist
- Right-wing
- Don't know

Did you vote in the 2019 general election?

- Yes
- No

Screen 31 (displayed if they reported having voted)

Which party did you support?

- The Conservative Party
- The Labour Party
- The Scottish National Party
- The Liberal Democrats
- Other

Screen 32

Many thanks for completing the survey! Please click the arrow below to return to Prolific and register your submission.

B.2 QUALITATIVE TREATMENT

<i>Consent block</i>

 \implies see Section [B.1](#)

Welcome block 2

Screen 2

What is your Prolific ID?

Note: your ID should automatically be generated in the text box below.

Screen 3

Welcome to the survey!

We are very grateful that you are participating since the answers that you give will form the basis of our academic research.

Screen 4

In surveys like these, there are occasionally participants who rush through the questions without reading them properly. Unfortunately, such participants can compromise the results of studies by providing essentially random answers. To show that you read questions carefully, please select 'turquoise' as the answer to the next question.

What is your favourite colour?

- Blue
- Red
- Purple
- Turquoise

Screen 5

Please click the arrow to proceed.

Screen 6

Please do not leave this browser tab at any point during the survey! For example, please make sure not to open a new browser, open a new browser tab, or change your browser tab. **If you leave this browser tab, you could lose your experimental earnings.**

Arguments block

[*Note: the order of policies was randomised within the block.*]

Screen 7

We would now like to ask you some questions about the **National Living Wage**.

As you may know, the National Living Wage is the minimum pay per hour that most workers in the UK are entitled to by law. More specifically, nearly all employees aged 23 or older must be paid the National Living Wage for every hour of work they do.

Screen 8a (argument treatment)

Some economists argue that, although minimum wages are intended to help working people, they can actually harm them by destroying their jobs.

Their logic can be illustrated with a simple example:

- Suppose that an employee generates £5 of revenue for every hour that they work.
- Suppose that a minimum wage is introduced which forbids them from working for less than £10/hour.
- Since the minimum cost of hiring the employee (£10/hour) now exceeds the amount of revenue that the employee generates (£5/hour), it is now in the firm's interests to fire the employee.

More generally, these economists argue that a sufficiently high minimum wage will not just destroy jobs but also prevent firms from posting job vacancies in the first place.

Screen 8b (no argument treatment)

Please click on the arrow to proceed.

Screen 9

Do you think that:

- The National Living Wage should be **increased**.
- The National Living Wage should be **decreased**.
- The National Living Wage should be kept just the same.

Screen 10

How confident are you in your answer?

- Very confident
- Quite confident
- Neither confident nor unconfident

- Quite unconfident
- Very unconfident

Screen 11

We would now like to ask you some questions about **income tax**. Income tax is a tax paid on most forms of income that an individual receives (including income from a job or self-employment).

Screen 12a (argument treatment)

If we were to raise taxes on those who earn above £130,000/year, some very rich individuals might need to make some small sacrifices. For example, they might need to buy less expensive wines, or to take fewer exotic holidays.

However, the money raised by these tax increases could be used to help people who are truly struggling and unable to pay for basic necessities like food or rent. This suggests that the benefits of raising taxes on the very rich greatly exceed the costs.

Screen 12b (no argument treatment)

Please click on the arrow to proceed.

Screen 13

Consider a British tax resident who lives in England and currently earns £160,000 per year.

Do you think that:

- This tax rate should be **increased**.
- This tax rate should be **decreased**.
- The tax rate should be kept just the same.

Screen 14

How confident are you in this answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 15

We would now like to ask you some questions about **refugees**, i.e. migrants who have left their country because of fear of persecution.

Screen 16a (argument treatment)

There is currently no legal way for the overwhelming majority of asylum seekers to claim refugee status in the UK.

For example, a gay person trying to escape persecution in Uganda, a person trying to escape civil war in Syria, and a Christian convert trying to escape religious persecution in Pakistan have no way to legally enter the UK for the purposes of making an asylum claim.

For this reason, the only way for such people to apply for asylum is to first enter the UK illegally, which usually requires making a dangerous crossing of the Channel by boat.

Screen 16b (no argument treatment)

Please click on the arrow to proceed.

Screen 17

Do you think that:

- The UK should accept **more** refugees.
- The UK should accept **fewer** refugees.
- The UK should keep accepting the same number of refugees.

Screen 18

How confident are you in this answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 19

We would now like to ask you about **government benefits** that are paid to people who are **unemployed**.

Screen 20a (argument treatment)

Unemployment benefits encourage people not to work: after all, these benefits are stopped (or reduced) the moment that a person starts working.

For this reason, cutting unemployment benefits would encourage some unemployed individuals to return to the workplace.

Screen 20b (no argument treatment)

Please click on the arrow to proceed.

Screen 21

Do you think that:

- The UK should **increase** unemployment benefits.
- The UK should **decrease** unemployment benefits.
- The UK should keep the level of unemployment benefits just the same.

Screen 22

How confident are you in this answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Demographics block ⇒ see Section [B.1](#)

B.3 CONTROL GROUP

Consent block ⇒ see Section [B.1](#)

Welcome block 1 ⇒ see Section [B.1](#)

Ideal points block

[Note: the order of policies was randomised within the block.]

Screen 7

We would now like to ask you some questions about the **National Living Wage**.

As you may know, the National Living Wage is the minimum pay per hour that most workers in the UK are entitled to by law. More specifically, nearly all employees aged 23 or older must be paid the National Living Wage for every hour of work they do.

Screen 8

What do you think that the UK's National Living Wage is right now (in pounds per hour)?

Screen 9

How confident are you in your answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 10

What do you think that the National Living Wage *should* be (in pounds per hour)?

Screen 11

How confident are you in your answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 12

We would now like to ask you some questions about **income tax**. Income tax is a tax paid on most forms of income that an individual receives (including income from a job or self-employment).

Screen 13

Consider a British tax resident who lives in England and currently earns £160,000 per year.

If that person were to earn an extra £100, what percentage of that £100 would they need to pay in income tax?

Screen 14

How confident are you in your answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 15

Consider again a British tax resident who lives in England and currently earns £160,000 per year.

If that person were to earn an extra £100, what percentage of that £100 *should* they pay in income tax?

Screen 16

How confident are you in your answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 17

We would now like to ask you some questions about **refugees**, i.e. migrants who have left their country because of fear of persecution.

Screen 18

In the year ending in September 2023, how many refugees do you think that the UK granted refugee permission to following an asylum claim?

Note: this number does not include those granted permission to remain in the UK under the Ukraine or Afghanistan schemes.

Screen 19

How confident are you in your answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 20

How many refugees do you think that the UK *should* grant refugee permission to each year?

Screen 21

How confident are you in your answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 22

We would now like to ask you about **government benefits** that are paid to people who are **unemployed**.

Screen 23

Specifically, please imagine a 25-year-old who is **out of work**. You can assume that they live with their parents, do not have children, do not have a disability or health condition, and do not have substantial savings.

How much do you think this person is eligible to receive every week in unemployment benefits? (Please answer in £/week.)

Screen 24

How confident are you in your answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Screen 25

Consider again a 25-year-old who is out of work. As before, you can assume that they live with their parents, do not have children, do not have a disability or health condition, and do not have substantial savings.

How much do you think this person *should* be eligible to receive every week in unemployment benefits?

Screen 26

How confident are you in your answer?

- Very confident
- Quite confident
- Neither confident nor unconfident
- Quite unconfident
- Very unconfident

Demographics block \implies see Section [B.1](#)