

Fetzer, Thiemo; Shaw, Callum; Edenhofer, Jacob

Working Paper

Informational Boundaries of the State

CESifo Working Paper, No. 10901

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Fetzer, Thiemo; Shaw, Callum; Edenhofer, Jacob (2024) : Informational Boundaries of the State, CESifo Working Paper, No. 10901, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/295990>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Informational Boundaries of the State

Thiemo Fetzer, Callum Shaw, Jacob Edenhofer

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Informational Boundaries of the State

Abstract

Formal conceptions of state capacity have mostly focused on indirect measures of state capacity – by, for instance, using the state’s fiscal or extractive capacity as a proxy for its overall capacity. Yet, this input or extractive view of state capacity falls short, especially since cross-country empirical evidence suggests that similar levels of fiscal capacity, measured by tax revenues as a percentage of GDP, can produce starkly different outputs – both in classic economic terms and in broader terms that citizens would recognize as desirable outcomes, including quality of life, health, security, equality of opportunity, and intergenerational mobility. This paper argues that a central step towards addressing these shortcomings of the conventional view is to account for a crucial and largely ignored boundary of the state or dimension of state capacity: its capacity to gather, process, and deploy information in its conduct of fiscal policy. Specifically, we study how the presence or lack of such informational capacity constrains governments in responding to crises, such as the recent energy price shock. Our framework provides the analytical toolkit to examine how the informational boundary of the state shapes the incentives for policymakers to resort to untargeted and/or distortionary policy instruments, as opposed to targeted and non-distortionary ones, in responding to crises. The policy response to the energy crisis following the invasion of Ukraine provides the empirical context upon which we bring this theoretical framework to bear on data, though the latter can be straightforwardly extended to other recent crises.

JEL-Codes: H110, O430, D630, D730, Q480, P160, C210, C550.

Keywords: state capacity, economic development, carbon taxation, political economy, pork-barrel politics.

Thiemo Fetzer

University of Warwick / United Kingdom

t.fetzer@warwick.ac.uk

Callum Shaw

*London School of Economics and Political
Science, London / United Kingdom*

c.j.shaw@lse.ac.uk

Jacob Edenhofer

University of Oxford / United Kingdom

jacob.edenhofer@some.ox.ac.uk

January 16, 2024

Fetzer is based at University of Warwick & Bonn and affiliated with CEPR, CAGE, NIESR, ECONtribute, Grantham Institute. Shaw is based at the London School of Economics and Political Science. Edenhofer is based at the University of Oxford. We thank Freddie Didier for excellent RA work and Peter Lambert for helpful comments throughout this research project. This is part of the wider Climate Crisis research agenda that can be followed on: <https://trfetzer.com/climate-crisis-research>. This work has benefited from support through the Leverhulme Trust Prize in Economics in 2022, the European Research Council Starting Grant (ERC, MEGEO, 101042703) and an ESRC Doctoral Training Partnership Studentship.

1 Introduction

Governments around the world have had to contend with a number of shocks and crises in the past three decades or so – ranging from the global financial crisis, the European debt and subsequent austerity crises, the 2015 migration crisis, (populist) uncertainty and de-globalization shocks, the COVID-19 crisis, and, most recently, the energy crisis. The ability of democratic political systems to respond effectively to these (often foreseeable) crises, such as climate change, has been increasingly called into question ([Stasavage, 2020](#); [Mittiga, 2022](#)). Many point to democracies' lack of state capacity or governments' inability to optimally deploy their existing capacity – which we refer to as lacking *performative* state capacity – as key reasons for concerns about their potential to produce good outcomes ([Schularick, 2021](#)).

This paper studies an important dimension of (performative) state capacity: namely, a state's ability to gather, analyze, and ultimately leverage policy-relevant information when designing and implementing (unconventional) fiscal policy, particularly in response to crises. Specifically, we examine the significance of this *informational* dimension of state capacity for demarcating the *boundaries of the state* – the set of shocks governments can effectively deal with if they choose to do so – in the context of the 2022 energy crisis that many governments had to grapple with in the wake of the Russian invasion of Ukraine.

We do so by contrasting two starkly different policy responses to the crisis, the British and German ones, through the lens of British data. Our analysis of the UK policy response demonstrates a democratic government's failure to optimally leverage its informational resources. In contrast, the German response – especially when compared against the backdrop of UK data – was much less distortionary, though its design and implementation were severely constrained by Germany's low informational capacity, particularly its lack of granular, high-quality administrative data. As a result, the German government had to rely on a rather coarsely targeted transfer scheme, whereas in the UK such data were available but not utilized. Given that Germany and the UK have, broadly speaking, similar energy demand profiles, comparing these two policy responses helps us elucidate key political economy

mechanisms that can drive suboptimal policy responses.

Different factors can limit the development or optimal use of the state’s informational capacity. A lack of readily available high-quality data, the complexity of crises, notably their heterogeneous impacts, and an absence of in-house statistical and data analysis expertise can prevent governments from responding to crises with the necessary agility and speed.¹ In the absence of such capacity, the resulting policy responses are likely to provide poor value for public money. Mistaking a failure of policy implementation for a failure of policy direction can then create incentives to pursue ‘public-service-focused’ austerity, thus further eroding the state’s agility to engage in effective real-time problem solving (Fetzer, 2019a; Fetzer and Feld, 2023; Hoddinott et al., 2022).²

Whatever the exact reasons for the lack of informational capacity in many otherwise high-capacity democracies, our paper focuses on the costs of insufficient informational capacity as well as the political dynamics underlying the under-utilization of such capacity. Specifically, we make four contributions. First, we provide descriptive cross-country and England-focused evidence on key design features of energy support policies. The former suggests that, across the board, countries responded to the energy crisis by deploying support measures that were both untargeted and distortionary. The latter suggests that the UK’s policy response, in addition to being regressive and inefficient, benefited (some) supporters of the incumbent party, the Conservatives, disproportionately.

Second, drawing on this evidence, we develop a theoretical framework to shed light on the ways in which informational capacity shapes policymakers’ incentives to respond to an energy crisis. We conceptualize this problem as a policymaker allocating a given budget for energy support measures between targeted lump-sum transfers and untargeted subsidies. In addition, consumers are heterogeneous and differ in their income levels, while the policymaker observes only a noisy signal of

¹For Covid-19-related examples, see e.g. Fetzer and Graeber, 2020; Fetzer, 2021a.

²Duque et al. (2024) suggest that austerity, lengthy administrative processes, and skill deficits – exacerbated by demography – often impede the public sector’s ability to adopt cost-saving technology. In fact, low-quality data and low skill levels can produce errors that erode trust, in part due to a media multiplier driven by availability heuristics (Besley et al., 2020).

true consumer types. The variance of the signal reflects the informational capacity of the state. Our central theoretical result demonstrates that only a policymaker who attaches greater weight to the welfare of high-income, as opposed to low-income, types opts for the untargeted policy instrument and does so only when the informational capacity at her disposal is low.

Third, we leverage granular data to juxtapose the British and German policy responses. The reason for comparing these two countries is that their policy responses differed starkly, among others, in the extent to which they preserved the signal value of prices. To demonstrate this point, we evaluate the actual policy choices – relative to a large vector of fiscally neutral and potentially superior alternatives, which would, however, have required a higher degree of informational state capacity. Our simulations demonstrate that the lack of informational capacity resulted in substantial inefficiency and that particularly (very) high-income Conservative supporters disproportionately benefited from this.

Fourth, guided by our theoretical framework, we attempt to rationalize the observed policy choices. The analysis suggests that a lack of informational state capacity – combined with political economy considerations arising from policymakers attaching different welfare weights to different groups in society – was likely an important factor in shaping policy choices in response to the energy price shocks. Therefore, we argue that sharp informational boundaries between the state, households, and firms, as well as the resulting stickiness in the flow of data between economic actors, are crucial and hitherto under-researched factors explaining the observed policy choices.³

Studying the effects of informational capacity (or lack thereof) and its sluggish deployment is of first-order importance for several reasons. First, the fiscal volume of support measures introduced in the wake of the energy crisis is nothing short of gigantic ([Arregui et al., 2022](#); [Sgaravatti et al., 2023](#)): across Europe, €768 billion has been earmarked to help consumers cope with rising energy prices, with German

³Low levels of trust in institutions, which have become particularly evident over the past decade or so in the Western world and have been tapped into by outside actors, may be another salient factor that undermines governments' ability to respond effectively ([Algan et al., 2017](#); [Besley and Dray, 2022](#); [Besley et al., 2023](#)).

and British energy support policies costing roughly €265 billion and £103 billion, respectively.⁴ Second, the support measures implied, in many cases, an expansion of fossil fuel subsidies, which, as of 2022, were roughly \$7 trillion globally (Black et al., 2023).⁵ This is especially jarring, given that these ‘dirty’ subsidies must be phased out quickly to avert the worst of the climate crisis.

Third, the underlying *design* of the support measures likely matters for their economic and political efficacy. Therefore, insufficiently effective measures could further undermine societal cohesion, at least to the extent that societal cohesion is a function of the (perceived) efficacy of crisis response. The British and German interventions stand in rather stark contrast in this respect.

Finally, unlike other sudden and unexpected crises, such as the COVID-19 pandemic, which brought to the fore significant differences in countries’ approaches to protecting lives and livelihoods and, more broadly, to distributing the economic burden of policies (Kaplan et al., 2020), the 2022 energy crisis, just like the climate crisis – was predictable well in advance. Consequently, it is possible to evaluate policy choices against alternative policies that were available to policymakers *at the time*, rather than with the benefit of hindsight.

Situating our contribution in the literature This paper is related to several strands of literature on state capacity and policy evaluation in economics and political science. First, our research is related to the literature on the determinants and boundaries of state capacity.⁶ Besley and Persson (2009, 2010, 2011) develop a theoretical framework for analyzing how different types of states develop, depending on rulers’ coercive power and societies’ cohesiveness. In their model, rulers can invest in both fiscal and legal capacity, which, in turn, affect the level of public good

⁴This is equivalent to approximately 7% of Germany GDP and 4.6% of UK GDP in 2021.

⁵Some hydrocarbon-rich economies managed to implement energy price reforms in a more equitable fashion by leveraging information capacity to address their distributional ramifications via lump-sum transfers (Aldubyan and Gasim, 2021).

⁶In exploring the informational constraints on policymaking, it builds on works examining how states developed the capacity to collect taxes via centralized state bureaucracies as well as legal systems. See, e.g. Levi, 1989; Tilly, 1993; Spruyt, 2002; Dixit, 2004, 2010; Fukuyama, 2011; Mann, 2012; Beramendi et al., 2019; Grzymala-Busse, 2020; Sánchez de la Sierra, 2020; Dahlström and Lapuente, 2022; Li et al., 2022; Garfias and Sellars, 2021, 2023; Albers et al., 2023.

provision.

We loosely draw on and extend this framework by considering an additional dimension of state capacity, namely informational capacity. The literature on state capacity – both in economics and political science – is strongly historically oriented, that is, it aims to explain the emergence of centralized states (Wang, 2021). As a result, this literature analyzes the informational capacity of states largely in non-digital form. In his seminal work, anthropologist and political scientist, James C. Scott (1999), coined the term *legibility* to capture the extent to which governments could *read* their citizens through the introduction of, for instance, cadastres. While Scott (1999) provides qualitative case-study evidence for the legibility, recent work has extended and modified Scott’s work both empirically and theoretically (Lee and Zhang, 2017; Ansell and Lindvall, 2020; Brambor et al., 2020; Bowles, 2023; Martin, 2023; Garifas and Sellars, 2023).

What little quantitative work exists on the informational dimension of state capacity focuses mainly on the use of data and information to develop the extractive capacity of states, their ability to collect taxes (Pomeranz, 2015; Brockmeyer et al., 2019; Naritomi, 2019; Slemrod, 2019; Slemrod and Keen, 2021; Weigel, 2020; Balan et al., 2022; Bergeron et al., 2023). Our paper, by contrast, focuses on the state’s role in gathering and analyzing (granular) data to deliver cost-effective and targeted fiscal interventions. This is particularly important in light of the rapid evolution of ICT and AI technologies⁷ and their potential impact on the ability of governments to implement such interventions.⁸ The role that information and data play in our analysis highlights the importance of debates around privacy, ‘digital’ property rights, and information governance. Poorly defined or enforced property rights relating to individuals’ data can undermine their ‘production’ or use (Posner and

⁷See Margalit and Raviv (2023) for experimental evidence on the willingness of bureaucrats to employ AI tools.

⁸A notable area in the literature discusses the means of modern autocrats or ‘spin dictators’, who use information heavily to both control and repress citizens (Weidmann and Rød, 2019; Guriev and Treisman, 2019, 2020, 2022; Dimitrov, 2022; Beraja et al., 2023). Narratives or stories can be used to engineer behavioral change, but they can also exacerbate or shape economic shocks (Besley et al., 2020) or can be used as hybrid weapons in a service sector trade escalation. See also Tirole (2021) for an analysis of the opportunities for social control these new technologies open up, even in democratic societies.

Weyl, 2018; Jones and Tonetti, 2020; Acemoglu et al., 2022; Fetzer, 2022b), leading to an underproduction of knowledge public goods, while likely resulting in rents in the form, for example, of the growth and concentration of superstar firms in some ICT sectors (Autor et al., 2020; Cunningham et al., 2021). The negative externalities of inadequately defined and enforced digital property rights are exacerbated by the public sector’s sluggishness in leveraging administrative and other data⁹ as well as deploying cutting-edge data science tools to design long-term policies and short-term responses to crises.¹⁰

On the methodological side, our paper relates to works that use micro- and macro-simulation methods for policy evaluation while seeking to account for heterogeneity between households. In macroeconomics, for instance, the growing popularity of heterogeneous-agent New Keynesian (HANK) models exemplifies this approach (Kaplan et al., 2018; Sargent, 2023).¹¹ In applied work, micro- and macro-simulation models can be fruitfully applied when rich household-level data is available and/or there is no plausibly exogenous variation in policy. This paper complements this research by (i) zeroing in on the importance of granular data, data science and statistical skills, and well-defined property rights over data in shaping (more) effective policy, and (ii) adding an implicit (spatial) political economy dimension to this literature.¹²

Finally, by presenting a political economy rationalization of some of the most important design features of countries’ responses to the energy price shock, the pa-

⁹This is sometimes because of legal barriers concerning the use of highly granular data, conflicts of interest, or low-quality data. Globally, there are two diametrically opposing views on the role that data or information should play in shaping welfare.

¹⁰High levels of privacy around personal data may facilitate tax evasion (Alstadsæter et al., 2019), which, even in the absence of actual evasion, can nevertheless shape narratives around growing inequalities (Piketty and Saez, 2003; Auten and Splinter, 2023) with low transparency and high distrust. This is especially true when people’s everyday environments worsen in ways they find difficult to comprehend or accept, or when they attribute these negative changes to the consequence of low (performative) state capacity (Fetzer, 2020, 2021b).

¹¹See Auclert et al. (2023); Bayer et al. (2023) for studies of the energy price shock in that tradition.

¹²These ‘technological’ dimensions strongly influence state capacity and policy outcomes. Athey and Wager (2021) examine how policymakers can optimally adjust policies when only observational data is available to them. Acemoglu et al. (2022) suggest that data oversharing can mean that it is welfare-enhancing to shut down data markets. Jones and Tonetti (2020) show how the non-rivalrous nature of data can provide a justification for granting users property rights over their data, which chimes well with the informally derived recommendations by Posner and Weyl (2018).

per speaks to the literature on populism and, more broadly, zero-sum politics (see, e.g., [Fetzer, 2019b](#); [Chinoy et al., 2023](#))¹³ in two ways. First, our empirical evidence suggests that the UK’s response to the energy crisis via the *Energy Price Guarantee* not only disproportionately benefited a small electoral group, but also did so in a regressive way. In addition to that, this policy also engendered significant negative social and economic externalities.¹⁴ Second, our findings are suggestive of an "unholy" electoral coalition between high- and low-income households. The high-income political right effectively designs self-serving, distortionary, and fiscally costly energy subsidies, such as a price cap that lower-income households may support because it may appeal to their perceptions of fairness¹⁵, or, if they are subject to their own mental models or cognitive constraints, they may find it especially easy to understand.¹⁶

The remainder of the paper proceeds by first presenting some motivating evidence and outlining the context in more detail. Next, we set out our conceptual framework to illuminate the trade-offs between equity, efficiency, and informational capacity. Drawing on that framework, the fourth section employs counterfactual policy simulations to bring this framework to bear on fine-grained, household- and individual-level data from the UK. The fifth section concludes.

2 Context, data and motivating evidence

We first present descriptive evidence on energy subsidies, as well as the relationship between the incidence of these subsidies and the voting behavior or political leanings of their winners and losers. Through this evidence we motivate our theoretical framework and interpret the results of our simulations of counterfactual policies.

¹³See [Fetzer \(2019a\)](#) and [Fetzer et al. \(2023b\)](#) for some examples.

¹⁴In the context of the energy crisis, [Fetzer \(2023a\)](#) documents how untargeted energy subsidies have likely caused a sharp increase in crime, in particular in areas with energy-inefficient homes.

¹⁵For a discussion relating to universal basic income, see e.g. [Ghatak and Maniquet \(2019\)](#).

¹⁶Theoretically, [Gabaix \(2020\)](#) introduces bounded rationality into a New Keynesian model. Alternative mechanisms could be: biased perceptions of within- and between-group inequality ([Hvidberg et al., 2023](#)) or a media multiplication channel, whereby tragic, albeit isolated, individual stories are amplified in a highly polarized (social) media ecosystem exacerbated by limited statistical literacy ([Besley et al., 2020](#)).

2.1 Time-series and cross-country evidence

Energy prices – in particular those of natural gas – increased sharply following both the COVID-19 pandemic and the Russian invasion of Ukraine. As Figure 1 bears out, natural gas prices started to increase in mid-2021 and then rose dramatically in early 2022. Although they have declined markedly since late 2022, natural gas prices are still more than 100% higher than their average during the period from 2019 to early 2022. This ‘hysteresis’ in global natural gas prices, viz. the move to a high-price regime, is demonstrated by fitting a Markov regime-switching model to the time series of gas prices.

(Figure 1)

Many governments, especially in Europe, responded to the surge in energy prices by putting in place vast support schemes to help households cope with the economic ramifications of higher energy prices. Classical welfare economics – reflected in the advice of international bodies, such as the International Monetary Fund (Arregui et al., 2022) – suggests that government support should (i) be targeted to households most in need to prevent regressive distributional effects, (ii) preserve price signals to incentivize energy-saving consumption behavior both in the short and long term, and (iii) be phased out after a period of time deemed suitable for households to make adjustments to their energy consumption by, for instance, investing in better insulation or acquiring heat pumps. Such first-best support policies might, however, not be feasible by virtue of, for example, insufficient administrative capacity to implement targeted transfers or lack of political will.

(Figure 2)

Figure 2 suggests that most of the transfer schemes implemented across Europe did not meet the (normative) standards of classical welfare economics, with many schemes being both untargeted and distortionary. The figure uses transformed data from Arregui et al. (2022), which measures the share of a country’s total expenditure that is untargeted and distortionary. The scatterplot shows that there is a strong

positive correlation between the extent to which energy support measures are untargeted and distortionary. In sections 3 and 4, we shed light on the ways in which trade-offs between efficiency and equity – a key political economy dimension in representative democracies – and informational capacity can shape policymakers’ best responses to energy price shocks, potentially leading them to deviate from the first best by opting for untargeted and distortionary interventions.

2.2 The UK and German responses: A tale of two energy support schemes

Figure 2 shows that there is substantial variation in the design of countries’ policy responses, particularly the degree to which they were distortionary. To understand this variation better, it is useful to compare the German and UK responses in somewhat more detail.

The German response to the energy price shock was devised by the *Gaskommission*, which was set up in late September 2022 by the government and comprised of representatives of labor unions and business organizations as well as academics.¹⁷ Given the informational and data availability constraints, the commission recommended implementing lump-sum transfers that crucially aimed to preserve the signal function of prices (*ExpertInnen-Kommission Gas und Wärme*, 2022). Specifically, 80% of household energy consumption, measured by consumption in the previous year, was subsidized. For the remaining 20% of energy consumption, households faced market prices. To avoid distorting societal and/or political preferences over equity across households, the lump-sum transfer was passed through the income tax system. In this way, the policy response proposed by the *Gaskommission* was close to that of an economically literate and benevolent social planner. Yet, implementing this proposal has proved challenging as the data necessary to deliver such a targeted transfer scheme were not – and still are not – available to policymakers. In fact, the German state, like many other governments around the world, has

¹⁷This reflects the techno-corporatist nature of (large-scale) policymaking in Germany, as is also borne out by the *coal commission*, the commission that designed Germany’s coal phase-out (*Furnaro et al.*, 2021).

no way of transferring funds to citizens directly, as, for instance, the recent debate about the recycling of revenues from carbon pricing illustrates (Konopka, 2022).¹⁸

In contrast, the British policy response, the *Energy Price Guarantee* (EPG), was to impose a cap on energy prices. The government effectively lowered prices by up to 50% relative to (expected) market prices. As a result, 100% of energy consumption was not subject to market prices, blunting the incentive for households to reduce their energy consumption (Fetzer, 2022a). In addition to distorting relative prices, the transfers made via the EPG were also entirely untargeted. One reason – albeit by no means the only or most important one, as we will discuss below – was that, similar to Germany, the UK government has no way of making individualized transfer payments to households and linking or indexing these transfers to individual energy consumption.

Given the scale of fiscal interventions and the fact that (unconventional) fiscal policy¹⁹ is often criticized on account of its potential distributional effects, it is vital to study how informational capacity shaped the design of policy responses and their distributional effects.

2.3 Within-country evidence

We focus on the UK to demonstrate the plausibility of these distributional considerations for the political economy of the adoption of untargeted and distortionary energy support schemes across countries. To this end, we leverage granular data on political preferences – measured by ward-level local election vote shares between 2008 and 2019. Furthermore, we construct a long-term average of Conservative Party vote share for around 7000 electoral wards across elections for local council-

¹⁸This stands in stark contrast to some other countries that have, in recent years, implemented energy price reforms – the distributional ramifications of which were addressed by direct household transfers (IMF, 2023).

¹⁹Seidl and Seyrich (2023) show how, given certain conditions, fiscal and monetary policy are perfect substitutes within a heterogeneous-agent New-Keynesian (HANK) framework Kaplan et al. (2018). Gong (2023) proposes a distributional decomposition framework for HANK models. Bachmann et al. (2021) provide empirical evidence for the likely progressive effects of VAT cuts.

lors.²⁰ By averaging across elections, we net out candidate- or election-level idiosyncratic factors. Therefore, our measure reflects longer-term political preferences and captures the *average* support of the Conservative Party’s core constituency across granular geographical units. We combine this measure of electoral support for the Conservatives with detailed, granular data on domestic energy consumption at the postcode level.²¹ The latter provides the median, mean, and total energy consumption, whose source is either electricity or gas, for all households living in an area. As with the election data, we remove idiosyncratic time factors by averaging the moments across the different years for which data is available.

Empirical specification Using the above data, we estimate a simple linear regression, capturing the extent to which energy consumption is structurally different across electoral wards with higher core support for the Conservative Party:

$$y_i = \alpha_{r(i)} + \beta \times x_{c(i)} + \nu \times p_w + \epsilon_d$$

We consider two different dependent variables, y_i . First, we use a measure of long-term mean or median energy consumption. A positive point estimate, $\nu > 0$, would indicate the extent to which energy consumption is higher in areas with a higher structural level of electoral support for the Conservatives. The second measure is a proxy for inequality in energy consumption: we construct the difference between median and mean energy consumption. In areas where mean energy consumption is significantly higher than median consumption, it is likely that the distribution of energy consumption across individual households is skewed to the right, with some households having very high consumption.

²⁰We focus on the Conservatives since they controlled government in 2022, and the UK’s majoritarian political system entails that, unlike in proportional democracies (e.g. Germany or Switzerland), control of governmental offices translates into almost uninhibited control over policy (Powell, 2000; McGann and Latner, 2013; Russell and Gover, 2017). Hence, the Conservatives were solely responsible for the design of *Energy Price Guarantee*.

²¹See here for the raw data: <https://www.gov.uk/government/publications/postcode-level-domestic-gas-and-electricity-consumption-about-the-data>.

Energy consumption, (local) consumption inequality and political preferences

The results from estimating the above regression are presented visually in the form of binned scatter plots in Figure 3.²² We focus on natural gas consumption since the latter is the predominant fuel used for heating across the UK (Stewart, 2023). Panel A in Figure 3 shows that higher levels of median or mean household-level natural gas consumption are associated with structurally higher levels of electoral support for the Conservatives across local elections. The implication is that Conservative Party voters likely benefit more from untargeted energy price subsidies in absolute terms, which is likely because of their higher energy consumption footprint.

(Figure 3)

Panel B of Figure 3 uses an alternative moment of the distribution of energy consumption to further examine the latent political economy dimension. In generating Panel B of Figure 3, we follow exactly the same approach as above, except that here the independent variable is – not the median (or mean) of energy consumption – but rather the difference between the mean and the median. This difference can be construed as a proxy for the skew in the distribution of energy consumption. If the mean household’s energy consumption is (significantly) higher than the median’s consumption, this suggests that there are some households with disproportionately high levels of energy consumption. Conversely, if the mean is lower than the median, there are likely some lower-income households in rather wealthy, high-energy-consumption areas in which some homes are possibly only partially occupied.

We find that electoral support for the Conservative Party is indeed predicted by this measure of local inequality in energy consumption. Support for the Tories is notably higher in areas with greater inequality in energy consumption. Such spatial inequality may be attributable to the effects of spatial institutions (with deep historical roots), as analyzed in great detail in Fetzer (2023b).²³ The notable V-shaped pattern between natural gas consumption skew and Conservative Party

²²Tabular results are presented in Table A1 in the appendix.

²³See also Koster (2024).

support is consistent with a pattern of policymaking that has been observed in the UK and some other Western countries, in particular those with majoritarian two-party systems, and has recently been referred to as zero-sum politics ([Drutman, 2019](#); [Fetzer, 2019b](#); [Gidron et al., 2020](#); [Boxell et al., 2022](#); [Hahm et al., 2023](#); [Chinoy et al., 2023](#); [Burn-Murdoch, 2023a](#)).

2.4 Individual-level data

To further substantiate these findings and work towards a more nuanced understanding of the political economy channel, we next show that the results obtained using aggregate data are robust. To that end, we use individual-level micro-data from the *Understanding Society* (USOC) study to document that there is a positive correlation between fuel and energy use, on the one hand, and support for the Conservative Party, on the other. ²⁴

Our dependent variable is the energy bills faced by consumers, which are self-reported in pounds per year at the household level h . Our key independent variable seeks to capture individuals' political preferences, specifically their support for the Conservatives. The USOC contains a question asking individuals to indicate which political party they feel closest to. The independent variable is then a dummy variable, which is equal to unity when individuals state that they feel closer to the Conservatives than to other political parties. For individuals who do not feel particularly close to any party, a broader measure of support for the Conservatives can be constructed by considering whether they would vote for the Conservatives if a general election were held tomorrow.

We estimate a range of specifications to explore the extent to which households of Conservative Party supporters are subject to higher energy bills compared to their peers. The specification follows very closely the one estimated above.

$$e_{h,i} = \alpha_{r(i)} + \beta \times x_{c(i)} + \nu \times p_i + \epsilon_d$$

²⁴The data was used previously in [Fetzer, 2019a](#); [Fetzer et al., 2023b](#) and is described in more detail there.

The results are presented in Table 1. On average, Conservative Party supporters – depending on the degree of saturation of the empirical specification – have annual energy bills that are between two to six percentage points higher than those of comparable individuals who do not support the Conservatives. Unsurprisingly, however, this result masks considerable heterogeneity.

(Table 1)

Figure 4 explores this heterogeneity by studying household income. We classify households into percentiles based on the empirical distribution of their self-reported income and then measure to what extent energy bills are (notably) higher among Conservative Party supporters, relative to the others in the same household income percentile. We categorize household income into 25 different bins, with each bin representing four percent of the total mass of survey participants. The resulting 25 point estimates are plotted in Figure 4, with a solid dot indicating that an estimated coefficient is statistically significant at the 5% level.

Consistent with the simple linear regression, we find that energy bills among those that lean Conservative are notably higher. Yet, this relationship is mostly flat across the household income distribution and increases markedly only for those in the top 5% of the distribution.²⁵

(Figure 4)

The inverted L-shaped relationship suggests that, in particular, high-income households that lean Conservative benefited disproportionately from untargeted energy price subsidies.²⁶ Figure A4 in the appendix shows results from an almost identical exercise, save for the fact that we express our dependent variable, energy bills, as a share of household income. For (almost) no specification do we

²⁵This chimes well with the findings relating to the relative skew in energy consumption (obtained using granular aggregated energy consumption data) that are presented in Figure 3.

²⁶This graph can be construed as showing *Piketty in space*; it shows that the hockey-stick-type pattern in the share of national income accruing to the richest – which [Piketty and Saez \(2003\)](#) document and [Auten and Splinter \(2023\)](#) dispute to some extent – is also present spatially, at least in the UK.

observe any significant differential between Conservative and non-Conservative-leaning individuals. Across the whole income distribution, the share of energy bills among Conservative-leaning individuals is thus similar to the share among non-Conservative-leaning ones. There is one sharp, though imprecisely estimated, exception: for Conservative-leaning individuals in households belonging to the lowest 5% of the income distribution, energy bills account for a notably higher share of monthly income. This suggests that, among the lowest-income households that lean Conservative, there are some households for whom energy costs are a disproportionately high share of monthly income. Note, however, that this relationship is not statistically robust.²⁷

Overall, the analysis suggests that Conservative-leaning high-income households benefit disproportionately from any transfer scheme that is not targeted on the basis of, for instance, household income or need. Given that, it is not particularly surprising that the Conservative UK government devised an untargeted energy support scheme that required minimal informational capacity and, at least nominally, *treated* all households equally. We posit that the simplicity of an untargeted price subsidy – which, in absolute terms, is more beneficial to households with particularly high levels of energy consumption – might have helped garner (relatively) broad popular support. That is, the Tories managed to implement an inefficient and distortionary policy with little public backlash by building an *unholy coalition* between high-income Conservative-leaning households – who benefited materially from untargeted price subsidies – and those who supported the EPG on account of its *simplicity* and/or greater perceived *fairness*.

²⁷In a fragmented and highly polarized (social) media landscape, with a large share of individuals subject to cognitive constraints, availability heuristics, or (endogenous) filter bubbles, individual narratives or stories can produce widespread (political) multiplier effects (Besley et al., 2020). In this way, distracting narratives can be generated that political opponents can exploit, thereby imposing a constraint on policymakers. Graeber et al. (2023) present evidence on the relative importance or stickiness of stories over statistics in human cognition.

3 Conceptual framework

In this section, we formally model the policymaker's problem in designing a policy response to the energy crisis. Consumers are heterogeneous in their income and, therefore, in their demand for energy, which is a normal good.²⁸ The policymaker has to split a fixed, exogenously determined budget between lump-sum transfers and a subsidy program. Furthermore, the policymaker has exogenous preferences over the transfer that each consumer type receives and observes each consumer's type only imperfectly. This provides a tractable framework for analyzing the trade-offs between efficiency, distributional objectives, and the benefit of higher informational capacity, and how these trade-offs are affected by policymakers' preferences, which themselves derive from the political constraints they face.

Consumer demand There are two consumer types, $\theta_i \in \{H, L\}$. Consumers have income m_θ (with $m_H > m_L$) and can purchase two goods: energy (good x) and another good (good y), representing all other consumption. The price of good y is normalized to 1; the price of energy is p^{-1} before the energy crisis, and rises to $p^0 > p^{-1}$ as a result of the crisis.

Each consumer receives support from the government in the aftermath of the crisis via a mix of lump-sum transfers and a price subsidy. We measure the utility consumers derive from the policy response in money-metric terms by calculating their equivalent variation: the transfer of wealth to a given consumer that would be required, at prices $(p^0, 1)$, for her to achieve the same utility as the level she achieves from the mix of a lump-sum transfer and price subsidy that the government actually implements.

Let g denote the lump-sum transfer given to a particular consumer, and let g_s be the government's total expenditure on the subsidy program. In the appendix, we derive the following expression for the "equivalent transfer" of a consumer with type θ , which we denote U_θ :

²⁸Figure A3 in the appendix presents empirical evidence that demonstrates the plausibility of this assumption.

$$U_\theta = g + f_\theta(g_s)$$

With $f'_H(g_s) + f'_L(g_s) \leq 1 \ \forall g_s$, $f_H(g_s) > f_L(g_s) \ \forall g_s$, $f'_H(g_s) > f'_L(g_s) \ \forall g_s$, $f_H(g_s) < 1 \ \forall g_s$, $f_H(0) = f_L(0) = 0$, $f'_H(0) + f'_L(0) = 1$, $f'_H(0) > \frac{1}{2}$, $f'_L(0) < \frac{1}{2}$, and $f'_H(g_s) + f'_L(g_s) < 1$ for $g_s > 0$.

U_θ , $f(\cdot)$ and the properties of $f(\cdot)$ are fully derived in the appendix; here, we simply explain the intuition.

Higher expenditure on lump-sum transfers to any consumer increases the consumer's equivalent variation in a one-to-one fashion, which is why g enters the function U_θ linearly. In contrast, the total government expenditure on subsidies, g_s , enters the function U_θ through the "subsidy benefit" function $f_\theta(\cdot)$, which is heterogeneous across the two consumer types.

The total benefit accruing from a pound spent on the subsidy program is always less than or equal to 1 ($f'_H(g_s) + f'_L(g_s) \leq 1 \ \forall g_s$). This is because of the **deadweight loss of commodity subsidization**, which is the subsidy analogue of the more familiar concept of deadweight loss of commodity taxation. In short, an additional pound (from the policymaker's point of view) of subsidy is worth less than a pound to the consumer because the subsidy distorts prices and prevents the consumer from correctly optimizing her consumption against the true price vector, which reflects the underlying supply.

A pound spent on the subsidy program always benefits the high-income consumer more than the low-income consumer ($f'_H(g_s) > f'_L(g_s) \ \forall g_s$ and $f_H(g_s) > f_L(g_s) \ \forall g_s$) because high-income consumers, all else equal, have a higher energy demand than their low-income counterparts.

The first pound spent on energy consumption provides a benefit equal to a pound ($f'_H(0) + f'_L(0) = 1$), since, at $g_s = 0$, the marginal deadweight loss of commodity taxation is zero. For all $g_s > 0$, the deadweight loss of commodity taxation is strictly positive; so $f'_H(g_s) + f'_L(g_s) < 1$ for $g_s > 0$.

$f'_H(0) > \frac{1}{2}$ and $f'_L(0) < \frac{1}{2}$ follow from $f'_H(0) + f'_L(0) = 1$ and $f_H(g_s) > f_L(g_s) \ \forall g_s$, and make explicit the fact that the first pound spent on the subsidy program (which

will result in a total consumer benefit equal to 1) benefits the high-income consumer by more than 50p and the low-income consumer by less.

$f_H(0) = f_L(0) = 0$ is trivially true.

The policymaker's problem There are two consumers in the population, indexed by i . The vector of consumer types (θ_1, θ_2) can take one of two values, (H, L) or (L, H) . The policymaker does not observe the actual types of consumers, but only a noisy vector of signals (ω_1, ω_2) . This vector can take either of two values, (h, l) or (l, h) . That is, the policymaker correctly believes that one consumer is a high type and the other a low type, while only receiving a noisy signal as to which one is which,²⁹ with:

$$\begin{aligned} Pr\left((\theta_1, \theta_2) = (H, L) \mid (\omega_1, \omega_2) = (h, l)\right) &= Pr\left((\theta_1, \theta_2) = (L, H) \mid (\omega_1, \omega_2) = (l, h)\right) = \beta \\ Pr\left((\theta_1, \theta_2) = (L, H) \mid (\omega_1, \omega_2) = (h, l)\right) &= Pr\left((\theta_1, \theta_2) = (H, L) \mid (\omega_1, \omega_2) = (l, h)\right) = 1 - \beta \end{aligned}$$

β , with $\beta \in [0.5, 1]$, can be interpreted as the informational capacity of the state, which we consider exogenous.

The policymaker seeks to maximize the sum of the expectation of some concave function of consumer benefit U_θ subject to an exogenously determined budget G . Therefore, the policymaker solves the following problem:

$$\begin{aligned} \max_{g_s, g_\omega \forall \omega} \quad & \sum_i^2 \mathbb{E}_{\theta_i} \left[\Delta_{\theta_i} c(U_{\theta_i}) \mid \omega_i \right] \\ \text{subject to} \quad & g_s + \sum_i^2 g_i \leq G \end{aligned}$$

Here g_s is the expenditure on the subsidy program, g_i is the lump-sum transfer given to consumer i and G is the exogenously determined budget. Δ_{θ_i} is the weight

²⁹We model the problem as one involving two consumers of different types so as to make the proofs clearer. This approach is analytically identical to a model with N consumers, divided evenly between high and low types, where the policymaker receives a different binary signal for each consumer.

the policymaker places on the utility of consumer with type θ_i (with $\Delta_H + \Delta_L = 1$), and $c(\cdot)$ is some function which satisfies $c'(\cdot) > 0$ and $c''(\cdot) < 0$.³⁰

The policymaker can only condition the lump-sum transfers g_i on the signal she receives. So, we can define g^l and g^h as the transfer given to the consumer emitting the l signal, and the transfer given to the consumer emitting the h signal, respectively. Using the probabilities defined above and the definition of U_{θ_i} , and given that the budget constraint will bind, the problem can thus be rewritten as:

$$\begin{aligned} \max_{g_s, g^h, g^l} \quad & \left(\beta \Delta_H c(g^h + f_H(g_s)) + (1 - \beta) \Delta_L c(g^h + f_L(g_s)) \right. \\ & \left. + \beta \Delta_L c(g^l + f_L(g_s)) + (1 - \beta) \Delta_H c(g^l + f_H(g_s)) \right) \\ \text{subject to} \quad & g_s + g^l + g^h = G \end{aligned}$$

Solving the policymaker's maximization problem yields the following three first-order conditions (derived in full in the appendix):

The uncertain lump-sum redistribution condition

$$\begin{aligned} & \beta \left(\Delta_H c'(g^h + f_H(g_s)) - \Delta_L c'(g^l + f_L(g_s)) \right) \\ & = (1 - \beta) \left(\Delta_H c'(g^l + f_H(g_s)) - \Delta_L c'(g^h + f_L(g_s)) \right) \end{aligned} \tag{1}$$

This condition describes the policymaker's incentive to distribute lump-sum transfers in such a way as to disproportionately benefit her most favored consumer type, while making clear that the policymaker is constrained in attempting to do

³⁰As income is exogenous in this model, it is appropriate to consider the policymaker as having concave preferences over the *equivalent transfer* each consumer receives, rather than having linear preferences over each individual's concave utility function, as is common in analyses of a social planner's maximization problem, where wealth and income are normally endogenous. A caveat worth mentioning is that, naturally, an *equivalent transfer* to a lower-income consumer increases utility by more than an equal *equivalent transfer* to a higher-income consumer. This consideration is reflected here in the relative values of the Δ parameters. Additionally, the concavity of $c(\cdot)$ generates risk aversion in the policymaker's decisions, which drives some important results in the model.

this by her uncertainty about consumers' true types. This mechanism can be elucidated by comparing the situation, where $\beta = 1$, to one where $\beta < 1$. If we were to assume $\beta = 1$, the condition reduces to a "certain" lump-sum redistribution condition:

$$\frac{\Delta_H}{\Delta_L} = \frac{c'(g^l + f_L(g_s))}{c'(g^h + f_H(g_s))} \quad (2)$$

That is, under certainty about consumer types, the policymaker uses lump-sum transfers to exactly achieve her preferred distributive outcome, whatever the values of $f_L(g_s)$ and $f_H(g_s)$. In an uncertain world, however, $\beta < 1$ and the right-hand side (RHS) of equation (1) becomes relevant. When equation (2) holds and provided that $\Delta_H \neq \Delta_L$, we get:

$$\frac{\Delta_H}{\Delta_L} \neq \frac{c'(g^h + f_L(g_s))}{c'(g^l + f_H(g_s))}$$

As a result, the policymaker has to deviate from the "preferred" condition of (2) in order to satisfy (1). This makes explicit that the policymaker has to shade her lump-sum transfers closer to each other than she would optimally like in order to insure against the risk that she has incorrectly identified the consumer's type. We refer to this effect as the "uncertainty cost of redistributing via lump-sum".³¹

The subsidy balance conditions:

The high-type transfer subsidy balance condition

$$\begin{aligned} & f'_H(g_s)\beta\Delta_H c'(g^h + f_H(g_s)) + f'_L(g_s)(1 - \beta)\Delta_L c'(g^h + f_L(g_s)) \\ & + f'_L(g_s)\beta\Delta_L c'(g^l + f_L(g_s)) + f'_H(g_s)(1 - \beta)\Delta_H c'(g^l + f_H(g_s)) \\ & = \beta\Delta_H c'(g^h + f_H(g_s)) + (1 - \beta)\Delta_L c'(g^h + f_L(g_s)) \end{aligned} \quad (3)$$

³¹Consider, for instance, a policymaker with $\Delta_H > \Delta_L$. Such a policymaker would like to set lump-sum transfers such that $g^h > g^l$. Under uncertainty, however, she has to consider the world where she mistakes one type of consumer for another, and so will shade the lump-sum transfers closer together than she would otherwise like in order to insure against the possibility that it is, in fact, the H types who end up with the g^l transfer and vice versa.

The low-type transfer subsidy balance condition

$$\begin{aligned}
& f'_H(g_s)\beta\Delta_H c'(g^h + f_H(g_s)) + f'_L(g_s)(1 - \beta)\Delta_L c'(g^h + f_L(g_s)) \\
& + f'_L(g_s)\beta\Delta_L c'(g^l + f_L(g_s)) + f'_H(g_s)(1 - \beta)\Delta_H c'(g^l + f_H(g_s)) \\
& = \beta\Delta_L c'(g^l + f_L(g_s)) + (1 - \beta)\Delta_H c'(g^l + f_H(g_s))
\end{aligned} \tag{4}$$

(3) equates the marginal benefit to the policymaker of a pound spent on the subsidy and a pound spent on the lump-sum transfer to the high type, and (4) does the same for the lump-sum transfer to the low type. Note, the marginal benefit of subsidy expenditure is one that accrues by raising the equivalent variation of both consumer types in both states of the world.

Proposition 1. *When there is no uncertainty over consumer types, the policymaker will not spend any budget on the subsidy program.*

Proof. Let $\beta = 1$. Then, (3) becomes:

$$\begin{aligned}
& f'_H(g_s)\Delta_H c'(g^h + f_H(g_s)) + f'_L(g_s)\Delta_L c'(g^l + f_L(g_s)) \\
& = \Delta_H c'(g^h + f_H(g_s))
\end{aligned}$$

and (4) becomes

$$\begin{aligned}
& f'_H(g_s)\Delta_H c'(g^h + f_H(g_s)) + f'_L(g_s)\Delta_L c'(g^l + f_L(g_s)) \\
& = \Delta_L c'(g^l + f_L(g_s))
\end{aligned}$$

Summing these two equations yields

$$\begin{aligned}
& 2\left(f'_H(g_s)\Delta_H c'(g^h + f_H(g_s)) + f'_L(g_s)\Delta_L c'(g^l + f_L(g_s))\right) \\
& = \Delta_H c'(g^h + f_H(g_s)) + \Delta_L c'(g^l + f_L(g_s))
\end{aligned} \tag{5}$$

As pointed out above, when $\beta = 1$, the uncertain lump-sum redistribution con-

dition reduces to the certain lump-sum redistribution condition:

$$\frac{\Delta_H}{\Delta_L} = \frac{c'(g^l + f_L(g_s))}{c'(g^h + f_H(g_s))} \quad (6)$$

Combining (5) and (6) implies:

$$\begin{aligned} & 2\left(f'_H(g_s)\Delta_H c'(g^h + f_H(g_s)) + f'_L(g_s)\Delta_H c'(g^h + f_H(g_s))\right) \\ & = 2\Delta_H c'(g^h + f_H(g_s)) \end{aligned} \quad (7)$$

Which simplifies to

$$f'_H(g_s) + f'_L(g_s) = 1$$

The only value of g_s for which the above condition holds is $g_s = 0$, which, in turn, follows from the properties of $f_\theta(\cdot)$. \square

This result is a manifestation of the deadweight loss of commodity taxation. If the policymaker has full information, she can achieve her redistributive objectives solely through lump-sum transfers and thus only stands to lose by introducing distortionary subsidies. In this special case, $g_s = 0$ and lump-sum transfers are pinned down by:

$$\frac{\Delta_H}{\Delta_L} = \frac{c'(g^l)}{c'(g^h)} \quad (8)$$

Proposition 2. *Under uncertainty, the policymaker will spend some budget on the subsidy program only if $\Delta_H > \Delta_L$.*

Proposition 2 is made clear by the following argument. Consider a policymaker with $\Delta_L > \Delta_H$ under a situation of uncertainty ($\beta < 1$). Such a policymaker will balance her expenditure on lump-sum transfers g^l and g^h to satisfy her uncertain

lump-sum redistribution condition (1). This condition can be rearranged to:

$$\frac{\Delta_H}{\Delta_L} = \frac{\beta c'(g^l + f_L(g_s)) - (1 - \beta)c'(g^h + f_L(g_s))}{\beta c'(g^h + f_H(g_s)) - (1 - \beta)c'(g^l + f_H(g_s))} \quad (9)$$

Suppose $g_s = 0$ and, therefore, that the policymaker only uses the policy instruments g^h and g^l . As we saw above, with complete certainty ($\beta = 1$), the policymaker would want to achieve a low $c'(g^l + f_L(g_s))$ and a high $c'(g^h + f_H(g_s))$ by setting g_l high and g_h low. The presence of uncertainty ($1 - \beta > 0$), however, means that the policymaker must consider the second terms of the numerator and the denominator, respectively. A high g_l results in a low $c'(g^l + f_H(g_s))$ and a low g_h results in a high $c'(g^h + f_L(g_s))$, which prevents the policymaker from satisfying the uncertain lump-sum redistribution condition by setting g_l and g_h as far apart as she would have done under full certainty. She, instead, has to bring the two closer together. This is simply a restatement of the "uncertainty cost of redistribution via lump-sum" described above.

Does such a policymaker have anything to gain by increasing g_s and diverting funds away from the lump-sum expenditures g^l and g^h ? The marginal gain for the policymaker by increasing g_s is given by

$$\begin{aligned} & f'_H(g_s)\beta\Delta_H c'(g^h + f_H(g_s)) + f'_L(g_s)(1 - \beta)\Delta_L c'(g^h + f_L(g_s)) \\ & + f'_L(g_s)\beta\Delta_L c'(g^l + f_L(g_s)) + f'_H(g_s)(1 - \beta)\Delta_H c'(g^l + f_H(g_s)) \end{aligned}$$

This shows that increasing the subsidy increases the equivalent variation of both consumers in both states of the world, whilst increasing the equivalent variation of the H consumer by more than that of the L consumer (by $f'_H(g_s) > f'_L(g_s) \forall g_s$). Diverting funds away from the lump-sum transfers towards subsidies also imposes the deadweight loss of commodity subsidization. Recall that $f_L(g_s) + f_H(g_s) \leq 1 \forall g_s$ and $f_L(g_s) + f_H(g_s) < 1$ for $g_s > 0$. Because $f_H(g_s) > f_L(g_s) \forall g_s$, this means that, for a policymaker with $\Delta_L > \Delta_H$, increasing g_s above 0 can never be optimal. Even under maximum uncertainty ($\beta = (1 - \beta) = \frac{1}{2}$), a pound spent on g^l will always make the policymaker better off than a pound spent on the subsidy. While

there is (at most) a $\frac{1}{2}$ chance that the pound spent on g_l will be misdirected to the H type, rather than the L type, this is still better than the $< \frac{1}{2}$ benefit that goes to the L type if the pound is instead spent on the subsidy program, not to mention the deadweight loss entailed by this policy the moment that $g_s > 0$.

The same argument, however, does not hold for a policymaker with $\Delta_H > \Delta_L$. For such a policymaker, under maximum uncertainty, it is still the case that a pound spent on g^h could have as much as a $\frac{1}{2}$ chance of being misdirected to the L type. Yet, a pound spent on subsidy benefits this policymaker's favored type (H) by an amount greater than $\frac{1}{2}$. This means that it can, in fact, be optimal for such a policymaker to incur the deadweight loss of commodity subsidization, given that doing so better helps her achieve her distributional goals.

This phenomenon can be illustrated by more closely inspecting (9). By increasing expenditure on g_s above 0 and decreasing the "risky" expenditure on g^h and g^l , the policymaker can satisfy (9) via control of the $f_L(g_s)$ and $f_H(g_s)$ parts of each $c(\cdot)$ function, rather than having to depend on controlling the g^l and g^h parts. That is, policymakers attaching greater weights to high-income consumers than low-income ones ($\Delta_H > \Delta_L$) can, to some extent, mitigate the "uncertainty cost of redistribution via lump-sum" by means of the subsidy lever. This is not the case for policymakers favoring low-income consumers ($\Delta_L > \Delta_H$) because the subsidy level only allows for redistribution toward the consumer with higher energy consumption.

Corollary 1. *Under maximum uncertainty, a policymaker with $\Delta_H > \Delta_L$ is better off than a policymaker with $\Delta_L > \Delta_H$.*

Proof. See appendix. □

The mechanism described in Proposition 2 shows that, given certain conditions, the harm resulting from low informational capacity depends crucially on the policymaker's distributional preferences. While we consider β exogenous in our framework, it could easily be extended to make explicit an additional corollary of the proposition: that the incentive to invest in informational capacity is a function of the policymaker's distributional preferences. In the specific example of energy cost support, we generate the result that a policymaker who favors

high-energy-consumption/high-income consumers (as a result of, say, pork-barrel incentives and a correlation between their voter base and high energy expenditure) has a lower incentive to invest in informational capacity than a policymaker with the opposite preferences. This is because that type of policymaker can use a policy instrument that allows her to avoid the costs of uncertainty associated with attempting to target lump-sum transfers and achieve her distributional aims in a way that the policymaker with opposite distributional preferences cannot. It is, moreover, straightforward to see that this low incentive to invest in informational capacity gives rise to a higher deadweight loss of commodity subsidization, which, in turn, is because the alternative policy instrument is itself inefficient.

4 Counterfactual policy scenarios

We next describe how we map the theoretical framework to granular (synthetic) data on millions of individual households, drawing on simulation techniques at scale. These techniques effectively allow us to evaluate the actually implemented policies – focusing in particular on the British and German examples as two almost diametrically opposite policy choices (see section 2). We perform an ex-ante evaluation of the policies’ characteristics and compare them to a broad bundle of ex-ante fiscally equivalent policies that could have been implemented, while implicitly varying β across the exercises.

For the simulation exercise, we hold a broad set of dimensions constant. We assume an exogenous market price p_m , while taking the (stylized) policy paths chosen by the UK and Germany as given. The UK’s policy response, the *Energy Price Guarantee*, the EPG, effectively set a subsidized price p_s , determined by $p_s = \tau p_m$, with $\tau = 55\%$ (Fetzer, 2022a; DESNZ, 2023). As a result, consumers faced energy prices that were 45% lower than market prices. The equivalent of the German policy response can be thought of as effectively introducing a subsidized energy consumption quota at the individual household level.³² With this in mind, we now discuss how we operationalize the key features of our model in the simulations.

³²See Appendix A for more discussion.

4.1 Mapping model to simulation

Types of households In the model, we conceive of the policymaker as choosing between an untargeted energy price subsidy at a financial cost, g_s , and lump-sum transfers targeted at two types of consumers at a financial cost of g^l and g^h , with a given total budget, $g_s + g^l + g^h = G$. Instead of assuming only two types of consumers, as in the model, we allow for greater heterogeneity in consumer types in the simulations, as measured by their level of energy consumption, q_i . We leverage both household-level energy consumption data (q_i) on all properties introduced in [Fetzer et al. \(2023a\)](#) and individual-level panel survey data. For each counterfactual energy support policy, we can measure the correlation between the bills faced by consumers – that is, households’ likely (imputed) energy bills minus the subsidy received – and, for example, their political preferences. This is similar to the exercises that were presented in section 2 as we, in essence, vary β , the government’s informational capacity. This is because different levels of disaggregation correspond to more or less fine-grained ways of assigning energy consumption quotas, with more fine-grained assignments implying greater informational capacity and vice versa.

Policy space In the feasible set of policy bundles, we simulate an untargeted price subsidy, g_s , as implemented via the *Energy Price Guarantee*, and a broad menu of different two-tier tariffs, in combination with lump-sum transfers. In the two-tier tariff structure, a consumer faces the subsidized, i.e. lower, price p_s for a quota of q_m units of energy consumption. Consumers face market prices p_m for any energy consumption above that quota.³³ We do not allow prices to vary and set p_m equal to the market price that the energy regulator would have allowed energy suppliers to charge consumers had there not been an intervention. In determining p_s , we follow the price that the energy regulators set in the year prior to the energy price shock. The alternative subsidy schemes are simulated so that the total energy bill that households face is, on average, held constant. In simulating alternative policies, we allow the total subsidy expenditure to vary within 20% above or below the actual

³³See Figure A5 in the appendix.

implemented policy in both the UK and the UK equivalent of the German proposal.

Before we set out how we operationalize the other components of our model in the simulations, it is worth dwelling on the way in which we simulate two-tier tariffs and lump-sum transfers. Two-tier tariffs can be thought of as a combination of a blanket subsidy and a lump-sum transfer. That is, they operate identically to lump-sum transfers for individuals whose actual consumption is higher than the subsidized quota.³⁴ However, the transfer to any consumer whose consumption is less than the quota will be less than under a lump-sum transfer schedule with the same ex-ante cost. Hence, the two-tier tariff policy scenarios allow the policymaker to resort to a range of policies that are less extreme than a full shift to lump-sum transfer payments.

Note, too, that in the simulations, we consider exposure to market prices as a proxy for the deadweight loss of commodity taxation, rather than attempting to measure the latter explicitly. All else equal, a policy that exposes a consumer to market prices on a greater share of their consumption is superior to one that reduces the market-exposed share of consumption. This is because of the deadweight loss associated with commodity taxation. While in the model the deadweight loss is explicitly recognized via the properties of the function $f(\cdot)$, in the simulations we acknowledge this by measuring the number of consumers who face market price at the margin.

Capturing informational capacity In our model, the parameter β captures the government's informational capacity; we map this into the simulation as an inference problem. With $\beta = 1$, a policymaker observes the actual level of household energy consumption q_i . A two-tier tariff structure would then imply that each household has an individualized quota, $q_{m_i} = s \times q_i$, where s captures the fraction of consumption that is subsidized. The household would face a subsidized price on the first $\min\{q_i, q_{m_i}\}$ units of consumption and the market price on the remaining $\max\{q_i - q_{m_i}, 0\}$ units of energy consumption.

³⁴Consumers face market prices on the marginal unit, as with a lump-sum transfer, and the total transfer to the consumer is fiscally equivalent.

For $\beta < 1$, we consider different ways of *estimating* the level of $q_{m|x_i} = s \times E(q_i|x_i)$, where informational capacity is captured by the degree to which a government is able to produce accurate estimates of household-level energy consumption using publicly available data. Crucially, these data were also available to the government and, more broadly, the public sector when designing and implementing the response to the energy price shock. Using data on (q_i, x_i) , we can achieve two objectives. First, we can simulate variously complex ways of estimating $E(q_i|x_i)$. This is because expanding the size (dimensionality) of the vector x_i means that we take into account more characteristics, allowing more granular targeting. As a result, a higher-dimensional x_i can be considered the empirical analogue of the theoretical construct of increasing β . Second, we can evaluate the performance of policy alternatives – such as blocked two-tier tariffs or blocked lump-sum transfers – relative to an untargeted energy subsidy, the policy that was actually implemented.

We simulate a broad range of alternative policies that vary in the degree to which they leverage different (household) characteristics x_i in estimating consumers' energy consumption. We assume that each consumer's energy consumption is equal to their consumption in the previous year, $q_{i,0} = q_{i,-1}$. Consequently, the policymaker's inference problem is to estimate each consumer's period -1 consumption.

A consumer i has period -1 energy consumption $q_{i,-1} = E(q_{i,-1}|x_i) + \epsilon_{i,-1}$, where $\epsilon_{i,-1} \sim N(0, \sigma)$. Naturally, σ will vary with the dimension of the vector³⁵ x_i . That is, the average error in the policymaker's estimate of a consumer's energy consumption is lower when more information is used to estimate that energy consumption. As a result, σ can be construed as the empirical analogue of β in our model, though the parameters are inversely related. When the policymaker conditions her decision on a larger number of household characteristics – that is, when the dimension of the vector x_i increases – σ decreases, which is analogous to increasing β . In the extreme case, where $x_i = \emptyset$, σ is equal to the variance of energy consumption in the population ($\beta = 0.5$). In the opposite extreme case, where $q_{i,-1} \in x_i$, $\sigma = 0$ ($\beta = 1$).

The policymaker can then condition the policy response to consumer i on $E(q_{i,-1}|x_i)$.

³⁵The characteristics we condition on in the simulations are described in the appendix.

In the case of a lump-sum transfer, this involves a payment to the consumer of an amount $L_i = s \times E(q_{i,-1}|x_i)p_m$, for some s . In the case of a two-tier tariff, it is the quota of subsidized energy consumption that the consumer is granted that is conditioned on $E(q_{i,-1}|x_i)$ – i.e., the policymaker will subsidize the first $q_{m,i} = t \times E(q_{i,-1}|x_i)$ units of energy consumption, for some t .

Capturing distributional aims Let us next turn to the interpretation of our model’s distributional parameters: Δ_H and Δ_L . One would expect the standard policymaker type to be one with $\Delta_H < \Delta_L$. Assuming policymakers to be *progressive* also seems consistent with the stated objective of the response to the energy crisis, which was to prevent the poorest in society from experiencing significant hardship.³⁶ In our context, however, the theoretical results that are of most interest are those that compare such a prioritarian policymaker (Adler, 2019; Adler and Norheim, eds, 2022) to one for whom $\Delta_H > \Delta_L$.

In the simulations, we assume that the Conservative Party has a set of pork-barrel objectives it seeks to pursue via its policy decision – namely, it aims to achieve some positive correlation between the transfer a household receives and the probability of that household supporting the Conservatives. The descriptive evidence discussed in section 2 – in particular the positive correlation between higher levels of energy consumption and Conservative Party support – provides the rationale for treating a policymaker with *conservative* pork-barrel incentives as one for whom $\Delta_H > \Delta_L$.

This discussion hopefully makes clear how, in the model, the balance between Δ_H and Δ_L can, broadly speaking, be considered a reflection of the policy position the policymaker believes will maximize her chances of re-election.

We also evaluate the degree of equity measured as the share of households that would be left better off with the two-tier block tariff vis-à-vis the UK policy and the

³⁶A fair amount of the media coverage of the crisis focused on the risk that some consumers would be “forced to choose between heating and eating” (Viner, 2023). The pre-analysis plan for Fetzer (2023a), posted on <https://osf.io/vhnjz/>, documents how a national food bank charity effectively did not consider it in its strategic interest to share data with the research team to shed light on these questions.

UK equivalent of the German proposal. We consider this to be a measure of the *equity* dimension within our modeling framework.

To evaluate the degree of political targeting, we estimate a regression that captures the degree to which there is a correlation between either an individual's or an area's support for the Conservative Party and the energy bills faced by consumers. The energy bills facing consumers are calculated as the bills that households would face net of each specific transfer schedule, that is, a pair $(s, q_{m|x})$. In essence, this allows us to examine to what extent the different subsidy schedules partition out the partisanship that is indirectly captured through the energy consumption schedule data that were documented in Section 2. We estimate:

$$bill_i^{p_m} - subsidy_{i,q|m_x} = \sum_{\tau}^9 \beta_{\tau} \times conservative_i \times \mathbb{1}_{Q_{\tau}}(y_i) + \sum_{\tau}^9 \xi_{\tau} \times \mathbb{1}_{Q_{\tau}}(y_i) + \epsilon_i$$

where Q_{τ} is the range of incomes between the τ -th and $(\tau + 1)$ -th decile of the income range, and, as such, $\mathbb{1}_{Q_{\tau}}(y_i) = 1$ if y_i is in this range.

When evaluating the properties of different policies, the coefficient estimates β_{10} and the average β , are of particular interest to us. Even when the average effect of $conservative_i$ on the energy bills consumers face (net of transfers) is held constant, it may be the case that different policies affect different parts of the income distribution differently, i.e., entail a different balance between β and β_{10} . Indeed, Figure A6 illustrates the distributional properties of a set of simulated counterfactual policies, relative to the UK's actual policy. The distributional properties shown in Figure A6 lend further evidence to the idea that the implemented policy reflected the preferences of policymakers for high-income consumers.

Capturing the deadweight loss of commodity subsidization In the model, efficiency is captured by the $f(\cdot)$ function, whose properties reflect the fact that a pound spent on subsidy raises total welfare by an amount less than a pound spent on lump-sum transfer. This is because of the deadweight loss entailed by commodity subsidization. In our simulations, we capture efficiency by measuring the

amount of consumption that is subject to market price at the margin.

4.2 Implementation of simulations

To implement the simulations, we require data on a set of features or characteristics, x_i , that policymakers can take into account in designing alternative (targeted) transfer schedules. To identify, within a broad set of transfer schemes, those that have fiscal costs similar to the one that was actually implemented, we carry out a grid search. We briefly describe these here, with more details provided in section C of the appendix.

Feature sets considered We consider a range of features, x_i , which are used to construct estimates of $q_{i,-1} = E(q_{i,-1}|x_i)$. These are then used to simulate alternative energy support measures (lump-sum transfers or two-tier tariffs). The features are, broadly speaking, categorical in nature; they are not only relevant for energy consumption, but are also salient in policymaking in the UK more generally, such as different administrative units, measures of deprivation, local tax bands, and property characteristics. We focus mainly on vectors (x_i) consisting of features that are readily available to policymakers. Policy complexity is ultimately determined by the number of different features x_i that are considered. This policy space grows exponentially in the number of features. Given the full set of characteristics considered here, there are 57,344 different possible combinations, and thus potential support policies. Of these feasible policies, we focus on a (stratified) random sample of 1,593 different combinations of features x_i that are used to estimate $q_{i,-1} = E(q_{i,-1}|x_i)$. The key reason is to reduce computational complexity.

Grid search We perform a grid search across the sample of 1,593 different combinations of features x_i that are used to estimate $q_{i,-1} = E(q_{i,-1}|x_i)$. We do so by varying the generosity of each transfer, with $s \in (0.70, 1.3)$, in the simulations. That is, we allow the amount of consumption that is subsidized, $q_{m|x}$, to vary around the median, where the range of these values is given by: $q_{m|x} = s \times E(q_{i,-1}|x_i)$.

This will help identify transfer schedules that are cheaper to implement than those actually implemented while performing similarly in all other relevant dimensions (equity, efficiency, and political targeting). For the set of 1,593 blocks, we thus allow generosity to vary across a search grid of size 60,³⁷ which means that we compute $1,593 \times 60 = 95,580$ potential transfer schedules.

While we vary the generosity of transfer schedules across the grid search, we retain only those schedules – out all 95,580 ones – that lie within 20% of the aggregate fiscal cost of the EPG, which serves as our benchmark here.³⁸

4.3 Discussion of simulation results

In discussing the central findings yielded by our simulation exercises, we proceed in two steps, first discussing the results obtained using (pseudo) near population-level household data and then turning to those obtained using individual-level survey data. As indicated above, we construct four empirical moments or measures for each counterfactual policy that is roughly fiscally neutral to the EPG, with these moments capturing key political economy dimensions of the UK government’s response to the 2022 energy crisis. Crucially, these moments allow us to interpret the theoretical results through the lens of our model.

4.3.1 Near population data

Univariate characterization We begin with a univariate characterization of the four dimensions – political targeting, efficiency, informational demandingness, and welfare – that are central to our analysis, i.e., we show how the various (counterfactual) two-tier tariff energy support policies that we simulated in Figure 5 fare

³⁷There are 60 steps of size 0.01 between 0.7 and 1.3.

³⁸The simulations using the household-level data are slightly more complicated since the feature set x_i is not fully equivalent and, in some cases, there are missing values. We adopt an equivalent approach for the individual-level data, but the policy space is more constrained. Due to missing data, we allow the total spending to vary marginally. Mechanically, this only arises because the survey data do not contain information for all of the household features x_i in all potential blocks. This results in slightly different samples for each of the transfer schedules.

with respect to these dimensions.³⁹ Panel A examines the degree of political partisanship, measured as long-term support for the Conservatives in local elections, which was used as motivating evidence in Figure 3. This figure plots the empirical distribution of the correlation between the energy bills faced by consumers (net of any subsidy or transfer) and the long-term average Conservative Party vote share across local elections. The vertical lines indicate the correlation generated by the UK’s EPG (its untargeted price subsidy) and the equivalent of the German support scheme, respectively. The latter can be thought of as a *de facto* individualized two-tier tariff (see section 2). We also include a vertical line that indicates the correlation between energy bills and Conservative Party support in the ‘no-intervention’ case.

(Figure 5)

Comparing the ‘no intervention’ and ‘UK/DE’ lines shows that, perhaps unsurprisingly, both the EPG and a German-type tariff drastically weaken the correlation between households’ net energy bills (net of transfers) and Conservative Party support. Yet, the high density in the neighborhood of the ‘UK/DE’ line, as indicated by the kernel density, shows that many of the simulated policies would have achieved a similar degree of *political targeting*, i.e., would have yielded a similar gradient in support for the *average Conservative Party* supporter. Such alternative block tariffs, by utilizing more informational capacity, could have even produced a greater degree of indirect partisan targeting than the actual policy. This finding is in keeping with the logic outlined in Proposition 1: when informational capacity is fully utilized and there is thus no uncertainty about energy consumption or consumer types, distributional aims can be better achieved than they could be using untargeted subsidies.

Panel B of Figure 5 focuses on the distribution of the share of consumption that is subject to market prices under different simulated policies, with this share serving as our proxy for *efficiency*. This figure documents that it would have been possible to design more efficient policies – policies that would have preserved the signal

³⁹We focus on the menu of two-tier tariffs in this discussion. Figure A6 in the appendix presents the equivalent graphs for the lump-sum transfer implementation.

value of prices for a much higher share of consumption. This is true even for the German support scheme, which was explicitly designed to do this ([ExpertInnen-Kommission Gas und Wärme, 2022](#)). Therefore, panel B illustrates a central result of our formal model: fully leveraging informational resources allows policymakers to achieve their distributional aims in ways that do not incur the deadweight loss associated with commodity subsidization.

Given the discussion about privacy and information governance that we alluded to in the introduction, Panel C of Figure 5 shows that there were many alternative policies that would have allowed for *better* targeting without raising concerns about privacy. The figure plots the empirical distribution of the number of consumption blocks that contain fewer than 10 households, as identified by the set of features x_i considered in the simulations. This measures the *informational demandingness* of support schemes, with the latter being high and therefore requiring high informational capacity when the number of blocks with fewer than 10 households is high and vice versa. The shape of the kernel density line suggests that the privacy risk of fully utilizing informational capacity is likely not as severe as feared, at least on the privacy dimension.

Finally, Panel D aims to capture (some of) the *welfare* effects of different support schemes – the extent to which fiscally neutral, more targeted alternative policies would have produced winners and losers, relative to the actually implemented measures. The plot suggests that such alternative policies could easily have made 70% households better off without, as in the British case, incurring the deadweight loss of commodity taxation. This demonstrates that many alternative policies would have at least improved average consumer welfare.

Multivariate characterization We next examine these distributions in a bivariate setting in Figure 6. In this way, we can shed some light on the trade-offs (or their absence) between the four dimensions discussed in the previous section. Panel A highlights that alternative two-tier tariffs could have been implemented that would have achieved the same level of political targeting as the actual policies, as indicated by the horizontal line, while, simultaneously, being more efficient, i.e. a greater

share of consumption would have been subject to market prices. Similarly, Panel B shows that there is *no* trade-off between the share of households that could have been made better off and the degree of political targeting. Panel C, moreover, suggests there is no trade-off between the privacy dimension and the degree of political targeting of the average core Conservative constituency.

(Figure 6)

This analysis suggests that it is hard to rationalize the actual policy responses to the energy crisis, especially the UK's EPG, from a narrow political economy perspective – a perspective solely focused on the average Conservative support gradient. Other support schemes could have been implemented (on the basis of publicly available data) that would have been similarly politically efficient – they would have achieved a similar degree of political targeting – while being more efficient and avoiding the deadweight loss of commodity taxation. These alternative policies would, however, have required a higher degree of informational capacity to be used in the design and implementation of fiscal policy.

4.3.2 Individual-level data

We next discuss the results obtained using survey data. These have the advantage that we can measure both partisan leanings and the underlying income and energy consumption data at the individual level. This allows us to comment on how the highly non-linear relationship between energy consumption and household income among Conservative supporters, documented in Figure 4, can help to further rationalize the observed policy choices. We focus specifically on the correlation between energy bills faced by consumers and income among Conservative-leaning voters.

Univariate characterization The univariate distribution plots are presented in Figure 7. Panel A measures the differential net-of-transfer energy bills among Conservative-leaning voters.⁴⁰ The blue vertical line visualizes the reference point or benchmark,

⁴⁰This is, in essence, a version of the average differential in energy bills among Conservative leaning voters presented in Table 1.

namely the actual schemes that were implemented. The figure reveals that alternative two-tier tariffs or lump-sum transfer schemes, with even greater benefits for the average Conservative supporter, were conceivable. This can be interpreted in two ways. First, policymakers could have resorted to superior (more efficient) measures, even when holding the degree of political targeting roughly constant (implying a low β). Second, and alternatively, they could have implemented policies that were far less partisan. Overall, this suggests that the degree of partisanship was a choice by policymakers (implying $\Delta_H > \Delta_L$).

(Figure 7)

Panel B makes this point even more sharply. It plots the differential in energy bills faced by consumers between Conservative-leaning households in the 10th income decile and non-conservative-leaning households. The figure suggests that fiscally neutral counterfactual policies – while producing an on average similar degree of political targeting across the board (i.e. ignoring income), as shown in Panel A – would have left those Conservative-leaning individuals in the top 10 income decile whose energy consumption is particularly high notably worse off. Their net-of-transfer energy bill differential is notably higher under all counterfactual policies compared to the implemented support measures in Germany and the UK.

Multivariate characterization Figure 7 suggests that Conservative policymakers may indeed have faced a trade-off revolving around the group of voters whose welfare they wanted to improve: providing more targeted energy bill support could have left Conservative-leaning households with (very) high energy consumption markedly worse off, while the average Conservative-leaning household could have been made better off (or at least not worse off). This trade-off is visualized in Figure 8, where we plot the joint distribution. The actually implemented support scheme produces a subsidy that disproportionately benefits the highest earners – and there is no fiscally neutral alternative that can produce a similar outcome. This reinforces the notion that $\Delta_H > \Delta_L$. Not only does the UK’s policy favor Conservative voters, but the extent to which it benefits them is higher for those in the highest income

bracket. This point is supported by Figure 8, which shows that many fiscally neutral policy alternatives achieve a similar benefit to the average Conservative Party supporter. Crucially, however, these alternatives imply significantly less support for the highest income decile than that entailed by the EPG, the UK’s actual policy.

(Figure 8)

Recall that Propositions 1 and 2 can be seen as the key theoretical counterparts of this simulation exercise. Proposition 1 shows that a policymaker would only use subsidy schemes in the absence of sufficiently high informational capacity. A policymaker who more effectively leverages information would, by contrast, have no incentive to incur the efficiency losses associated with such a (distortionary) policy, opting instead for targeted lump-sum transfers. This fits very well with the results of the simulations, which demonstrate that the policymaker can achieve her pork-barrel incentives without having to incur the efficiency losses of subsidy (the deadweight loss of commodity subsidization) if they make full use of their information (i.e. if they have a high β).

Proposition 2 makes explicit that the trade-off between efficiency and distributional objectives exists only for a policymaker with specific pork-barrel incentives, which arise because of their political aims ($\Delta_H > \Delta_L$). To repeat what we stated above: The incentive to provide a transfer that is skewed towards a specific voter base may induce a policymaker to resort to policy instruments that may be inherently less efficient and could reduce overall welfare. Corollary 1, moreover, highlights that pork-barrel incentives and, more generally, political economy dynamics can prevent a government from investing in its informational capacity, despite such investment improving overall welfare. The simulations demonstrate this point powerfully.

5 Conclusion

In this paper, we examine – both conceptually and empirically – one crucial *boundary of the state*, namely its capacity to gather and process information, in the context

of the UK and German responses to the energy price shock in late 2022. Our conceptual framework goes beyond conventional economic reasoning – which focuses on the trade-off between equity and efficiency – by introducing a constraint on the informational capacity policymakers face in the real world.

A key result of our model is that limited informational capacity – which might exist by virtue of stringent data protection laws or a lack of technical ability to gather and process highly granular, high-frequency information – forces policymakers to rely on relatively broad, i.e. poorly targeted or completely untargeted, subsidies in trying to cushion the energy price shock, and that this incentive is particularly strong for policymakers whose core constituency primarily consists of high-income households.

Indeed, our conceptual framework can help us rationalize two almost diametrically opposed responses to the energy price shock – the highly targeted, relatively non-distortionary German response and the UK’s highly distortionary, untargeted one. Our model provides one potentially important explanation for these starkly different responses of two advanced industrial democracies. In the UK, as our simulations further bear out, the government’s *Energy Price Guarantee* is, given some level of (non-perfect) informational capacity, only optimal when assuming a strong preference for high-income over low-income households on the part of policymakers. Since the Conservative’s core constituency is, to a significant extent, drawn from that group ([Burn-Murdoch, 2023b](#)), our theoretical and empirical considerations elucidate the political rationale behind a policy that looks very different from the first-best policy, as suggested by conventional welfare economic reasoning. The German response was much closer to that policy, with deviations reflecting, at least partly, the fact that the German government’s notoriously low informational capacity meant that individualized transfers could not be easily implemented.

Turning from the specific context of our analysis to the broader context, our theoretical results and empirical findings speak to debates about improving governments’ ability to boost state capacity by improving data access, the bureaucracy’s digital literacy – its capacity for analyzing high-frequency data in real time – and coordination between public and private actors. Our analysis demonstrates the costs

of failing to improve governments' informational capacity and, more broadly, its performative state capacity – its ability to deliver policies that achieve its intended objectives and thus minimize negative unintended consequences. Given that the energy crisis is likely to rear its ugly head again this winter and that governments will face other, but structurally similar, crises as countries seek to decarbonize their economies, the ability to design and implement effective relief measures will certainly remain crucial and perhaps even grow in importance.

Finally, our analysis can help inform debates about the best strategies for governments to harness the possibilities opened up by technological change, particularly the emergence of (generative) artificial intelligence, without sacrificing accountability, and without enabling bureaucrats and/or politicians to infringe on civil liberties and political rights in nefarious ways. Putting in place regulation and institutions, or reforming existing ones, to navigate the tension between the embrace of new technological possibilities and the importance of civil liberties, as well as political rights, is of first-order importance for increasing liberal democracy's chances of survival in the 21st century.

References

- Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asu Ozdaglar,** “Too Much Data: Prices and Inefficiencies in Data Markets,” *American Economic Journal: Microeconomics*, November 2022, 14 (4), 218–256.
- Adler, Matthew D.,** *Measuring Social Welfare: An Introduction*, New York: Oxford University Press, October 2019.
- **and Ole F. Norheim, eds,** *Prioritarianism in Practice*, Cambridge: Cambridge University Press, April 2022.
- Albers, Thilo N. H., Morten Jerven, and Marvin Suesse,** “The Fiscal State in Africa: Evidence from a Century of Growth,” *International Organization*, January 2023, 77 (1), 65–101.
- Aldubyan, Mohammad and Anwar Gasim,** “Energy price reform in Saudi Arabia: Modeling the economic and environmental impacts and understanding the demand response,” *Energy Policy*, 2021, 148 (January 2021), 1–36.
- Algan, Yann, Sergei Guriev, Elias Papaioannou, and Evgenia Passari,** “The European Trust Crisis and the Rise of Populism,” *Brookings Papers on Economic Activity*, 2017, pp. 309–382.
- Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman,** “Tax evasion and inequality,” *American Economic Review*, 2019, 109 (6), 2073–2103.
- Ansell, Ben W. and Johannes Lindvall,** *Inward Conquest: The Political Origins of Modern Public Services*, Cambridge: Cambridge University Press, November 2020.
- Arregui, Nicolas, Oya Celasun, Dora Iakova, Aiko Mineshima, Victor Mylonas, Frederik Toscani, Yu Ching Wong, Li Zeng, and Jing Zhou,** “Targeted, Implementable, and Practical Energy Relief Measures for Households in Europe,” December 2022.
- Athey, Susan and Stefan Wager,** “Policy Learning With Observational Data,” *Econometrica*, 2021, 89 (1), 133–161.
- Auclert, Adrien, Hugo Monnery, Matthew Rognlie, and Ludwig Straub,** “Managing an Energy Shock: Fiscal and Monetary Policy,” August 2023.

- Auten, Gerald and David Splinter**, "Income Inequality in the United States: Using Tax Data to Measure Long-Term Trends," *Journal of Political Economy*, nov 2023, 0, 1–67.
- Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen**, "The Fall of the Labor Share and the Rise of Superstar Firms," *The Quarterly Journal of Economics*, may 2020, 135 (2), 645–709.
- Bachmann, Rüdiger, Benjamin Born, Olga Goldfayn-Frank, Georgi Kocharkov, Ralph Luetticke, and Michael Weber**, "A Temporary VAT Cut as Unconventional Fiscal Policy," *SSRN Electronic Journal*, 2021.
- Balan, Pablo, Augustin Bergeron, Gabriel Tourek, and Jonathan L Weigel**, "Local elites as state capacity: How city chiefs use local information to increase tax compliance in the democratic republic of the Congo," *American Economic Review*, 2022, 112 (3), 762–797.
- Bayer, Christian, Alexander Kriwoluzky, Gernot Müller, and Fabian Seyrich**, "Hicks in HANK: Fiscal Responses to an Energy Shock," October 2023.
- Beraja, Martin, Andrew Kao, David Y Yang, and Noam Yuchtman**, "AI-tocracy," *The Quarterly Journal of Economics*, August 2023, 138 (3), 1349–1402.
- Beramendi, Pablo, Mark Dincecco, and Melissa Rogers**, "Intra-Elite Competition and Long-Run Fiscal Development," *The Journal of Politics*, January 2019, 81 (1), 49–65.
- Bergeron, Augustin, Gabriel Z Tourek, and Jonathan L Weigel**, "The state capacity ceiling on tax rates: Evidence from randomized tax abatements in the drc," Technical Report, National Bureau of Economic Research 2023.
- Besley, Timothy and Sacha Dray**, "Trust as state capacity: The political economy of compliance," WIDER Working Paper 2022, UNU-WIDER November 2022.
- **and Torsten Persson**, "The Origins of State Capacity: Property Rights, Taxation, and Politics," *American Economic Review*, September 2009, 99 (4), 1218–1244.
- **and —**, "State Capacity, Conflict, and Development," *Econometrica*, 2010, 78 (1), 1–34.
- **and —**, *Pillars of Prosperity: The Political Economics of Development Clusters*, Prince-

- ton, NJ: Princeton University Press, August 2011.
- , **Anders Jensen, and Torsten Persson**, “Norms, Enforcement, and Tax Evasion,” *The Review of Economics and Statistics*, July 2023, 105 (4), 998–1007.
- , **Thiemo Fetzer, and Hannes Mueller**, “How Big is the Media Multiplier ? Evidence from Dyadic News Data,” *forthcoming, Review of Economics and Statistics*, 2020.
- Black, Simon, Antung A Liu, Ian Parry, and Nate Vernon**, “IMF Fossil Fuel Subsidies Data: 2023 Update,” 2023.
- Bowles, Jeremy**, “Identifying the Rich: Registration, Taxation, and Access to the State in Tanzania,” *American Political Science Review*, May 2023, pp. 1–17.
- Boxell, Levi, Matthew Gentzkow, and Jesse M. Shapiro**, “Cross-Country Trends in Affective Polarization,” *The Review of Economics and Statistics*, January 2022, pp. 1–60.
- Brambor, Thomas, Agustín Goenaga, Johannes Lindvall, and Jan Teorell**, “The Lay of the Land: Information Capacity and the Modern State,” *Comparative Political Studies*, February 2020, 53 (2), 175–213.
- Brockmeyer, Anne, Spencer Smith, Marco Hernandez, and Stewart Kettle**, “Casting a wider tax net: Experimental evidence from Costa Rica,” *American Economic Journal: Economic Policy*, 2019, 11 (3), 55–87.
- Burn-Murdoch, John**, “Are we destined for a zero-sum future?,” *Financial Times*, September 2023.
- , “The Conservatives risk driving a wedge between themselves and the electorate,” *Financial Times*, August 2023.
- Chinoy, Sahil, Nathan Nunn, Sandra Sequeira, and Stefanie Stantcheva**, “Zero-Sum Thinking and the Roots of U.S. Political Divides,” September 2023.
- Cunningham, Colleen, Florian Ederer, and Song Ma**, “Killer Acquisitions,” *Journal of Political Economy*, mar 2021, 129 (3), 649–702.
- Dahlström, Carl and Victor Lapuente**, “Comparative Bureaucratic Politics,” *Annual Review of Political Science*, 2022, 25 (1), 43–63.
- de la Sierra, Raúl Sánchez**, “On the Origins of the State: Stationary Bandits and

- Taxation in Eastern Congo," *Journal of Political Economy*, January 2020, 128 (1), 32–74.
- DESNZ**, "Energy Price Guarantee up until 30 June 2023," September 2023.
- Dimitrov, Martin K.**, *Dictatorship and Information: Authoritarian Regime Resilience in Communist Europe and China*, Oxford: Oxford University Press, December 2022.
- Dixit, Avinash K.**, *Lawlessness and Economics: Alternative Modes of Governance*, Princeton, NJ: Princeton University Press, March 2004.
- , "Democracy, Autocracy and Bureaucracy," *Journal of Globalization and Development*, January 2010, 1 (1).
- Drutman, Lee**, "American politics as a zero-sum game. new america resilience," October 2019.
- Duque, Ricardo Gabriel, Jacob Edenhofer, Thiemo Fetzer, and Sofia Pessoa**, "Origins and Consequences of Austerity," 2024.
- Fetzer, Thiemo**, "Did Austerity Cause Brexit?," *American Economic Review*, November 2019, 109 (11), 3849–86.
- , "Austerity and Brexit," *Intereconomics*, 2020, 55 (1), 27–33.
- , "Measuring the Epidemiological Impact of a False Negative: Evidence from a Natural Experiment," 2021.
- , "On the perceptions of local decline and populism," 2021.
- , "Beyond the Energy Price Guarantee. With or Without?," October 2022.
- , "Discussion of: COVID and Income Inequality," *Economic Policy*, January 2022, 37 (109), 201–203.
- , "Endogenous crime: Evidence from the policy response to the energy crisis in England," 2023.
- , "Regulatory barriers to climate action: Evidence from Conservation Areas in England," *CAGE Working Paper*, 2023.
- **and Immanuel Feld**, "Austerity and Climate (In)Action," *mimeo*, December 2023.
- **and Thomas Graeber**, "Does Contact Tracing Work? Quasi-Experimental Evidence from an Excel Error in England," December 2020. Pages: 2020.12.10.20247080.

- , **Ludovica Gazze, and Menna Bishop**, “Distributional and climate implications of policy responses to the energy crisis: Lessons from the UK,” 2023.
- , **Srinjoy Sen, and Pedro C L Souza**, “Housing Insecurity and Homelessness: Evidence from the United Kingdom,” *Journal of the European Economic Association*, April 2023, 21 (2), 526–559.
- Fetzer, Thomas**, “Nationalism in Political Economy Scholarship,” in T Fetzer and S Berger, eds., *Nationalism and the Economy: Explorations into a Neglected Relationship*, Prague: Central European University Press, 2019, chapter 1.
- Fukuyama, Francis**, *The Origins of Political Order: From Prehuman Times to the French Revolution*, Profile Books, May 2011.
- Furnaro, Andrea, Philipp Herpich, Hanna Brauers, Pao-Yu Oei, Claudia Kemfert, and Wesley Look**, “German just transition: A review of public policies to assist German coal communities in transition,” Technical Report, Resources for the Future 2021.
- Gabaix, Xavier**, “A behavioral new Keynesian model,” *American Economic Review*, 2020, 110 (8), 2271–2327.
- Garfias, Francisco and Emily A. Sellars**, “From Conquest to Centralization: Domestic Conflict and the Transition to Direct Rule,” *The Journal of Politics*, July 2021, 83 (3), 992–1009.
- and —, “State Building in Historical Political Economy,” in Jeffery A. Jenkins and Jared Rubin, eds., *The Oxford Handbook of Historical Political Economy*, Oxford: Oxford University Press, 2023.
- Garifas, Francisco and Emily A. Sellars**, “Fiscal Legibility and State Development: Theory and Evidence from Colonial Mexico,” October 2023.
- Ghatak, Maitreesh and Franccedilois Maniquet**, *Universal Basic Income: Some Theoretical Aspects*, Vol. 11 2019.
- Gidron, Noam, James Adams, and Will Horne**, *American Affective Polarization in Comparative Perspective* Elements in American Politics, Cambridge: Cambridge University Press, November 2020.
- Gong, Zheng**, “Decomposing HANK,” 2023.

- Graeber, Thomas, Florian Zimmermann, and Christopher Roth**, “Stories, Statistics, and Memory,” 2023.
- Grzymala-Busse, Anna**, “Beyond War and Contracts: The Medieval and Religious Roots of the European State,” *Annual Review of Political Science*, 2020, 23 (1), 19–36.
- Guriev, Sergei and Daniel Treisman**, “Informational autocrats,” *Journal of economic perspectives*, 2019, 33 (4), 100–127.
- and —, “A theory of informational autocracy,” *Journal of Public Economics*, June 2020, 186, 104158.
- and —, *Spin Dictators: The Changing Face of Tyranny in the 21st Century*, Princeton; Oxford: Princeton Univers. Press, April 2022.
- Hahm, Hyeonho, David Hilpert, and Thomas König**, “Divided We Unite: The Nature of Partyism and the Role of Coalition Partnership in Europe,” *American Political Science Review*, March 2023, pp. 1–19.
- Hoddinott, Stuart, Matthew Fright, and Thomas Pope**, “‘Austerity’ in public services: Lessons from the 2010s,” Technical Report, Institute for Government October 2022.
- Hvidberg, Kristoffer Balle, Claus Thustrup Kreiner, and Stefanie Stantcheva**, “Social Positions and Fairness Views on Inequality,” *Review of Economic Studies*, nov 2023, 90 (6), 3083–3118.
- IMF**, “Saudi Arabia: 2023 Article IV Consultation-Press Release; Staff Report; and Informational Annex; IMF Country Report No. 23/323; July 7, 2023,” *IMF Staff Report*, 2023, (23).
- Ito, Koichiro**, “Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing,” *American Economic Review*, February 2014, 104 (2), 537–563.
- Jones, Charles I. and Christopher Tonetti**, “Nonrivalry and the Economics of Data,” *American Economic Review*, September 2020, 110 (9), 2819–2858.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante**, “Monetary Policy According to HANK,” *American Economic Review*, March 2018, 108 (3), 697–743.
- , —, and **Giovanni L Violante**, “The Great Lockdown and the Big Stimulus: Trac-

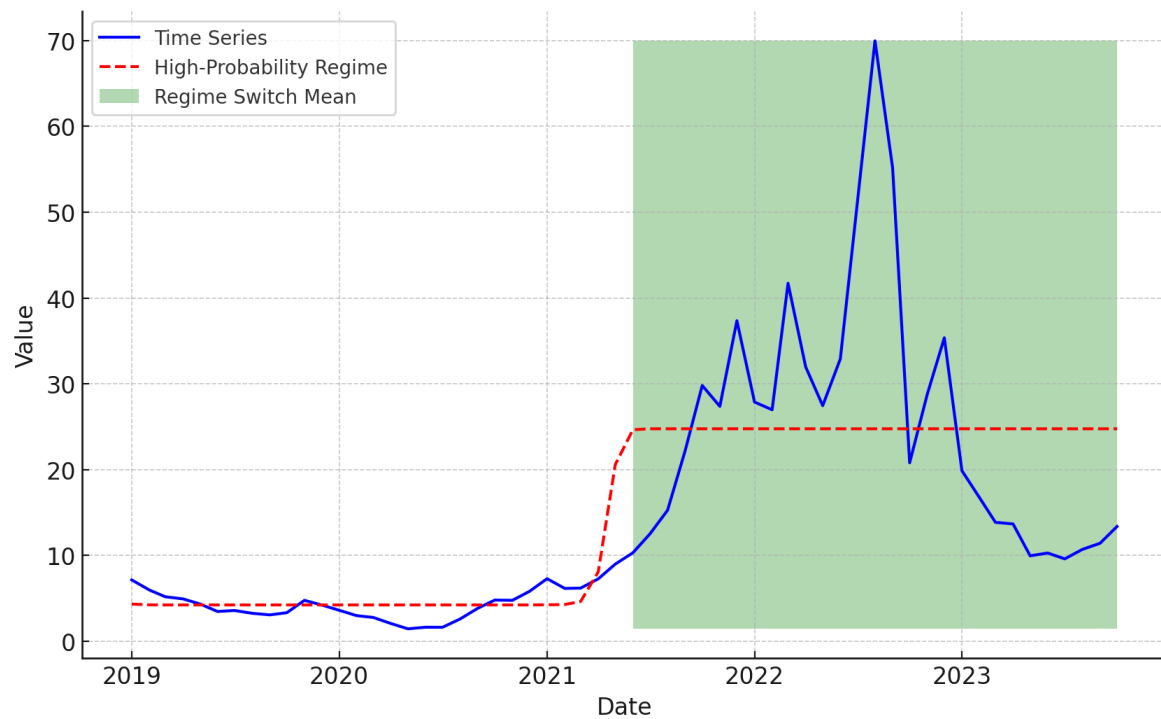
- ing the Pandemic Possibility Frontier for the U.S.," *NBER Working Paper Series*, 2020.
- Konopka, Boris**, "Klimageld in Deutschland," March 2022.
- Koster, Hans RA**, "The Welfare Effects of Greenbelt Policy: Evidence from England," *The Economic Journal*, 2024, 134 (657), 363–401.
- Lee, Melissa M. and Nan Zhang**, "Legibility and the Informational Foundations of State Capacity," *The Journal of Politics*, January 2017, 79 (1), 118–132.
- Levi, Margaret**, *Of Rule and Revenue*, Berkeley: University of California Press, August 1989.
- Li, Weijia, Gérard Roland, and Yang Xie**, "Erosion of State Power, Corruption Control and Fiscal Capacity," *The Economic Journal*, May 2022, 132 (644), 1542–1565.
- Mann, Michael**, *The Sources of Social Power*, Vol. 1, New York: Cambridge University Press, November 2012.
- Margalit, Yotam and Shir Raviv**, "The Politics of Using AI in Public Policy: Experimental Evidence," September 2023.
- Martin, Lucy E. S.**, *Strategic Taxation: Fiscal Capacity and Accountability in African States*, New York, NY: Oxford University Press, May 2023.
- Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R Green et al.**, *Microeconomic Theory*, Oxford: Oxford University Press, 1995.
- McGann, Anthony J and Michael Latner**, "The calculus of consensus democracy: Rethinking patterns of democracy without veto players," *Comparative Political Studies*, 2013, 46 (7), 823–850.
- Mittiga, Ross**, "Political legitimacy, authoritarianism, and climate change," *American Political Science Review*, 2022, 116 (3), 998–1011.
- Naritomi, Joana**, "Consumers as Tax Auditors," *American Economic Review*, sep 2019, 109 (9), 3031–3072.
- Ofgem**, "Energy price cap," 2023.
- Piketty, T. and E. Saez**, "Income Inequality in the United States, 1913-1998," *The Quarterly Journal of Economics*, feb 2003, 118 (1), 1–41.

- Pomeranz, Dina**, “No taxation without information: Deterrence and self-enforcement in the value added tax,” *American Economic Review*, 2015, 105 (8), 2539–2569.
- Posner, Eric A. and E. Glen Weyl**, *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*, Princeton: Princeton University Press, May 2018.
- Powell, G Bingham**, *Elections as instruments of democracy: Majoritarian and proportional visions*, New Haven: Yale University Press, 2000.
- Russell, Meg and Daniel Gover**, *Legislation at Westminster: Parliamentary Actors and Influence in the Making of British law*, Oxford: Oxford University Press, 2017.
- Sargent, Thomas J.**, “HAOK and HANK Models,” July 2023.
- Schularick, Moritz**, *Der entzauberte Staat: Was Deutschland aus der Pandemie lernen muss*, Munich: C.H. Beck, 2021.
- Scott, James C.**, *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*, revised ed., New Haven, CT London: Yale University Press, February 1999.
- Seidl, Hannah and Fabian Seyrich**, “Unconventional Fiscal Policy in a Heterogeneous-Agent New Keynesian Model,” *Journal of Political Economy Macroeconomics*, December 2023, 1 (4), 633–664.
- Sgaravatti, Giovanni, Simone Tagliapietra, Cecilia Trasi, and Georg Zachmann**, “National fiscal policy responses to the energy crisis,” September 2023.
- Slemrod, Joel**, “Tax Compliance and Enforcement,” *Journal of Economic Literature*, dec 2019, 57 (4), 904–954.
- **and Michael Keen**, *Rebellion, Rascals, and Revenue: Tax Follies and Wisdom through the Ages*, Princeton: Princeton University Press, 2021.
- Spruyt, Hendrik**, “The Origins, Development, and Possible Decline of the Modern State,” *Annual Review of Political Science*, 2002, 5 (1), 127–149.
- Stasavage, David**, “Democracy, Autocracy, and Emergency Threats: Lessons for COVID-19 From the Last Thousand Years,” *International Organization*, December 2020, 74 (S1), E1–E17.
- Stewart, Iona**, “Constituency Data: Central Heating, 2021 Census,” June 2023.

- Tilly, Charles**, *Coercion, Capital, and European States, A.D. 990-1990: Ad 990 - 1992*, Cambridge, Mass.: Wiley-Blackwell, March 1993.
- Tirole, Jean**, "Digital Dystopia," *American Economic Review*, 2021, 111 (6), 2007–2048.
- und Wärme, BMWK ExpertInnen-Kommission Gas**, "Sich durch den Winter: Abschlussbericht," Technical Report, BMWK, Berlin October 2022.
- Viner, Katharine**, "Millions in Britain Are Choosing between Heating and Eating. It's Not Too Late to Help Them," *The Guardian*, January 2023.
- Wang, Yuhua**, "State-in-Society 2.0: Toward Fourth-Generation Theories of the State," *Comparative Politics*, October 2021, 54 (1), 175–198.
- Weidmann, Nils B. and Espen Geelmuyden Rød**, *The Internet and Political Protest in Autocracies*, New York, NY: Oxford University Press, August 2019.
- Weigel, Jonathan L**, "The participation dividend of taxation: How citizens in Congo engage more with the state when it tries to tax them," *The Quarterly Journal of Economics*, 2020, 135 (4), 1849–1903.

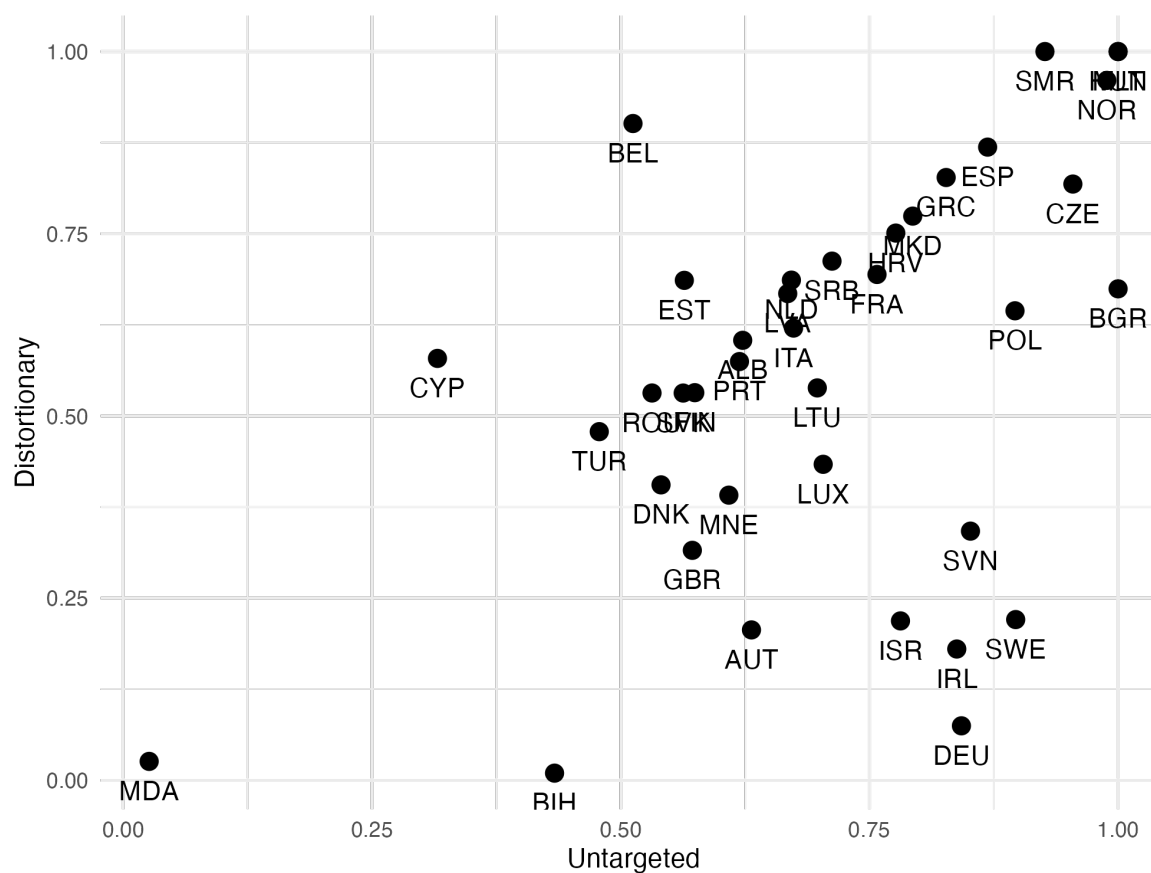
Figures and tables

Figure 1: Time series of global natural gas prices with fitted regime-switch model



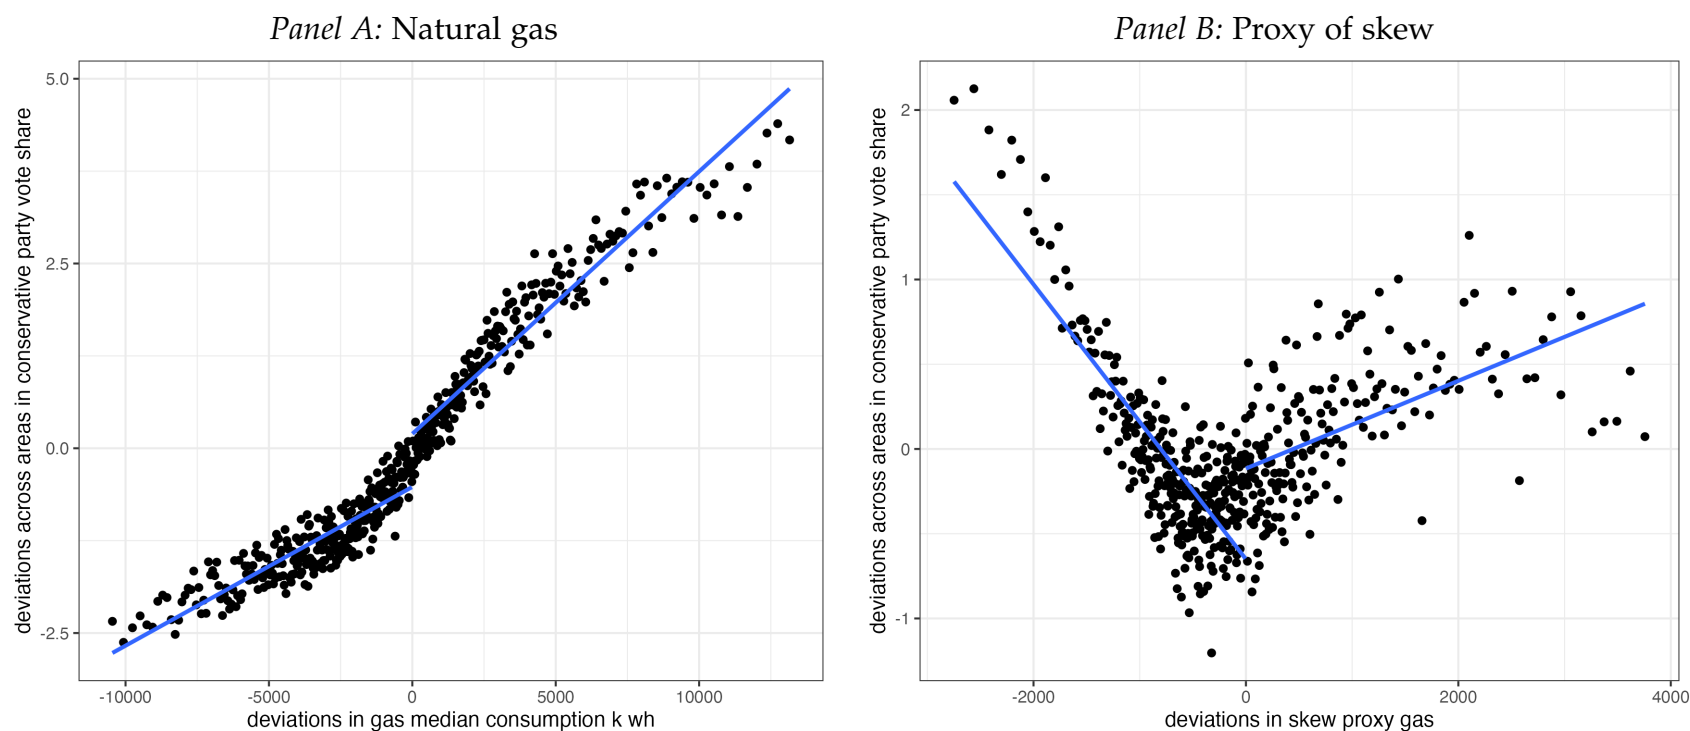
Notes: This figure plots the average natural gas price in US dollars per million metric British thermal unit, along with a Markov-regime-switching model indicating the high-price regime.

Figure 2: Untargeted versus distortionary policy measures in support of households



Notes: This figure plots the share of the fiscal response that is classified as using policy mechanisms that are not targeted and/or distorting the signal function of prices. The underlying data is taken from [Arregui et al. \(2022\)](#) and rescaled. A linear regression would yield an R^2 of approximately 22%.

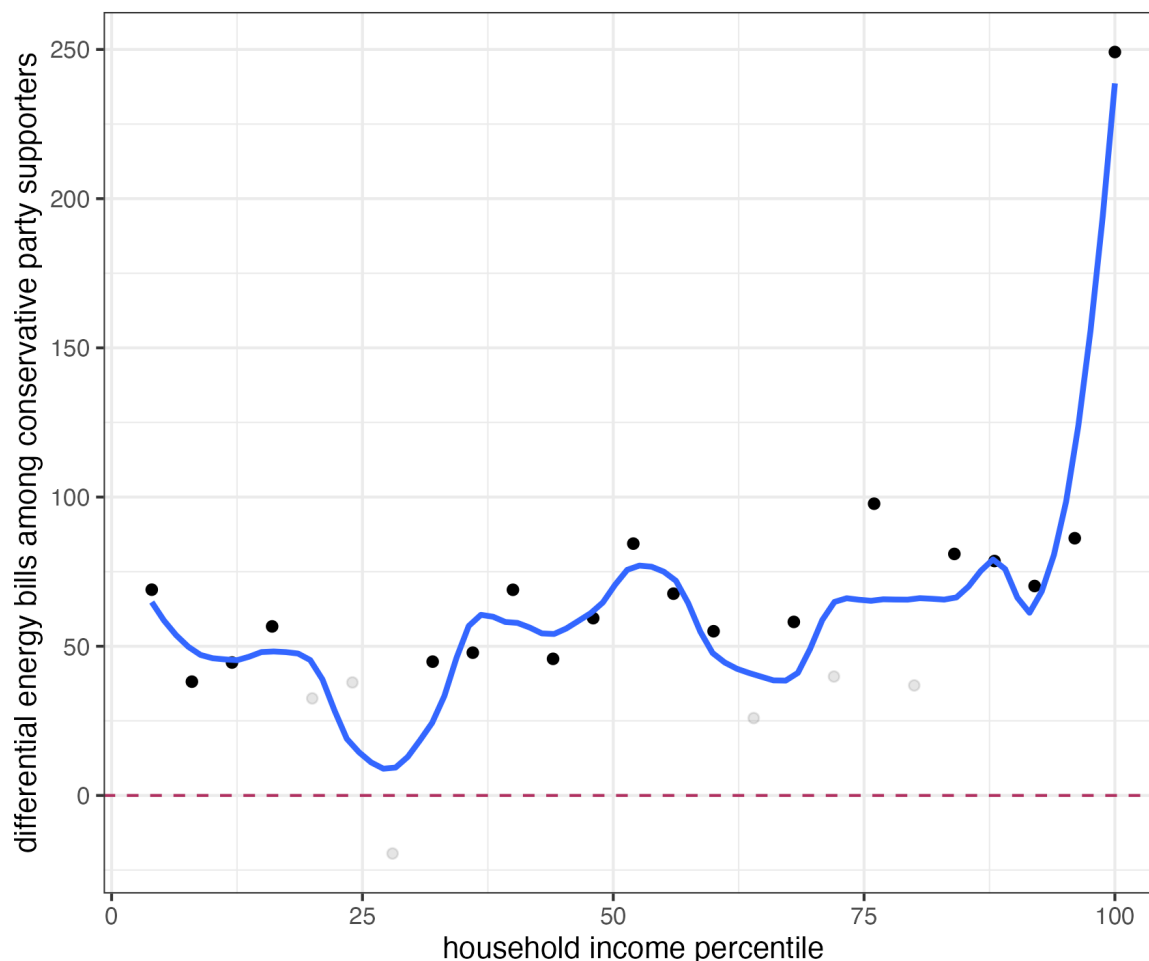
Figure 3: Relationship between energy consumption and political preferences, as measured by ward-level Conservative Party vote shares in local elections



Notes: This figure suggests a strong positive relationship between higher levels of energy consumption and political preferences. Energy consumption is measured as the long-term average of median household consumption in a postcode. There are around 1,102,781 postcodes for which this measure is available in the period from 2013 to 2020. The measure is demeaned by local-authority-level fixed effects to center the data. The vertical axis represents the ward-level vote share that Conservative Party candidates running for local council seats garnered from 2010 onward. The underlying micro-data was used previously in [Fetzer \(2019a\)](#). On the basis of the candidate-by-ward-by-year dataset, we compute the simple, long-run average of Conservative Party vote share to capture stable party-related preferences. By taking the average, we net out fluctuations in vote shares due to, for instance, variation in candidate-specific characteristics. This measure is available for 6,032 wards and, as with energy consumption, local-authority-level fixed effects (2021 boundaries) are removed. The combined dataset has 888,564 observations. For ease of visualization, we present a binned scatterplot with 500 bins, where the averages of the residualized measures are computed for the horizontal and vertical axes. Two linear regressions are fitted, allowing both the intercept and the slope to change around the center of the explanatory variable.

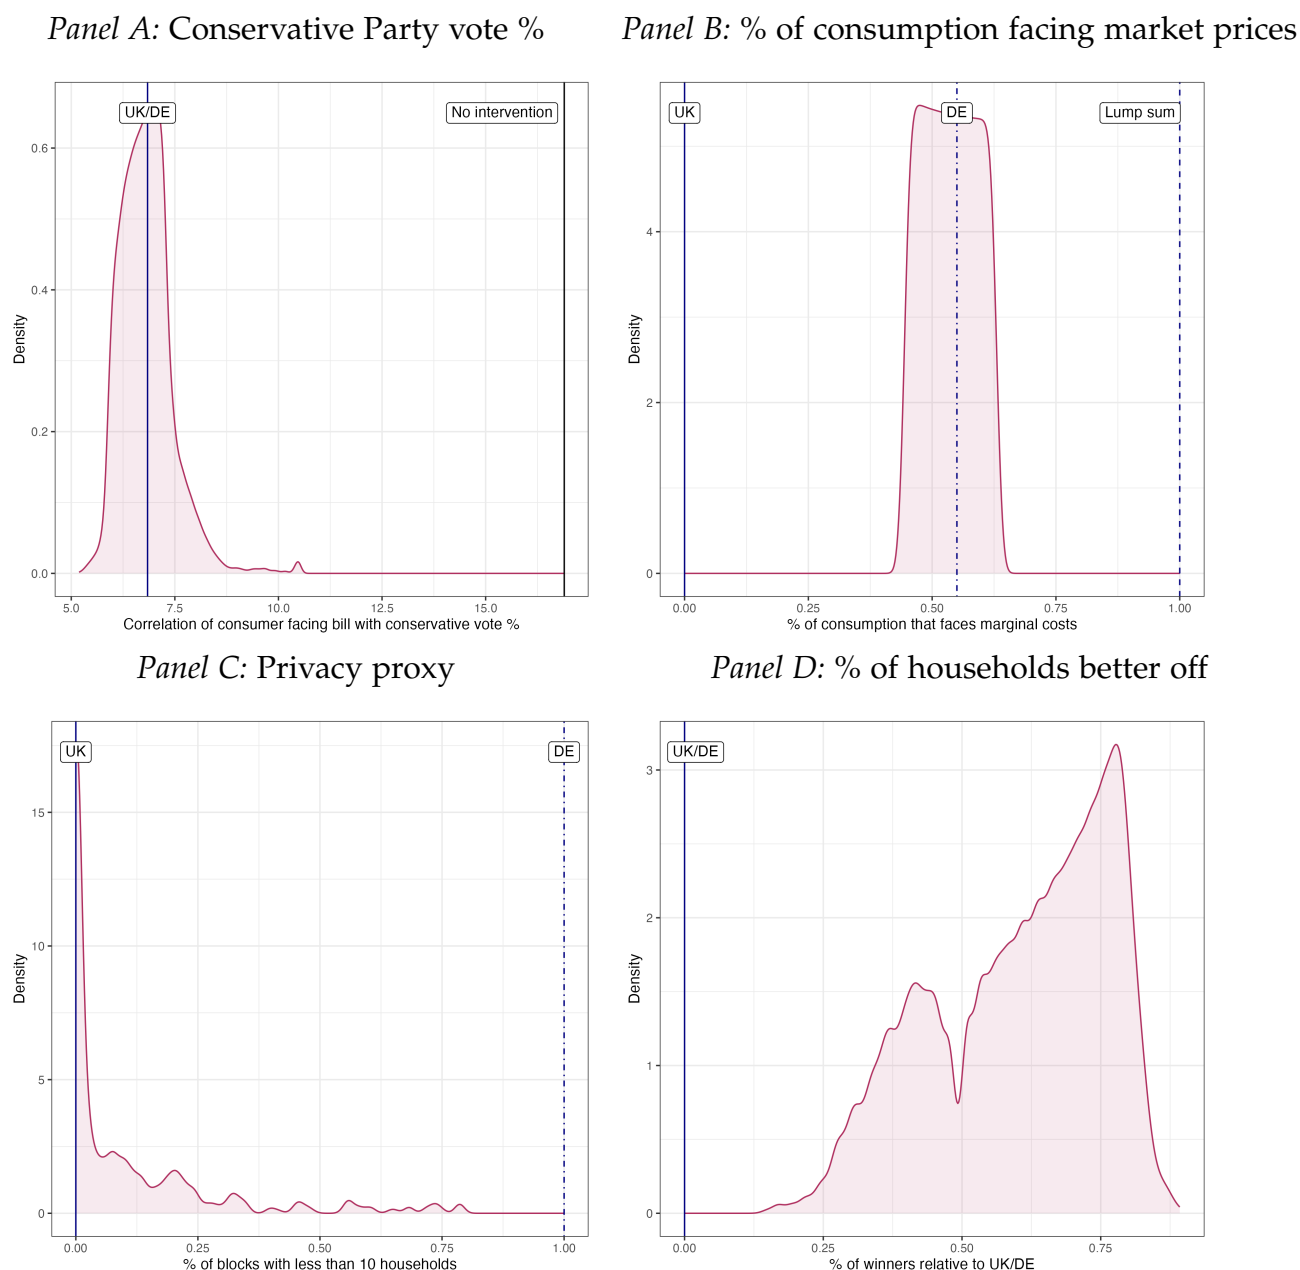
Figure 4: Heterogeneity in the relationship between household income, political preferences and energy consumption – proxied via energy bills

Energy bill differential among Conservative leaning households as a function of income



Notes: This figure plots the non-linear relationship between household income, support for the Conservative Party and energy bills. Each dot represents a point estimate associated with a specific household income percentile. 25 point estimates are shown. Dots are faintly colored if the specific point estimate is not significant at the 5% level. Dots are solid when they are statistically significant for at least the 5% level. A loess fit across the point estimates is presented to highlight the non-linear nature of that relationship. Each regression controls for the following fixed effects: local authority, household size, year and month of the interview, the household income percentile and the survey wave. The figure suggests that energy bills are strongly increasing in absolute terms among Conservative Party supporters across the income distribution, as evidenced by the positive level effect relative to the horizontal dashed line. There is notable heterogeneity, with energy consumption disproportionately higher among very high-income households.

Figure 5: Characterization of the empirical distribution of fiscally neutral two-tier tariff alternatives vis-à-vis equivalents of the UK and German policy responses, respectively

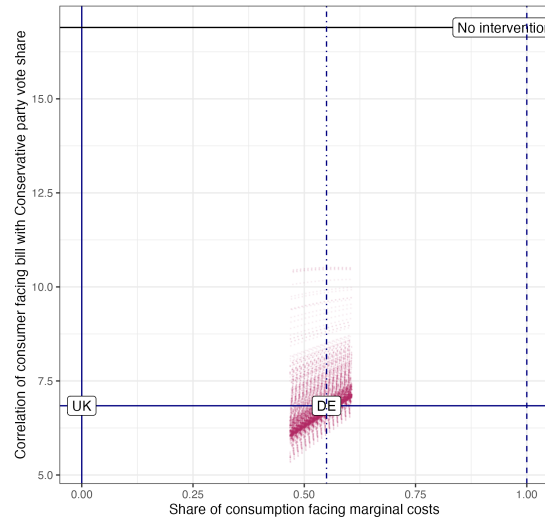


Notes: These figures show the empirical distribution that results from evaluating a broad range of fiscally neutral alternative two-tier block tariffs (based on a range of metrics), relative to the policies that were implemented in the UK (*Energy Price Guarantee*) and Germany (two-tier individualized tariff) respectively. Panel A documents the correlation between bills faced by consumers, net of the subsidy amount, and Conservative Party vote share, with the correlation derived via an exercise similar to what is presented in Table 1. Panel B visualizes the empirical distribution of the share of consumption facing market prices under different policies. Panel C shows the distribution of the share of two-tier tariff blocks that are estimated to include fewer than 10 households, which speaks to the privacy or informational capacity dimension. Panel D presents the empirical distribution of the share of households that, all else equal, would be better off vis-à-vis an individualized two-tier tariff or the untariffed price subsidy, as was implemented via the EPG in the UK.

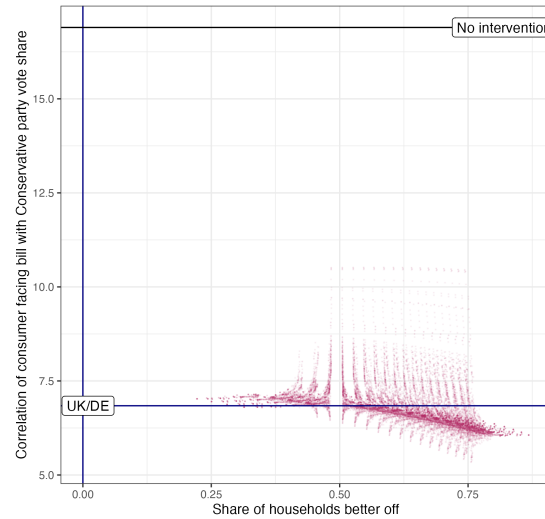
Figure 6: Characterization of trade-offs (or their absence) in the design of energy support policies

Trade-off between degree of political targeting and ...

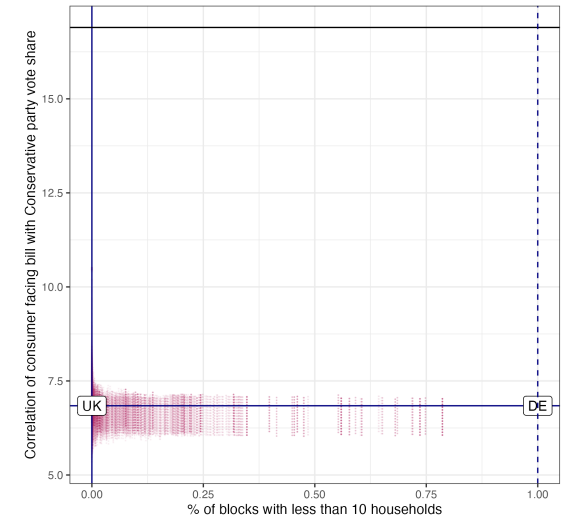
Panel A: Efficiency



Panel B: Equity

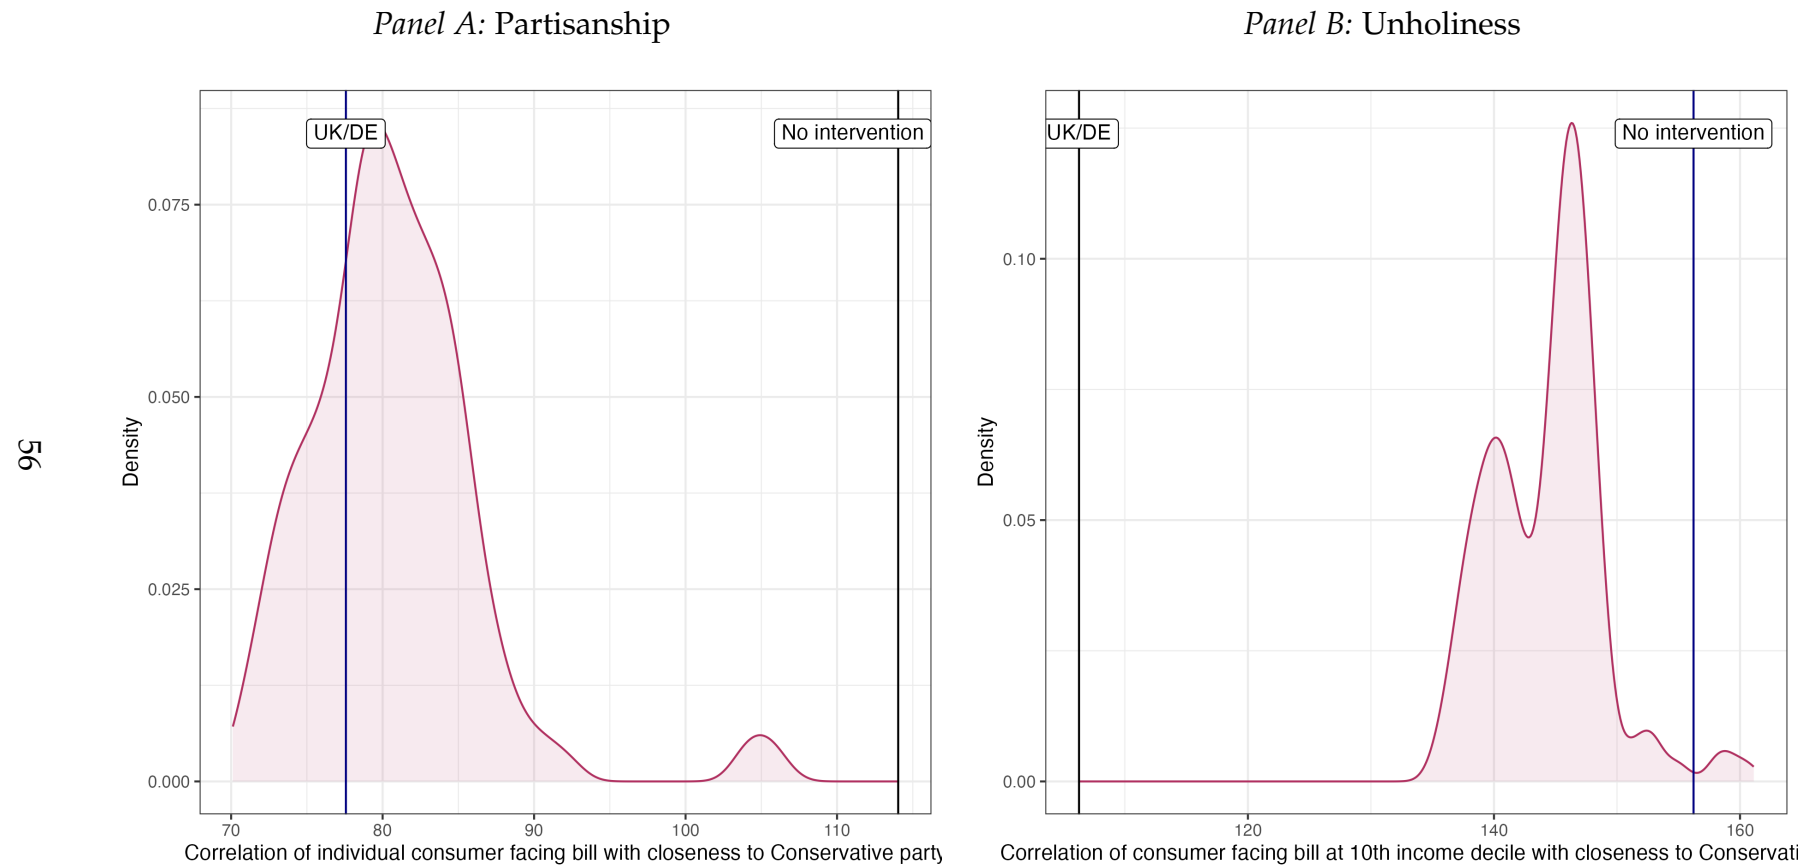


Panel C: Privacy



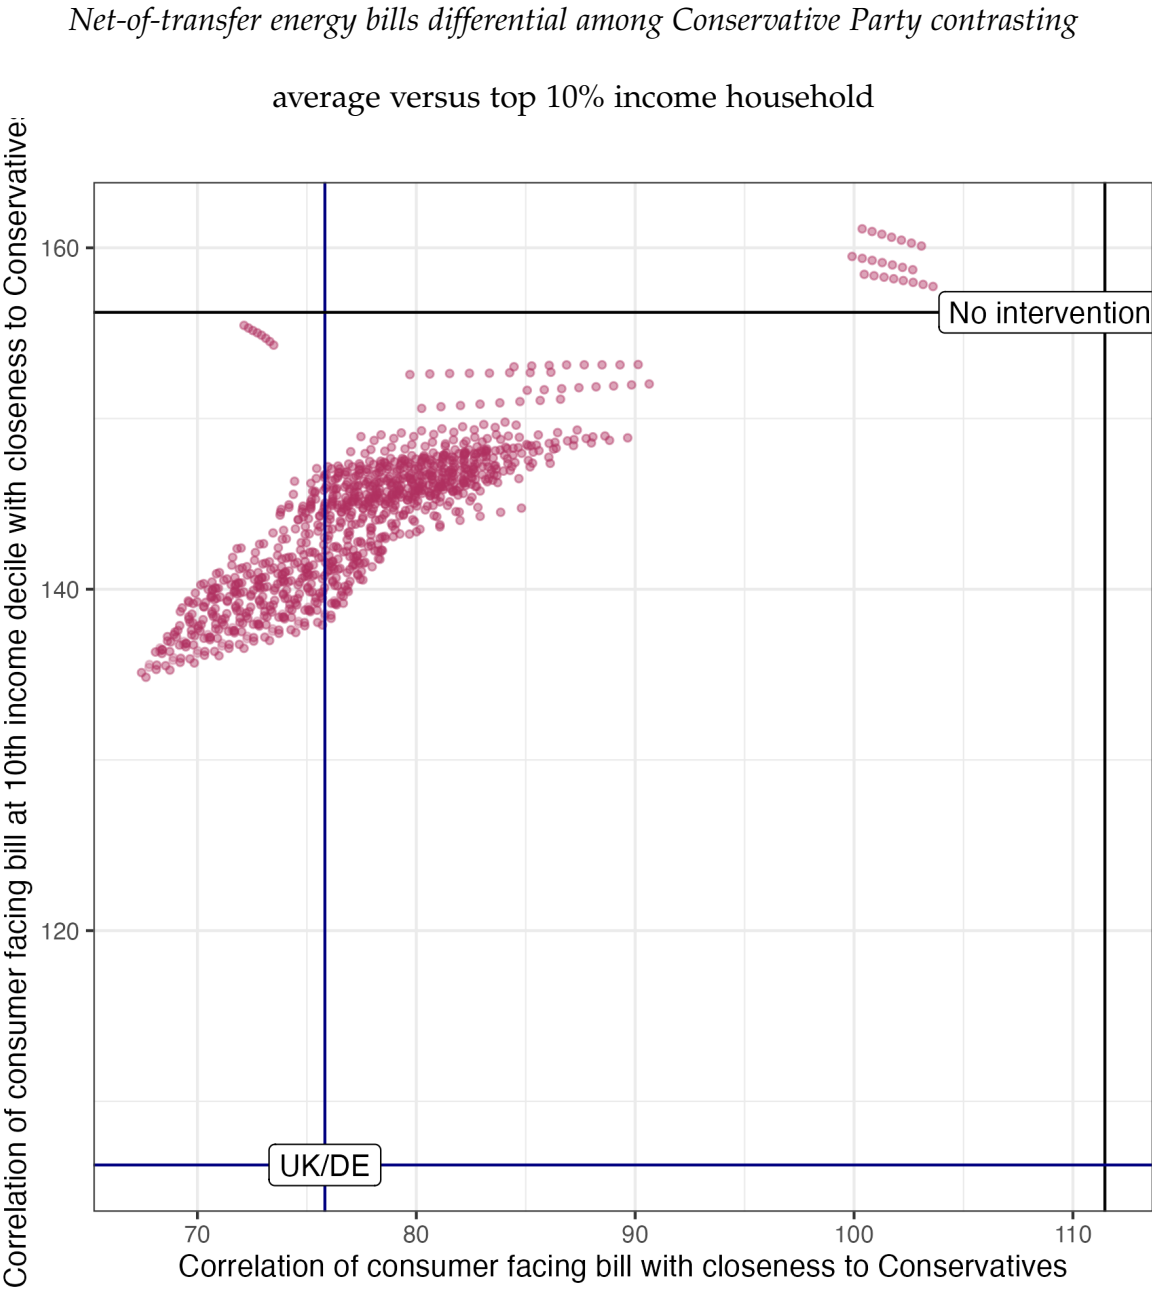
Notes: This panel of figures displays three trade-offs – between the degree of political targeting and equity, efficiency, and privacy – that are obtained by evaluating a broad range of alternative energy support measures, such as two-tier tariffs. Crucially, all simulated alternatives are both similarly costly to the EPG and could have been implemented on the basis of publicly available data. These alternatives are plotted in two dimensions, with each dot representing a policy alternative. They are plotted against the correlation with the degree of political partisanship in the micro-data. The location of the UK's actual policy, the EPG, in that space is indicated via 'UK'. The UK version of the German policy support is indicated as DE. The No-intervention benchmark is also illustrated. In total, four dimensions are shown: Panel A focuses on the trade-off between efficiency (the share of consumption facing market prices) and the degree of political targeting. Panel B displays the empirical distribution of the trade-off between the degree of political targeting and the share of households that would be better off. Panel C shows the trade-off between privacy, measured as the degree of statistical *inferability* of the degree of financial support to households based on the socio-economic characteristics considered, and the degree of political targeting.

Figure 7: Empirical distribution of the *average* degree of political targeting and partisanship for the top 10 income percentile



Notes: This panel of figures displays the empirical distributions of partisanship and unholiness, generated by evaluating a broad range of fiscally neutral two-tier alternative block subsidy schedules (designed around consumer blocks or archetypes), relative to the UK's EPG and an equivalent two-tier individualized tariff, as was implemented in Germany. Panel A presents the empirical distribution of the different correlation coefficients, capturing the correlation between the net-of-transfer energy bills faced by consumers and whether an individual is supporting the Conservative Party. Panel B presents the same relationship as Panel A, but focuses on the correlation among households in the top 10% of the income distribution.

Figure 8: Existence of trade-off between partisanship and degree of unholiness



Notes: Figure presents a bivariate plot of the distribution of the figures presented in Figure 7. The horizontal axis measures the correlation between net-of-transfer energy bills and whether an individual leans Conservative. The vertical axis measures the correlation between net-of-transfers energy bills and whether individuals lean Conservative among the top 10% income households.

Table 1: Individual-level analysis of relationship between energy consumption – proxied by bills – and political preferences

	(1)	(2)	(3)	(4)	(5)
<i>Panel A: Narrow political preferences</i>					
Close to Conservatives	62.88*** (9.839)	62.66*** (9.900)	70.65*** (8.389)	53.85*** (8.231)	17.20** (7.427)
Dependent variable mean	1,285.9	1,285.9	1,285.8	1,285.8	1,285.8
R ²	0.06310	0.09982	0.24690	0.25743	0.31639
Observations	90,589	90,589	90,553	90,553	90,553
<i>Panel B: Broader political preferences</i>					
Close to or would vote Conservatives	58.49*** (8.351)	57.55*** (8.365)	67.04*** (6.769)	49.90*** (6.645)	16.54*** (6.037)
Dependent variable mean	1,278.7	1,278.7	1,278.6	1,278.6	1,278.6
R ²	0.05269	0.08023	0.22331	0.23221	0.28306
Observations	157,061	157,061	157,008	157,008	157,008
Regression specification:					
Local authority & Year x Month of interview	Additive	Interacted	Interacted	Interacted	Interacted
Income and household size			X	X	X
Tenure				X	X
Property characteristics					X

Notes: This table presents results documenting the correlation between self-reported individual-level political preferences and the estimated energy bill, with the latter being a proxy for household energy consumption. Political preferences are measured as a dummy variable that is equal to unity if an individual feels close to the Conservative Party (Panel A), or if an individual feels close to the Conservatives or would vote for them if a general election was held tomorrow (Panel B). The different columns show the point estimates obtained from estimating specifications with a varying set of control variables. Standard errors, given in parentheses, are clustered at the district level, with stars indicating *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

Appendix to “Informational Boundaries of the State”

For Online Publication

A Policy alternatives

A.1 Representative agent’s equivalent two-tier tariff

We begin by describing how we arrive at a representative household’s equivalent two-tier tariff. The UK’s *Energy Price Guarantee* effectively reduced the energy prices that households faced by 50% compared to the market price p_m that would have prevailed from October 2022 onward ([Fetzer, 2022a](#); [DESNZ, 2023](#)). There is an equivalent formulation, whereby households could have received a subsidized price p_s set at the 2021 October energy prices on the first 50% of consumption q_i , while facing market prices p_m on the remaining 50% of consumption. The “market price” could have been set, as per the *Office of Gas and Electricity Market’s* (Ofgem’s) regular energy price cap that would have applied without intervention ([Ofgem, 2023](#)).

More formally, with any two-tier tariff, consumers face two different sets of prices p_s and p_m , where p_m denotes the market price, while p_s denotes the subsidized price, with $p_s < p_m$. For the simulations, we set p_s equal to the energy price that prevailed in the year prior to the Russian invasion of Ukraine, which, in turn, is equal to the energy price cap set by Ofgem in October 2022. The energy price cap sets the maximal unit price per kilowatt hour (kWh) that energy suppliers can charge customers and is reviewed on a quarterly basis ([Ofgem, 2023](#)). This price is designed to allow energy firms to cover their cost, while also allowing for a profit margin in the regulated industry. This is because the price p_m that constitutes the market rate is the energy price cap that was announced in October 2022, and thus represents the price that energy suppliers would have been able to maximally charge customers. The *Energy Price Guarantee* reduced that price cap by around 50%.

The subsidy is designed so that the representative household faces the subsidized price p_s on the first q_m units of consumption. That is, we can write the representative household's estimated bills under such a two-tier tariff scheme as follows:

$$E(C^{\text{Two Tier}}) = E(\min\{q_m, q_i\}) \times p_s + E(\max\{q_m - q_i, 0\}) \times p_m$$

The amount of subsidy that the representative household receives S_i can be written as

$$E(S^{\text{Two Tier}}) = E(\min\{q_m, q_i\}) \times (p_m - p_s)$$

The *Energy Price Guarantee* (EPG), rather than setting a two-tier price system, can be represented as a wedge $\tau \in (0, 1)$ that lowers the price consumers face relative to the market price p_m . That is, we can write the total bills that a representative household faces as

$$E(C^{\text{EPG}}) = E(q) \times \tau \times p_m$$

and the implicit subsidy as:

$$E(S^{\text{EPG}}) = E(q) \times (1 - \tau)p_m$$

We take τ as given based on the design parameters of the EPG ([DESNZ, 2023](#)). This allows us to identify the corresponding q_m threshold that would produce the same bills and subsidy volumes under a two-tier tariff. For ease of exposition, let $q_m = s \times q$, this implies that the two-tier tariff that is equivalent to the EPG can be computed by solving the following system of equations

$$E(C^{\text{EPG}}) = E(C^{\text{Two Tier}})$$

$$E(S^{\text{Two Tier}}) = E(S^{\text{EPG}})$$

Using the supplied data, we see that, with $s \approx 0.5$, the two-tier tariff produces, at the household-level, the same value of a subsidy and bills as the one generated by the EPG. That is: setting $q_m = 0.5 \times E(q)$, with the prices p_m and p_s exogenously given, a two-tier tariff would generate the same expected energy bill compared to the energy price cap. The notable difference, though, is the role that prices play: in the two-tier tariff solution, the signal value of market prices p_m is maintained.

Households' expectations could further be anchored in a dynamic fashion by announcing a dynamically declining threshold q_m over time, following best practices around subsidies in the form of a sunset clause.

A.2 Individualized two-tier tariff

The two-tier tariff solution that is fiscally equivalent – ex ante – to the rather blunt *Energy Price Guarantee* (EPG) also maps to an individualized two-tier tariff. A policy alternative that would be (much) more targeted than a blunt intervention in energy-price-setting behavior would introduce an individualized quota upon which a subsidized price is levied $q_{mi} = s \times q_{i,-1}$, where the quota is set based on, for instance, last year's energy consumption. The individual-level two-tier tariff solution that is – ex-ante – equally costly to the EPG would set $s = 0.5$ – due to the *law of iterated expectations*.

In terms of *implementability*, such a transfer system would require data on individual households' energy consumption. Such data may not be available to policymakers because of data protection laws or other privacy considerations. The German policy approach was to administer individualized transfers through private sector entities. Furthermore, to hold constant preferences over redistribution or inequality, the lump-sum transfer associated with the individualized quota was passed through the income tax system, meaning that high-income households – who receive a large implicit lump-sum transfer – have to pay income tax on that transfer, commensurate with their income.

The individualized tariff, by setting an individualized quota, has desirable properties: it preserves the signal value of prices for the bulk of consumption. Given

the existing research on two-tier tariffs and the evidence suggesting that consumers respond to average, rather than marginal, prices (Ito, 2014), both tariffs are ex-ante equal as they would produce the same *average increase* in bills. Yet, in the context of the two-tier tariff, the signal value of prices is maintained, while it is weakened in the case of the EPG.

Individualized tariffs also come with further implementation constraints or limitations since they require granular data that may not be in the public domain, or can only be gathered and/or used by infringing on data protection laws or other privacy regulations. As regards the implicit subsidy, all else equal, the fiscally neutral individualized two-tier tariff and the EPG would produce the same level of household subsidization.

A.3 Targeted lump-sum transfer

In addition to considering the two-tier tariff and the individualized two-tier tariff, we next consider an alternative – a targeted lump-sum transfer. This can be implemented especially easily, given that, in essence, it would simply require sending out physical checks. Support would be already more targeted if it took the form of providing council tax credits. In this way, support would consist in reducing another type of financial burden – council tax – that households face.

Naturally, the key distinction here is that, implicitly, with such a lump-sum transfer, households that have consumption *below* the given block that is subsidized $q_m|x$, implicitly are left *better off*. But it is a particularly easily implementable way to provide targeted support to households affected by rising energy bills.

$$E(S^{\text{lump-sum}}) = s \times E(q_m|x)$$

The amount of subsidy that a household receives is now just a constant. For the purposes of the aggregate comparisons across different transfer schedules, however, we consider a lump-sum transfer as effectively providing an energy price of 0 for the first q_m units of consumption and the market price for the rest for ease of

comparison.¹

B Further formal derivations

Deriving consumer equivalent variation functions in full: There are two consumer types, $\theta = \{H, L\}$. Consumers have income m_θ (with $m_H > m_L$) and can purchase two goods: energy (good x) and another good (good y) representing all other consumption. The price of good y is normalised to 1; the price of energy is p^{-1} before the energy crisis, and rises to p^0 as a result of the crisis.

Consumers have preferences $u(x, y)$ over the two goods, with $u_1(x, y) > 0$, $u_2(x, y) > 0$, $u_{11}(x, y) < 0$, $u_{22}(x, y) < 0$, $u_{21}(x, y) \geq 0$ and $u_{12}(x, y) \geq 0$. As the price of the non-energy good is fixed at 1, the price vector at any time can be completely described by the price of energy, p . Denote the Marshallian demand of consumer of type θ for energy $x(p; m_\theta)$ and the Hicksian demand of a consumer at utility level u^n for energy as $h(p; u^n)$. Consumers are identical in all ways apart from income and so have identical Hicksian demand functions, conditional on a particular utility level.

The consumer receives support from the policymaker in the aftermath of the crisis in the form of a mix of lump-sum transfers and a subsidy on price. We proceed by measuring the consumer's utility from the policy response in money-metric terms by calculating the consumer's equivalent variation. That is, we compute the transfer of wealth to the consumer at prices $(p^0, 1)$ that would be required for her to achieve the same utility at this price vector, compared to the utility she achieves from the mix of lump-sum transfer and subsidy the policymaker decides to implement.

The equivalent variation of any lump-sum amount g the policymaker transfers to the consumer is equal to the size of that lump-sum transfer.

Let $g_\theta^e(s)$ be the amount of money that the policymaker spends on subsidizing a consumer with type θ by imposing a subsidy of s , such that $g_s = g_H^e(s) + g_L^e(s)$ is the total amount the policymaker spends on the subsidy program. s is the subsidy

¹With a targeted lump-sum transfer, the share of consumption that faces marginal cost is 100%.

the policymaker places on the price of energy which results in the total expenditure g_s . If the policymaker gave g_θ^e as a lump-sum amount to the consumer with type θ , the equivalent variation of this transfer would be equal to g_θ^e . The actual equivalent variation of the subsidy program is, however, equal to:

$$g_\theta^e(s) - DWL_\theta(s)$$

where $DWL_\theta(s)$ is the deadweight loss of commodity subsidization that arises as a result of consumer with type θ being subsidized at s for their energy consumption. The deadweight loss is equal to $g_\theta^e(s) - EV(p^0, p^0 - s; m_\theta)$, where $EV(p^0, p^0 - s; m_\theta)$ is the equivalent variation of the policy, which decreases the price of energy from p^0 to $p^0 - s$ for a consumer with income m_θ .

Let u_θ^p be the utility a consumer with income m_θ achieves when the price of energy is p . Then, the cost to the policymaker of subsidizing consumer type θ with a subsidy of size s is

$$g_\theta^e(s) = sh(p^0 - s; u_\theta^{p^0-s})$$

The equivalent variation of this subsidy to this consumer is:

$$EV(p^0, p^0 - s; m_\theta) = \int_{p^0-s}^{p^0} h(p; u_\theta^{p^0-s}) dp$$

and thus the deadweight loss function can be written in full as

$$DWL_\theta(s) = sh(p^0 - s; u_\theta^{p^0-s}) - \int_{p^0-s}^{p^0} h(p; u_\theta^{p^0-s}) dp$$

As shown in [Mas-Colell et al. \(1995\)](#), as $h(p; u^n)$ is strictly decreasing in p , the deadweight loss is strictly positive for all $s > 0$. Further, the derivative of this loss

function with respect to s is equal to 0 for $s = 0$ and is strictly positive for all $s > 0$.

Thus, while the total cost to the government of the subsidy program is

$$g_s = g_H^e(s) + g_L^e(s),$$

the total benefit (measured as total equivalent variation) accruing to consumers is $\phi_H(s) + \phi_L(s)$, where $\phi_\theta(s) = g_\theta^e(s) - DWL_\theta(s)$, and $\phi'_H(s) > \phi'_L(s)$ (and $\phi_H(s) > \phi_L(s) \forall s$) because $u_H^{p^0-s} > u_L^{p^0-s}$. Further, following the above reasoning, the properties of the deadweight loss function, $\phi_\theta(0) = g_\theta^e(0)$ and $\phi'_\theta(0) = g_\theta^e(0)$ because $DWL'_\theta(0) = 0$.

As g_s is strictly increasing in s , $\phi'_H(s)$ and $\phi'_L(s)$ can be rewritten as implicit functions of g_s (which we define $f_H(g_s)$ and $f_L(g_s)$) that satisfy all of the properties described in the main body of the paper.

Deriving first-order conditions for the policymaker's problem: Defining the Lagrangian multiplier λ , the first derivatives of the Lagrangian of the problem are:

$$g^h: \quad \beta \Delta_H c'(g^h + f_H(g_s)) + (1 - \beta) \Delta_L c'(g^h + f_L(g_s)) = \lambda \quad (10)$$

$$g^l: \quad \beta \Delta_L c'(g^l + f_L(g_s)) + (1 - \beta) \Delta_H c'(g^l + f_H(g_s)) = \lambda \quad (11)$$

$$g_s: \quad f'_H(g_s) \beta \Delta_H c'(g^h + f_H(g_s)) + f'_L(g_s) (1 - \beta) \Delta_L c'(g^h + f_L(g_s)) \\ + f'_L(g_s) \beta \Delta_L c'(g^l + f_L(g_s)) + f'_H(g_s) (1 - \beta) \Delta_H c'(g^l + f_H(g_s)) = \lambda \quad (12)$$

Combining the equations yields the uncertain lump-sum redistribution condition and the two subsidy balance conditions:

- The uncertain lump-sum redistribution condition: Combining (10) and (11)

and rearranging yields:

$$\begin{aligned} & \beta \left(\Delta_H c'(g^h + f_H(g_s)) - \Delta_L c'(g^l + f_L(g_s)) \right) \\ &= (1 - \beta) \left(\Delta_H c'(g^l + f_H(g_s)) - \Delta_L c'(g^h + f_L(g_s)) \right) \end{aligned} \quad (13)$$

- The subsidy balance conditions: Combining (12) with (10) or (10), respectively, yield the high-type transfer subsidy balance condition:

$$\begin{aligned} & f'_H(g_s) \beta \Delta_H c'(g^h + f_H(g_s)) + f'_L(g_s) (1 - \beta) \Delta_L c'(g^h + f_L(g_s)) \\ &+ f'_L(g_s) \beta \Delta_L c'(g^l + f_L(g_s)) + f'_H(g_s) (1 - \beta) \Delta_H c'(g^l + f_H(g_s)) \\ &= \beta \Delta_H c'(g^h + f_H(g_s)) + (1 - \beta) \Delta_L c'(g^h + f_L(g_s)) \end{aligned} \quad (14)$$

and the low-type transfer subsidy balance condition:

$$\begin{aligned} & f'_H(g_s) \beta \Delta_H c'(g^h + f_H(g_s)) + f'_L(g_s) (1 - \beta) \Delta_L c'(g^h + f_L(g_s)) \\ &+ f'_L(g_s) \beta \Delta_L c'(g^l + f_L(g_s)) + f'_H(g_s) (1 - \beta) \Delta_H c'(g^l + f_H(g_s)) \\ &= \beta \Delta_L c'(g^l + f_L(g_s)) + (1 - \beta) \Delta_H c'(g^l + f_H(g_s)) \end{aligned} \quad (15)$$

Proof of Corollary 1: The $\Delta_L > \Delta_H$ policymaker will choose $g_h = g_l$ and $g_s = 0$, and the $\Delta_H > \Delta_L$ policymaker will choose $g_s \neq 0$. The corollary then follows from the fact that $\Delta_H = 1 - \Delta_L$.

C Further details on simulation implementation

C.1 Description of household characteristics that can be conditioned on (i.e. elements of x)

Spatial identifiers We consider a broad vector of 13 different spatial identifiers at which a representative household's energy consumption is estimated, ranging from the most granular postcode level – which includes more than one million different

households – to the much coarser region level that subdivides England into nine² different regions.

Indices of multiple deprivation The indices of multiple deprivation provide a higher-dimensional view of the relative social and economic deprivation of different parts of the country in a variety of domains: income, employment, education, skills and training, health and disability, barriers to housing and services, and a composite deprivation score. Each dimension of relative deprivation is measured at the level of the lower layer super output area (LSOA). The deprivation features we cover capture the overall multiple deprivation score as well as the main constituent components: income, employment, education and skills, health and disability, as well as barriers to housing and services deprivation. The use of such indices allows for the two-tier tariffs to be stratified by the relative socio-economic deprivation of an area's population. LSOA's are commonly used in public policy. They are designed as statistical geographies, built from census area blocks, to be comparable, having, on average, a similar number of residents. In total, there are 32,000 spatial units in England. For each index, we construct a measure of both the quintile and a dummy variable indicating whether or not a score is above or below the median. This will, for example, identify the most deprived areas based on specific domains across LSOA spatial units. In total, there are six quintile and six binary features.

Council tax band Council tax is a tax payable for the provision of local services. Each residential property in the UK is liable to pay council tax, with local authorities enforcing and collecting this tax. Notionally, council tax liability should be linked to the underlying property value. Yet, the underlying rating lists have not been updated since 1991. Accurate records of property values are particularly relevant since subsidies or social tariffs could be directly linked to local councils' existing tax collection or enforcement mechanisms. Indeed, these mechanisms could have been leveraged to provide more targeted energy bill support. The feature is cat-

²These nine regions are: London, the North East, North West, East Midlands, West Midlands, Yorkshire, East of England, South West, and South East.

egorical, which implies that a consumption quota could be designed to allow for differentiated levels of subsidized energy consumption for homes in an area.

Property characteristics For each property, we consider four additional features: the property’s main heating fuel (gas or electricity), its type (e.g. flat versus detached home), age, and size. This constitutes four additional features.

In order to obtain a comprehensive menu of possible alternative energy subsidy schedules, we construct all potential ways of combining these features. In this way, we arrive at a reference estimate $E(q_{i,-1}|x_i)$ on the basis of which the energy subsidy could have been handed out.

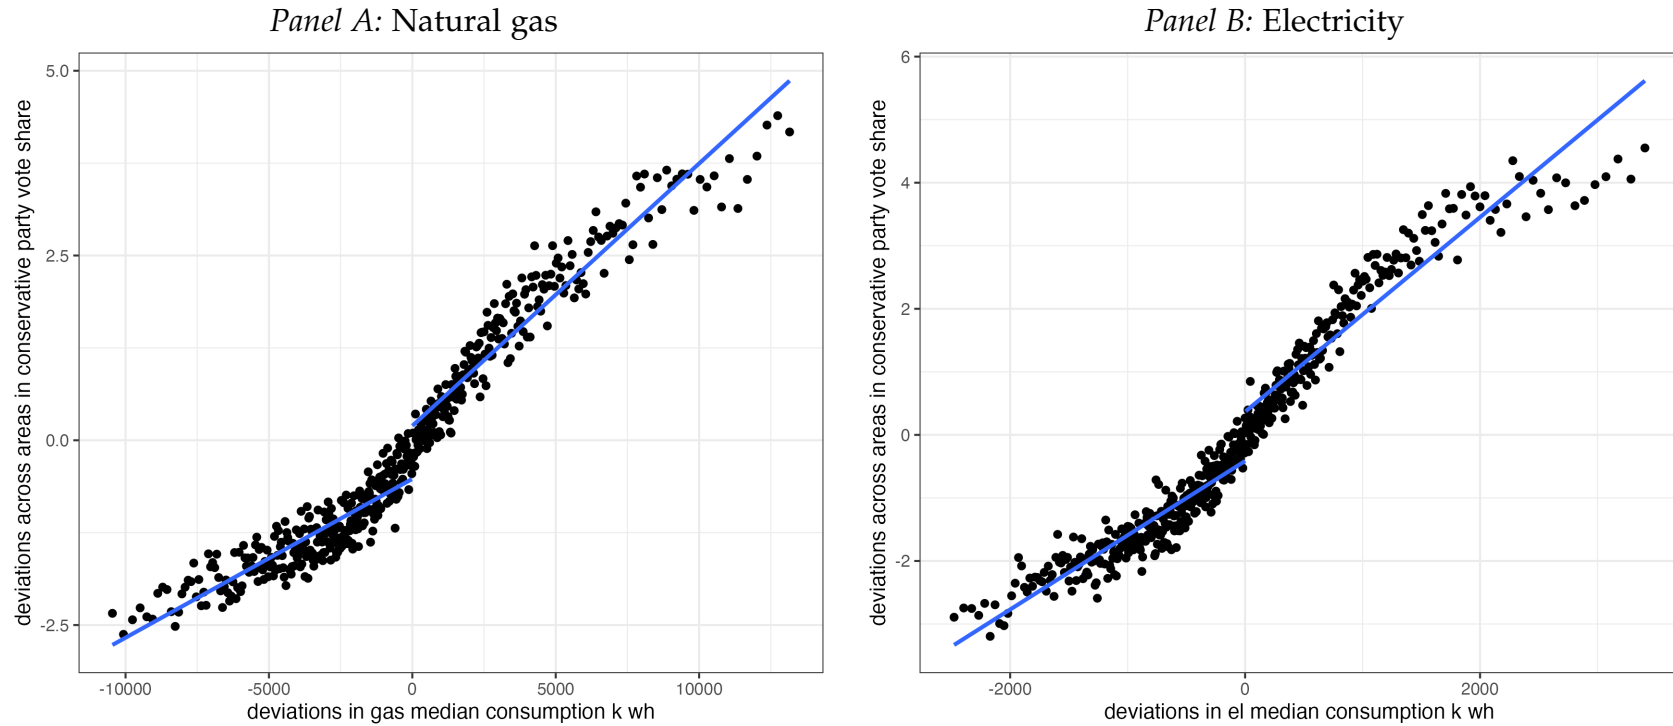
We consider 31 features. This implies $2^{31} = 2,147,483,648$ ways of selecting different subsets of features. For some features, however, such as spatial identifiers, we consider different blocks in the analysis, given that spatial identifiers are broadly nested when moving from more to less granular levels of aggregation. Instead of identifying 2^{31} potential combinations of features, we consider 14×2^{17} potential combinations. This still leaves us with 1,835,008 potential combinations of different features. Simulating this broad set of counterfactual policies is computationally infeasible. Thus, we seek to reduce the dimensionality further.

Near population data The indices of multiple deprivation scores are discretized into six binary features (above or below the median score) or organized as six features that capture the relative quintile of an LSOA in a given deprivation area domain. We consider each of these two groups of features – the binary and the quintile set – separately. In total, we consider 14×2^{11} ways of sampling features, with the binary set of deprivation indices and a further 14×2^{11} for the quintile feature set. The end result is a set of $2 \times 14 \times 2^{11} = 57,344$ combinations of features x_i that can be considered to construct an estimate $E(q_{i,-1}|x_i)$. We further restrict the subset of possible policies to a stratified random sample that covers 5% of the feasible policies. The sample is stratified by policy complexity – as measured by the absolute count of the number of features that are considered.

Individual-level data For the simulation results that use individual-level survey data, we use data for a smaller subset of features, as individual-level data only contain a smaller policy space. The advantage of these data is that they provide us with sharp measurements of political preferences at the individual level.

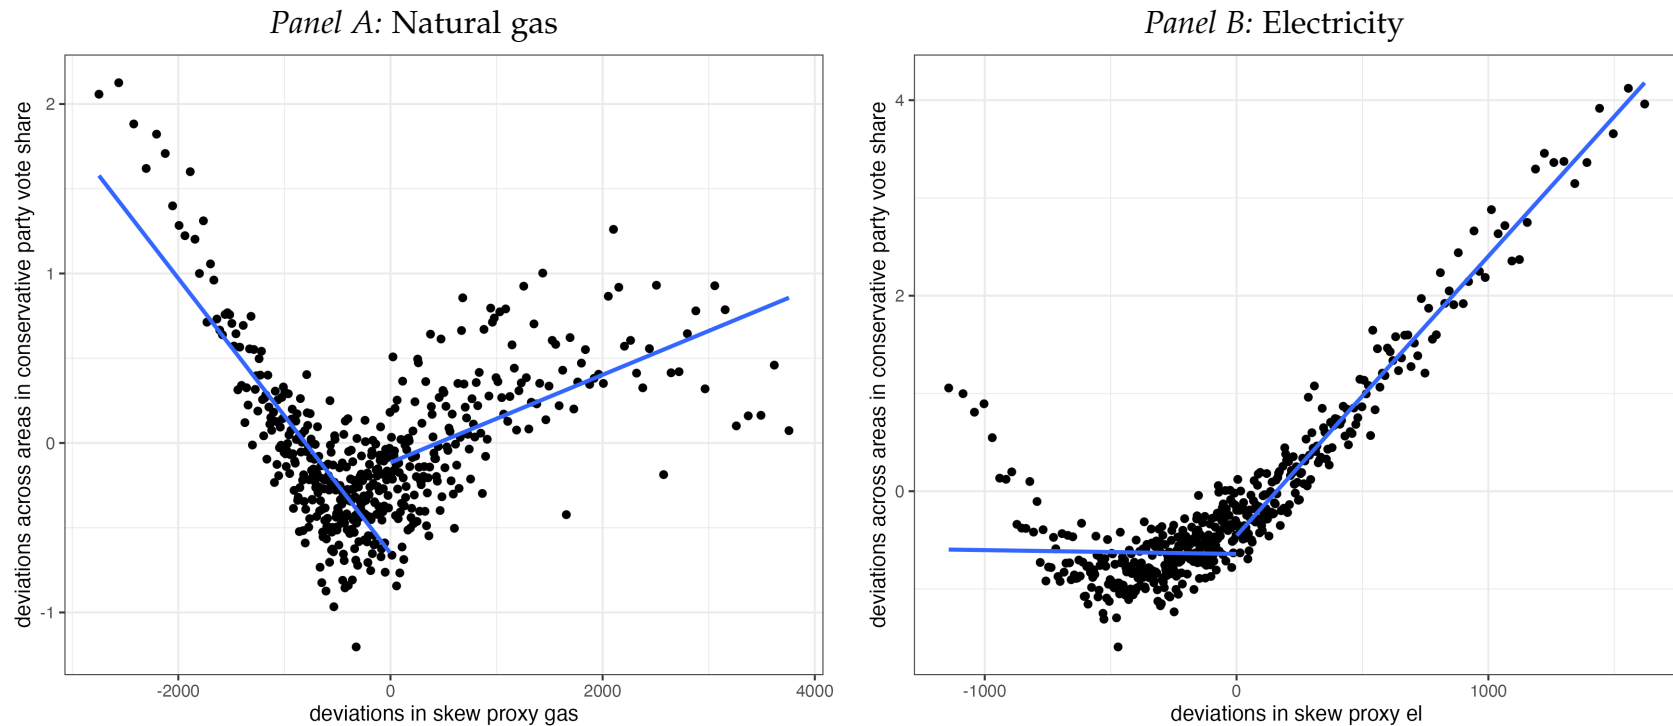
C.2 Additional Figures

Figure A1: Relationship between energy consumption and political preferences, as measured by ward-level Conservative Party vote shares in local elections



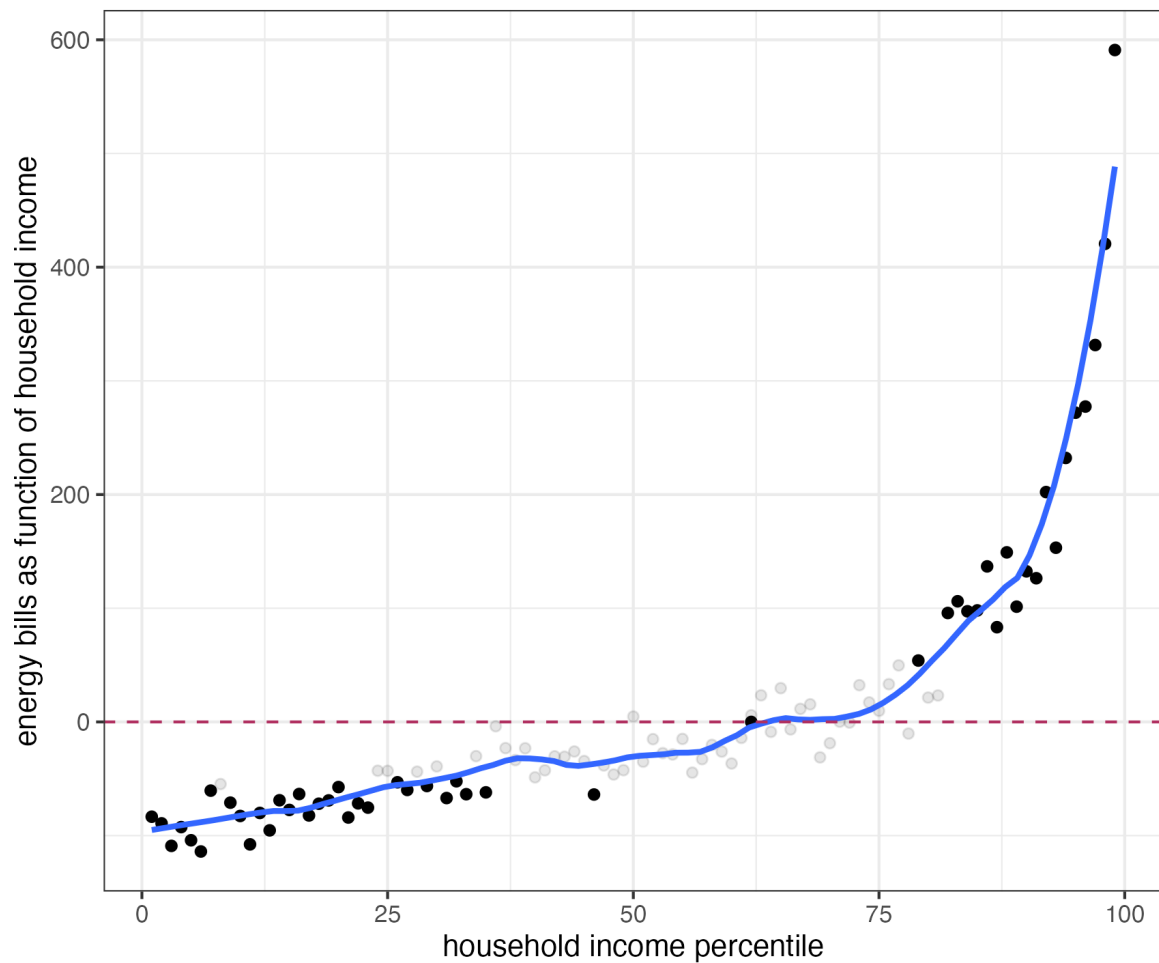
Notes: This figure suggests a strong positive relationship between higher levels of energy consumption and political preferences. The figure presents the results from regressions that capture the relationship between energy subsidies and the Conservative Party vote share. Energy consumption is measured as the long-term average of median household consumption in a postcode. There are around 1,102,781 postcodes for which this measure is available in the period from 2013 to 2020. The measure is demeaned by local-authority-level fixed effects to center the data. The vertical axis represents the ward-level vote share that Conservative Party candidates running for local council garnered from 2010 onwards. The underlying micro-data was previously used in [Fetzer \(2019a\)](#). The simple average of Conservative Party vote share is computed to generate a longer-term measure, capturing stable political preferences, on the basis of the candidate-by-ward-by-year dataset. This is done to net out fluctuations in vote shares due to, for instance, variation in candidate-specific characteristics. This measure is available for 6,032 wards and, as with energy consumption, local-authority-level fixed effects (2021 boundaries) are removed. The combined dataset has 888,564 observations. For ease of visualization, we present a binned scatterplot with 500 bins, where the averages of the residualized measures are computed both for the horizontal and vertical axes. Two linear regressions are fitted, allowing both the intercept and the slope to change around the center of the explanatory variable.

Figure A2: Relationship between proxy measure of energy consumption inequality, measured as the difference between mean and median, and political preferences, as measured by ward-level Conservative Party vote shares in local elections



Notes: This right-hand panel of this figure visualizes the V-shaped pattern between natural gas consumption inequality and support for the Conservative Party. The left-hand panel, by contrast, shows that the relationship between electricity consumption inequality and Conservative Party support roughly follows the shape of an inverted L. The variables on the x-axes are proxies for the skew of the distribution of electricity and natural gas consumption, respectively, with positive values indicating a right-ward skew (mean greater than median) and negative values indicating a left skew. There are around 1,102,781 postcodes for which this measure is available from 2013 to 2020 inclusive. The measure is demeaned by local-authority-level fixed effects to center the data. The vertical axis shows the ward-level Conservative Party vote share that Conservative candidates running for local council were able to achieve from 2010 onwards across elections. The underlying micro-data was previously used in [Fetzer \(2019a\)](#). The simple average is computed to produce a longer-term measure of stable political preferences from the candidate-by-ward-by-year level dataset. This measure is available for 6,032 wards and, as with energy consumption, local-authority-level fixed effects (2021 boundaries) are removed. The combined dataset has 888,564 observations. For ease of visualization, we present a binned scatterplot, with 500 bins in which the averages of the residualized measures are computed both for the horizontal and vertical axes. Two linear regressions are fitted, allowing both the intercept and slope to change around the centered data.

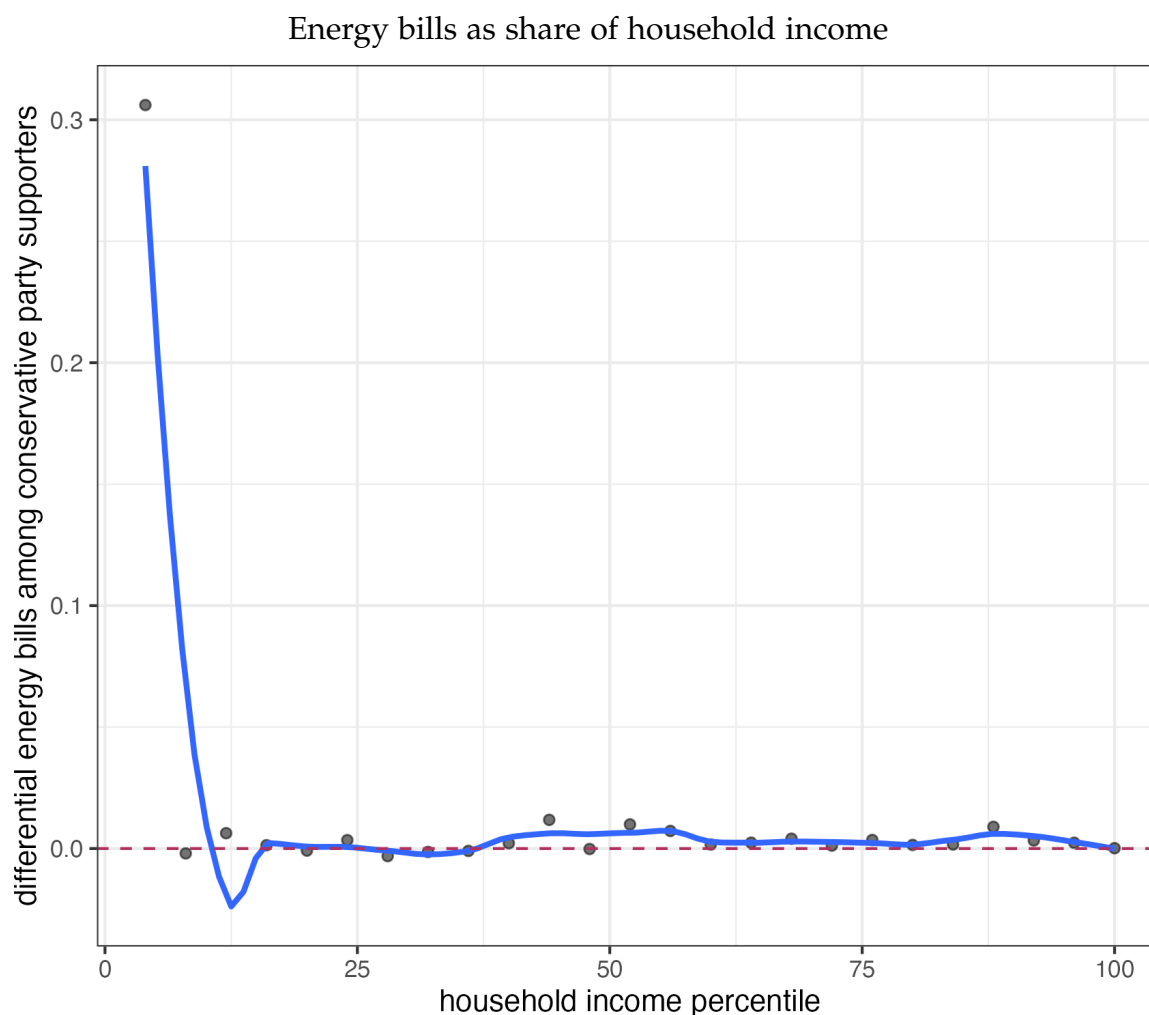
Figure A3: Energy bills as a function of household income



Notes: This figure plots the non-linear relationship between household income and energy consumption across households. Each dot represents a percentile in the income distribution. The omitted category is the 50th percentile meaning the estimates on the vertical line represent the difference of energy consumption of a household relative to the household with the median income. Dots are solid when they are statistically significant for at least the 5% level. A loess fit across the point estimates is presented to highlight the non-linear nature of that relationship. Each regression controls for the following fixed effects: local authority, household size, year and month of the interview, and the survey wave.

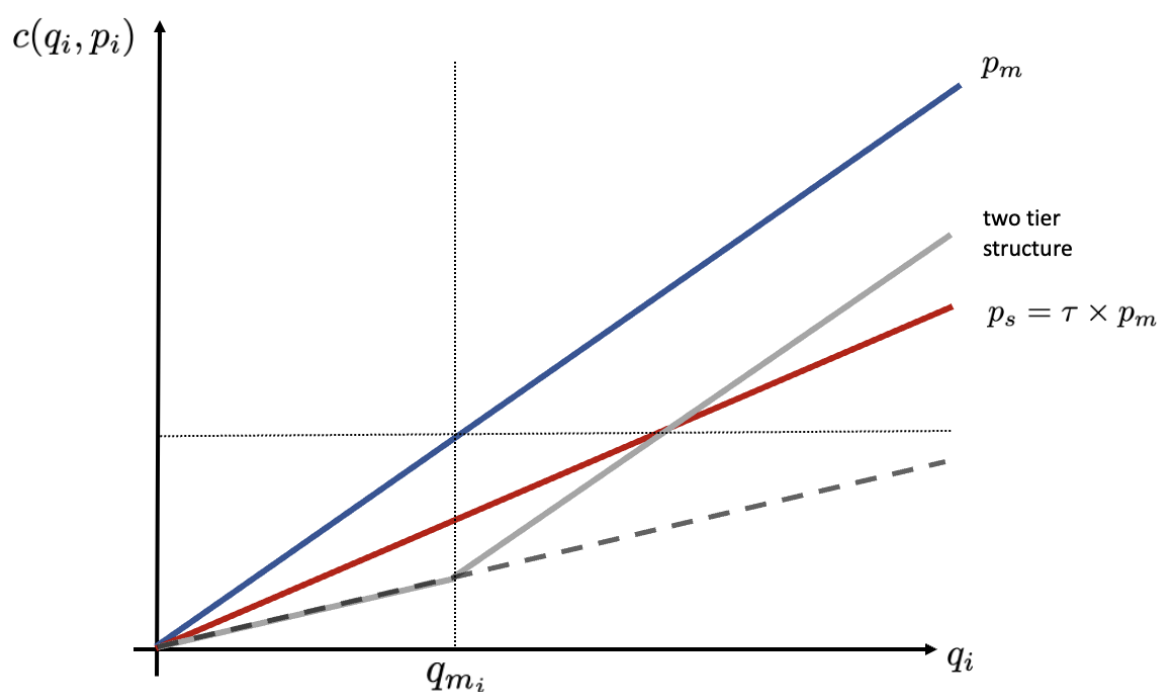
C.3 Additional Tables

Figure A4: Heterogeneity in the relationship between household income, political preferences and energy consumption – proxied via energy bills



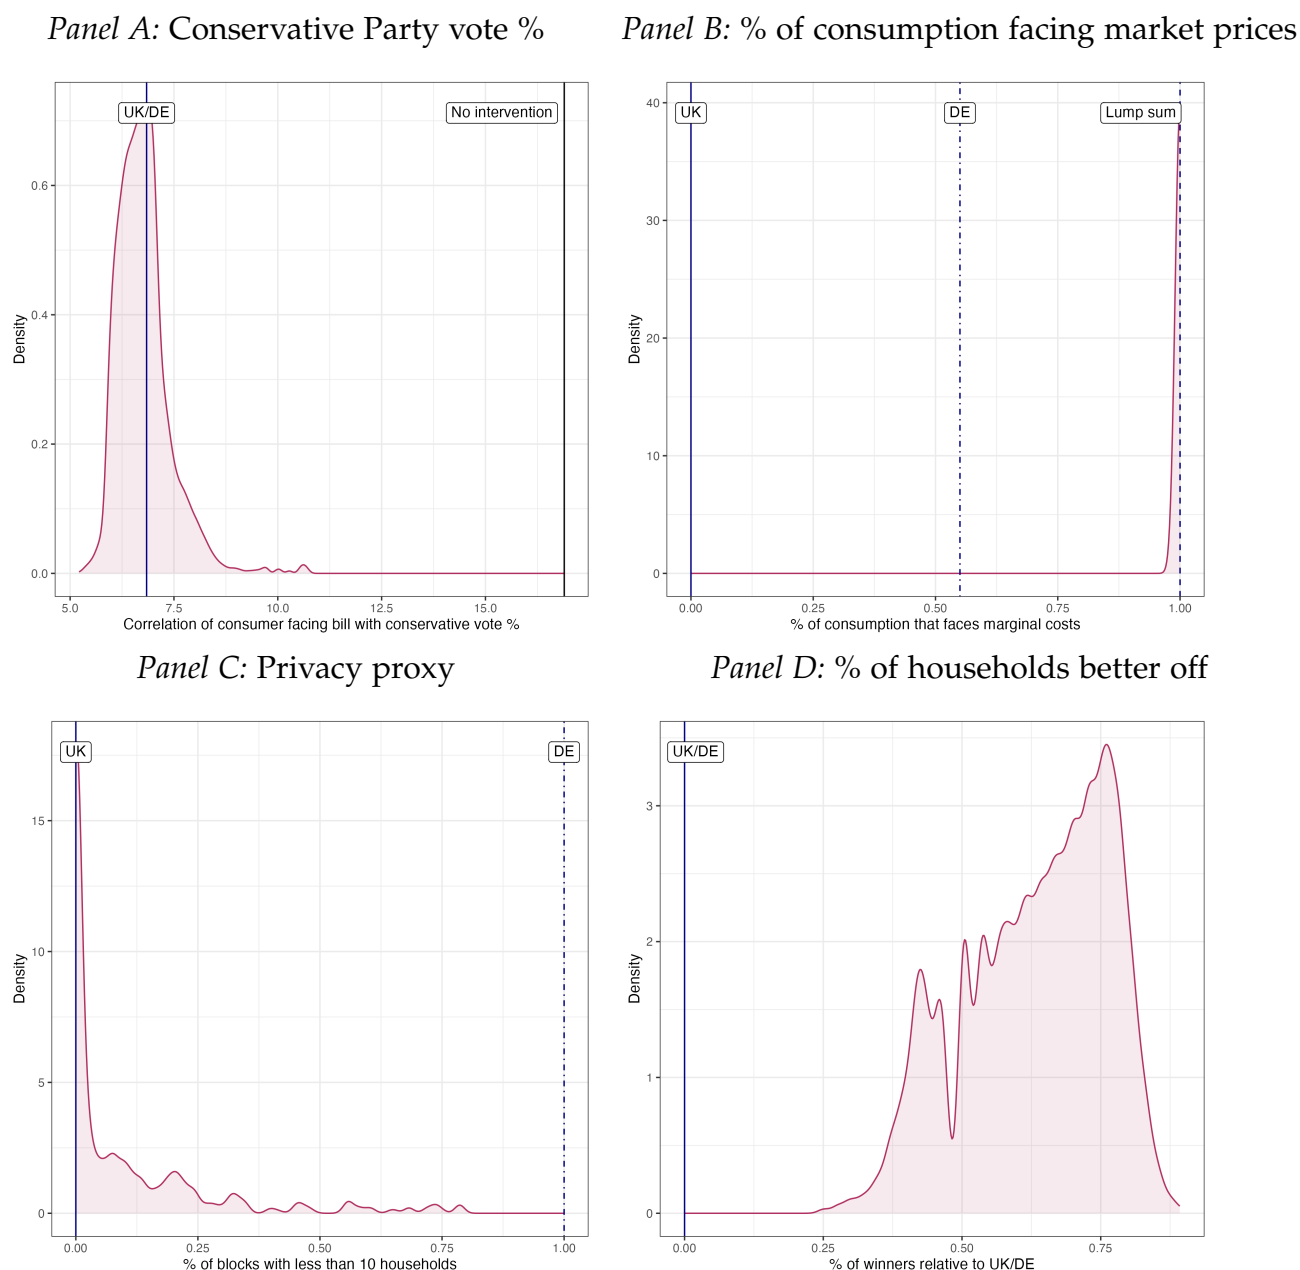
Notes: The figure plots the non-linear relationship between household income, support for the Conservative Party, and energy bills. Each dot represents a point estimate associated with a specific household income percentile. There are 25 point estimates presented. Dots are faintly colored if the specific point estimate is not significant at the 5% level. Dots are solid when they are statistically significant for at least the 5% level. A loess fit across the point estimates is presented to highlight the non-linear nature of that relationship. Each regression controls for the following fixed effects: local authority, household size, year and month of interview, household income percentile, and the survey wave. The figure suggests that energy bills are strongly increasing in absolute terms among Conservative Party supporters across the income distribution, as evidenced by the positive level effect relative to the horizontal dashed line. There is notable heterogeneity, with energy consumption disproportionately higher among very high-income households.

Figure A5: Hypothetical energy bills under different energy tariffs with perfectly price-inelastic energy demand



Notes: The figure plots energy bills for a household i under different energy prices and energy tariffs. The price p_m represents the market price, while p_s , displayed here by the solid red line, represents a subsidized price – a price that is subsidized by a factor of τ , relative to the market price. The dashed gray line represents energy bills under pre-war energy prices. The solid gray line visualizes the energy bills consumers would face under a two-tier tariff structure. There exists a quota q_{m_i} such that for each individual household i the energy costs are given as follows under the two-tier tariff. Up to q_{m_i} , each household pays the subsidized price, p_s . Beyond that level of consumption, the household will pay the market price, p_m .

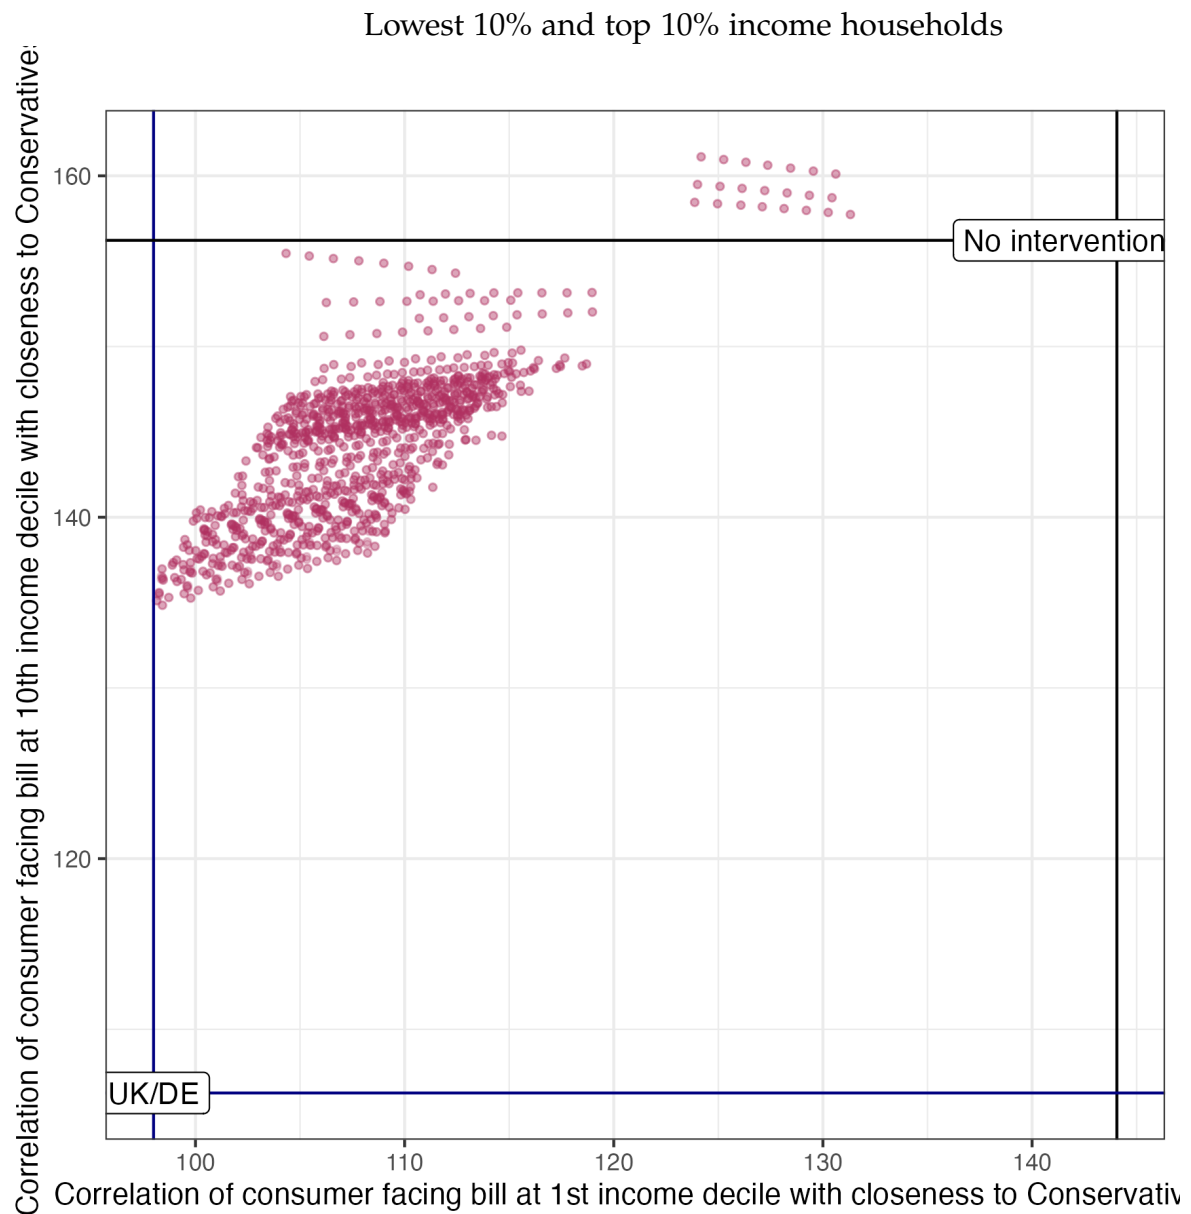
Figure A6: Characterization of the empirical distribution of fiscally neutral lump-sum transfer alternatives vis-à-vis equivalents of the UK and German policy responses, respectively



Notes: This panel of figures displays empirical distributions generated by evaluating a broad range of fiscally neutral alternative lump-sum transfer policies (designed around consumer blocks or archetypes) on a range of metrics – relative to the UK’s *Energy Price Guarantee* and an equivalent two-tier individualized tariff, as was implemented in Germany. Panel A documents the correlation between the net (net of lump-sum transfer amount) energy bills faced by consumers and Conservative Party vote share, using an approach similar to that presented in Table A1. Panel B presents the empirical distribution of the % of consumption that faces market prices under the policy alternatives. Panel C presents the distribution of the share of two-tier tariff blocks that are estimated to include fewer than 10 households, which speaks to the privacy dimension. Panel D presents the empirical distribution of the share of households that, all else equal, would be better off vis-à-vis an individualized two-tier tariff or the untariffed price subsidy, as was implemented via the EPG.

Figure A7: Absence of trade-off between partisanship and degree of unholiness

Differential in net-of-transfer energy bills among Conservative Party supporters by household income for



Notes: This panel of figures displays the empirical distribution that results from evaluating a broad range of fiscally neutral (alternative) lump-sum transfer policies, designed around consumer blocks or archetypes, on a range of metrics vis-à-vis the UK's actual policy, the *Energy Price Guarantee*, and an equivalent two-tier individualized tariff, as was implemented in Germany. Panel A documents the correlation between the net (net of lump-sum transfer payments) bills consumers faced by consumers and Conservative Party vote share, with the correlation derived via an exercise akin to what is presented in Table A1. Panel B presents the empirical distribution of the % of consumption that faces market prices under the policy alternatives. Panel C presents the distribution of the share of two-tier tariff blocks that are estimated to include fewer than 10 households, which speaks to the privacy or informational capacity dimension. Panel D visualizes the empirical distribution of the share of households that, all else equal, would be better off, relative to an individualized two-tier tariff or the untargeted price subsidy, as was implemented via the EPG.

Table A1: Relationship between energy consumption proxies and core Conservative Party electoral support across local elections in England

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Average gas consumption per meter</i>						
Average Natural Gas consumption kwh	0.3204*** (0.0301)	0.3184*** (0.0301)	0.3389*** (0.0231)	0.3498*** (0.0237)	0.2754*** (0.0175)	0.2246*** (0.0152)
Dependent variable mean	20.787	20.787	20.787	20.787	20.787	20.787
R ²	0.18076	0.19198	0.28157	0.34535	0.53262	0.59157
Observations	888,564	888,564	888,564	888,564	888,564	888,564
<i>Panel B: Average electricity consumption per meter</i>						
Average Electricity consumption kwh	1.536*** (0.1175)	1.509*** (0.1151)	1.443*** (0.0956)	1.414*** (0.0887)	1.131*** (0.0702)	0.9823*** (0.0628)
Dependent variable mean	21.810	21.810	21.810	21.810	21.810	21.810
R ²	0.19555	0.21150	0.29301	0.36074	0.52743	0.57941
Observations	1,000,456	1,000,456	1,000,456	1,000,456	1,000,456	1,000,456
Regression specification:						
Area FE	Region	Health care	Enterprise Zone	Travel to work	Local Authority	Constituency

Notes: This table presents results documenting the positive correlation between long-term Conservative Party vote share across local elections and the level of energy consumption. Panel A focuses on natural gas consumption, while panel B studies electricity consumption. Across the columns, we iteratively add more granular area-level fixed effects, moving from least granular region-level controls to most granular. Standard errors are provided in parentheses and are clustered at the district level, with stars indicating *** p < 0.01, ** p < 0.05, and * p < 0.1.