Institute for Fiscal Studies

Peter Levell
Martin O'Connell
Kate Smith

Working paper

25/03

The welfare effects of price shocks and household relief packages: evidence from an energy crisis

Economic and Social Research Council

The Welfare Effects of Price Shocks and Household Relief Packages: Evidence from an Energy Crisis

Peter Levell, Martin O'Connell and Kate Smith*

February 2025

Abstract

Aggregate price shocks can lead to significant inequality in losses both across and within income groups. This creates a trade-off between supporting households through subsidies, which target those most affected but introduce inefficiencies, and transfers, which are less distortionary but harder to target precisely. We develop a framework to quantify this trade-off, using rich panel data on energy spending and income, alongside price and policy variation from the 2022-23 European Energy Crisis. We show that, in the absence of policy intervention, average household welfare losses would have been equivalent to 6% of income, with some households facing much larger losses. A combination of an energy price subsidy and universal transfers reduced both the mean and dispersion of household losses but incurred efficiency costs equivalent to 12% of the total funds spent on the relief package. Under a range of social preferences, better-targeted transfers would have reduced the optimal subsidy rate, though not eliminated it entirely.

Keywords: Energy prices, subsidies, transfers, targeting support, household energy demand
JEL classifications: D12, D31, D61, H20, H31, H53, Q41, Q48
Acknowledgments: This project was supported by the ESRC through the ESRC-funded Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (grant references ES/M010147/1 and ES/W005891/1). We thank Richard Blundell, Andrew Ellis, Lucie Gadenne, Rachel Griffith, Xavier Jaravel, Cormac O'Dea, Krishna Pendakur, Imran Rasul, Thomas Waters and numerous seminar participants for valuable comments. All errors and omissions remain the responsibility of the authors.

*Levell is at the Institute for Fiscal Studies; O'Connell is at the University of Wisconsin-Madison and the Institute for Fiscal Studies; Smith is at the London School of Economics and the Institute for Fiscal Studies.
Correspondence: peter_l@ifs.org.uk, moconnell9@wisc.edu and k.e.smith2@lse.ac.uk.

1 Introduction

How should governments respond to rapid increases in the cost of living driven by shocks to the prices of staple goods? The unequal impacts of inflation,[1] combined with many households' limited ability to smooth temporary shocks, often prompt substantial fiscal interventions. Recent examples include the price subsidies and transfers implemented by more than 25 countries in response to the European Energy Crisis, amounting to nearly 2% of European GDP in 2022 (European Central Bank, 2024). Despite the recurring nature of energy shocks and the widespread use of fiscal measures, we know relatively little about the distribution of households' responses and how this informs optimal policy design.

This paper provides novel evidence on the distribution of welfare losses from a large energy price shock and develops a framework for evaluating alternative policy responses to aggregate shocks. Our approach leverages price and policy variation from the sharp rise in European energy prices in 2022–23. We fully account for behavioural responses in welfare measurement, which we show are substantial and varied. It is common to assume that these responses do not have a first-order effect on individuals' welfare. Although this is a reasonable approximation for small changes, it works much less well when shocks are large. A key advance we make is to model the challenge of efficiently targeting support to the most vulnerable. We show that the incidence of energy shocks is only weakly correlated with income, meaning that conventional redistribution through the tax and transfer system is ill suited to compensating households. Fully accounting for variation in shock exposure and behavioural responses – both across and within income groups – leads to quantitatively different conclusions about the welfare effects of the crisis and qualitatively alters the optimal policy prescription.

We study the UK during the European Energy Crisis, an ideal setting for several reasons. First, retail prices for residential gas and electricity are subject to a binding regulatory cap, adjusted periodically in discrete increments based on international wholesale costs. This creates an exogenous source of price variation, allowing us to estimate consumption responses. Second, the UK government implemented a substantial relief package – costing the equivalent of 1.3% of annual GDP in just 6 months – that combined a price subsidy with direct energy-support transfers.[2] We use this variation to estimate household responses to the most common policy interventions during the crisis. Third, we have access to detailed

[1]See Klick and Stockburger (2024) and Jaravel (2024), and Chen et al. (2024) for recent evidence of heterogeneity in inflation rates in the US and UK. See Jaravel (2021) for a review of evidence on inflation inequality.

[2]In October 2022, the UK government introduced a 39% residential energy price subsidy and a universal £400 "energy-support" transfer to all households, delivered in monthly installments by energy suppliers. The fiscal cost of these measures over October 2022 – March 2023 was approximately 1.3% of annual GDP (Office for Budget Responsibility, 2024).

bank account and credit card data, capturing income and spending for a panel of a quarter of million UK households from 2019 to 2023. A key strength of this dataset is that it enables us to observe energy spending and income at the household level at high frequency, sidestepping the limitations of energy consumption datasets that typically do not include income (Borenstein, 2012).

Together, these features allow us to provide novel evidence on households' exposure to shocks, as well as their responses to price rises and accompanying policy interventions. In the pre-crisis period, energy budget shares varied significantly, with some households particularly vulnerable to price rises. Among households in the bottom income decile, 10% allocated more than 20% of their budget to residential energy. Although energy spending is correlated with income – higher-income households spend more in absolute terms but less as a share of total expenditure – income explains only 7% of the variation in energy spending. This suggests that relying solely on income-based transfers to support households in an energy crisis would leave many facing significant losses.

Households' willingness to switch away from energy in response to price rises mitigates the welfare impact of shocks but raises the efficiency costs of price subsidies.[3] We estimate demand responsiveness by exploiting large, periodic jumps in the regulated energy price, along with our high-frequency consumption data, and detailed seasonal and weather controls. Our results are similar under an alternative specification incorporating data from Northern Ireland, which was not subject to the same regulatory cap. In April 2022, a regulatory cap adjustment led to a 45% increase in the real price of residential energy, triggering a near-immediate average reduction of 14% in household energy consumption – implying an average own-price elasticity of 0.31. We show that households with higher pre-crisis energy spending respond more strongly to price rises, while higher-income households are slightly less price responsive than lower-income households. We thus contribute new evidence to the existing literature[4] on how residential energy demand – both on average and across incomes and prior energy usage – responds to large, salient price increases.

Providing support to households through transfers avoids the inefficiencies associated with distorting price signals but can be difficult to target towards those most affected. Additionally, transfers can introduce other inefficiencies, such as overconsumption due to la-

[3]The mitigating effect of even relatively small substitution elasticities is highlighted in Moll et al. (2023) who study the macroeconomic effects of the cut-off from Russian gas in Germany.

[4]This literature documents a wide range of elasticity estimates (see survey in Labandeira et al. (2017)). Factors driving this variation include differences between short- and long-run responses (Deryugina et al., 2020), the source of price variation (e.g., cross-state price differences in the US (Dergiades and Tsoulfidis, 2008) or changes from non-linear contracts (Reiss and White, 2005)), and households' failure to respond to marginal price incentives embedded in complex non-linear contracts (e.g., Ito, 2014; Shaffer, 2020). Our finding of rapid consumption responses is consistent with households' electricity usage response to the 2000-01 energy crisis in California (Reiss and White, 2008).

belling effects. For households on pay-in-advance energy contracts, we estimate an average marginal propensity to consume energy (MPCE) out of energy-support transfers of 33% – substantially higher than the 4% MPCE we estimate from a non-energy-related government transfer. This finding adds to the literature documenting evidence of a "flypaper effect" (Hines and Thaler, 1995), where agents do not always treat money as fungible.[5] If this reflects distortions to household decision-making – such as mental accounting (Shefrin and Thaler, 1988) – then observed choices may diverge from underlying welfare.

To evaluate welfare effects, we estimate a household energy demand model with two key features. First, we develop a method to assess welfare effects when households make privately suboptimal choices. We extend the choice-based approach to welfare developed by Bernheim and Rangel (2009) and applied by Chetty et al. (2009) in their study of tax salience, to a rich empirical demand model and a setting where price and policy changes are non-marginal. This approach circumvents the need to fully specify the positive model underlying departures from optimal choices, meaning our results are robust to a variety of behavioural biases. Second, we employ a flexible empirical specification based on the EASI demand system (Lewbel and Pendakur, 2009), leveraging the panel structure of our data to incorporate rich preference heterogeneity. The model performs well in predicting, out-of-sample, the distribution of energy demands and its variation with income.

We embed the demand model within a social welfare framework to characterise the policy trade-off between targeting assistance to the most affected households and minimising inefficiencies. Our specification of social preferences allows for both vertical equity concerns – where a given financial loss is more burdensome for lower-income households – and aversion to loss inequality within income groups – conditional on income, hardship may increase non-linearly with loss size. This formulation nests and extends a standard social welfare function by incorporating a proportional version of the equal sacrifice principle (see Fleurbaey and Maniquet, 2018). In our context, this principle is natural, as it frames energy-support policy as aiming to restore households as closely as possible to their pre-shock circumstances. The ideal policy would distribute public funds as personalised lump-sum transfers to equalise losses as a proportion of income. However, *feasible* policies must balance the trade-off between targeting and efficiency costs. We use our framework to quantify the welfare consequences of energy price shocks and evaluate alternative policy responses, drawing four key lessons.

[5] Examples include the allocation of the "Winter Fuel Payment" to energy spending (Beatty et al., 2014), education-support transfers (Benhassine et al., 2015), households' spending response to receipt of SNAP vouchers (Hastings and Shapiro, 2018), and variation in marginal propensities to consume based on mode of transfer delivery (Boehm et al., 2025).

First, energy price shocks can lead to large and highly unequal welfare losses. Without government intervention, the European Energy Crisis would have resulted in an average equivalent variation loss of 6% of after-tax income, with losses reaching 11% at the 95th percentile. Lower-income households were disproportionately exposed: those at the 10th income percentile faced average proportional losses 5 percentage points higher than those at the 90th income percentile. However, variation in losses within income groups is even larger: among the poorest 10% of households, the gap between the 10th and 90th loss percentiles would have been 11 percentage points. To the best of our knowledge, we provide the first comprehensive analysis of the incidence of the 2022-2023 European Energy Crisis that accounts for how households adjusted their behaviour in response to the shock.[6] Failure to account for these behavioural responses overstates the welfare losses from the shock by nearly 60%.

Second, we show that, given the wide variation in losses within income groups, a policy of income-based transfers alone would have led to larger declines in welfare than the UK's implemented package. The UK allocated two-thirds of public funds used to an energy price subsidy, which effectively targets support at households facing large welfare losses – reducing the 95th percentile of losses from 11% to 3% of income. However, this approach created inefficiencies equivalent to 12% of public funds spent on the package. These costs arose from substantial household substitution responses triggered by subsidised prices – this led to excess energy consumption and higher social costs from carbon emissions. Policy-induced inefficiencies are often measured by the total fiscal externality, i.e., the difference between a policy's cost under no behavioural response and its actual cost (see Finkelstein and Hendren, 2020). In our setting, which is characterised by large policy changes, this approach overstates inefficiencies by almost double. This is because households themselves benefit from their behavioural adjustments, offsetting some of the externality this imposes on public costs.

Third, transfers can introduce distortions that amplify other inefficiencies. Had energy-support transfers not induced a flypaper effect for pay-in-advance households, efficiency costs would have been one-fifth lower. Around one-sixth of these costs arise directly from distortions in agents' choices. The remaining share arises from a fiscal spillover associated with increased consumption of a subsidised good. These additional efficiency costs have not been previously documented in the literature, and yet, given that policies rarely operate in isolation, are likely to be important in other settings.

[6] A handful of papers focus on its macroeconomic effects (e.g., Auclert et al., 2023; Pieroni, 2023), or estimate the distributional effects using only pre-crisis spending data (e.g., Bachmann et al., 2022; Fetzer et al., 2023). More broadly, much of the literature on energy price shocks focuses on macroeconomic and average effects (Kilian, 2008).

Fourth, setting transfers as a proportional function of *both* income and past energy usage enables policy to better target households most vulnerable to large proportional losses.[7] Nonetheless, we estimate that fluctuations in energy demand over time are substantial enough to justify a non-zero subsidy even in this case. The optimal policy package – combining a subsidy with transfers based on income and past energy usage – closes more than 60% of the gap in the social value of losses between the first-best personalised transfer policy and the implemented UK response. This approach could have delivered a comparable level of support to the UK's package at 17.5% lower public resource cost.

Our work contributes to a broad literature evaluating the effectiveness of public policy in directing assistance to those most in need. A large body of research examines the equity-efficiency trade-off in redistributing resources to lower-income households (Piketty and Saez, 2013), where the focus is naturally on vertical equity. For instance, Hahn and Metcalfe (2021) show that using energy subsidies to support low-income households can be welfare-reducing. However, policy responses to shocks must account for wide heterogeneity in vulnerability within income groups (see, for example, Borusyak and Jaravel (2024) in the trade context). A nascent literature highlights that such heterogeneity can justify the use of distortionary fiscal instruments to protect households from idiosyncratic shocks, as these instruments can better target the most affected households (e.g., Lieber and Lockwood, 2019; Gadenne et al., 2024). Another strand of research emphasises how heterogeneity in consumption baskets can overturn standard monetary policy prescriptions in response to sectoral shocks (Olivi et al., 2024). A key takeaway from our analysis is that differential vulnerability to aggregate shocks creates a trade-off between targeting and efficiency – one that improved transfer design can mitigate, though not eliminate.

The rest of this paper is structured as follows. Section 2 describes the dataset we use and the distribution of exposure to energy price rises. In Section 3 we discuss pricing in the UK energy market, and present our results on how energy consumption responds to price rises and evidence of a flypaper effect from transfers used to support households. Section 4 outlines our model of household energy demand. Section 5 contains results on the welfare effects of a large energy price shock and evaluates alternative policy responses. A final section concludes.

[7] This approach is feasible for energy support, since energy suppliers retain consumption records. For example, the German and Austrian governments based their support during the European Energy Crisis on past energy usage, as did Brazil in response to a previous energy crisis (Costa and Gerard, 2021). The role of information in shaping government responses in these situations is explored in Fetzer et al. (2024), who argue that the UK policy response could have reduced losses for most households by incorporating additional available data, such as property characteristics or prior energy consumption.

2 Data and Exposure to Energy Price Shocks

In this section we describe our data and provide evidence on the distribution of exposure to energy price rises. We provide supplementary details on the data construction and representativeness in Appendix A.

2.1 Dataset

Our primary dataset contains detailed information on spending and after-tax-and-transfer incomes derived from individual-level bank account and credit card statements. This information is provided by analytics business ExactOne and is collected by ClearScore, a fintech company that helps users monitor and manage their finances.[8] Users who signs up to ClearScore link in their bank account details, from which ClearScore extract their entire transaction histories over the preceding three years. ExactOne then translates these histories into a fully anonymous transaction-level dataset – mirroring bank statements – for just under a quarter of a million UK users over the period 2019 to 2023.

A key advantage of the ExactOne data, compared to similar datasets from individual banks, is that users are encouraged to link all of their accounts and credit cards, including those jointly held with their spouse. We refer to users as households. The dataset gives us a relatively comprehensive picture of households' incomes and spending patterns.

We focus on households that have at least one account that records spending on energy, aggregating the data (including across accounts) to the household-year-month level. We construct measures of monthly energy (gas and electricity combined) spending,[9] monthly non-durable spending, and monthly income. All financial variables are expressed in 2022 prices, unless otherwise specified. Younger individuals and those that live in northern regions of England are over-represented in the data – we show robustness of our results to reweighting our sample along these dimensions to match UK population figures. In the appendix, Figure A.4 shows that the distributions of monthly non-durable and energy spending in the ExactOne data align closely with those in the UK's main consumer spending survey, the Living Costs and Food Survey (LCFS, Office for National Statistics (2024)).

[8]ClearScore provides its services free of charge. It generates revenue by earning commissions from financial product providers when users apply through its platform.

[9]We focus on combined gas and electricity spending as 70% of UK electricity consumers and 80% of UK gas consumers have duel bills, meaning they pay for both gas and electricity together. This means we cannot distinguish between spending on gas and electricity. 99.9% of UK households use electricity and 85% use gas (Department for Business Energy and Industrial Strategy, 2022c).

Analysis sample

Households pay for their energy in different ways. We use these differences to define a sample of households whose spending closely tracks their energy usage.

The majority of households in the UK (nearly 80%) pay for their energy via direct debit, where money is automatically withdrawn from their account each month to pay for their energy usage.[10] Of those paying by direct debit in the ExactOne sample, 91% pay a fixed direct debit, meaning the amount they pay each month is smoothed over the year based on their expected energy usage. Energy suppliers review direct debit payment amounts every 3-6 months, based on households' actual consumption measured from meter readings. For the remaining 9% of direct debit households, their bills correspond more closely to their actual energy usage each period – we refer to these as "variable" direct debits. This applies to households that either have a smart meter, which automatically transmits usage to the supplier, or that send their supplier regular meter readings. In the appendix we show that (i) the energy spending of variable direct debit households is much more seasonal than those on fixed direct debits (Figure A.1) and (ii) the distributions of income, average monthly energy spending, age and region are very similar for the variable direct debit households compared with the fixed direct debit households (Figure A.5). This gives us confidence that the energy spending of variable direct debit households tracks their usage, and that these households are representative of the wider population on fixed billing.[11]

15% of households have a prepayment meter, which operates on a pay-in-advance basis, requiring topping up (either online or in shops) before the credit can be used for energy consumption. The amount of credit that can be held on such meters is limited and prepayment customers typically top up their meters regularly.[12] For these households it is therefore also the case that their spending each month closely corresponds to their energy usage. Prepay households have lower incomes and spend less on energy than those on direct debits.

We distinguish between these different groups in the data based on the payment method they use (e.g., direct debit or card payment) and the variability in their energy spending. We also require that households are present in the data for at least six months over the period of the crisis (June 2021 to March 2023).[13] Our full ExactOne sample (across all billing

[10] 78% of households in the LCFS use direct debit to pay their gas, electricity, or combined bills.

[11] Approximately 7% of households pay via standard credit, which means they only pay for energy once they have received a bill for their actual use (either monthly or quarterly). We group these together with the fixed direct debit households.

[12] Prepay households in the ExactOne data top up their meters more than 5 times per month, on average, and 85% of those that recording spending on gas or electricity prepayment meters in the Living Costs and Food Survey report expecting their payment to cover one month or less.

[13] In almost all our analysis, we exclude observations from Northern Ireland, as it is covered by a separate energy regulator and has different policies to the rest of the UK (which, for expositional simplicity we abbreviate to "the UK"). The only exception is a robustness check for our estimated price elasticities – further details are provided in Appendix B.2.

types) comprises over 250k households from January 2019 to December 2023. Our analysis sample (those that pay for all their energy – gas and electricity – either through variable direct debit or prepayment) covers 1.2m observations for 75k households over the period June 2021 to December 2023. Both the full and analysis samples contain a higher proportion of prepay households than the UK population (25% in the full sample, vs. 15% as measured in the LCFS). We therefore reweight to account for this difference in all our analysis. Figure 2.1 validates our data and sample construction by showing that the trends in energy spending over the crisis in the ExactOne data (analysis sample) accord with those from the UK National Accounts.

Figure 2.1: *Trends in energy spending in ExactOne data and National Accounts*

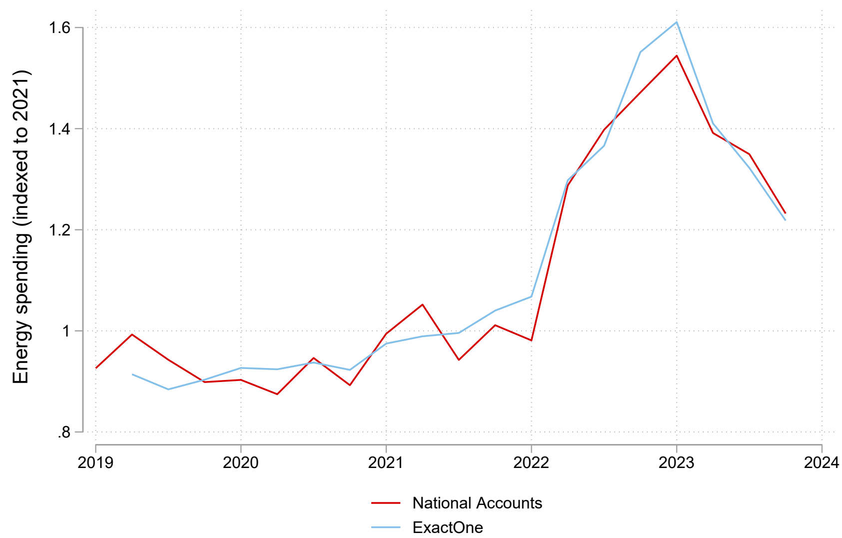

Notes: The figure shows deseasonalised energy spending (indexed to 1 in 2021) in the National Accounts and ExactOne data. We deseasonalise the ExactOne data by subtracting calendar month effects estimated using the pre-crisis period (2019 and 2020); we use the deseasonalisation of the National Accounts provided by the Office of National Statistics. The ExactOne line is estimated on the analysis sample, controls for household fixed effects and reweights by payment type, age band and region to the UK population. Spending is expressed in 2022 prices.

2.2 Pre-Crisis Energy Spending

In Figure 2.2 we describe energy spending over the two year period prior to the crisis (2019-20). In the appendix, Figure A.6 shows that the patterns are very similar when focusing on 2019 alone, meaning Figure 2.2 does not reflect patterns specific to the COVID-19 pandemic.

Panel (a) shows how the mean, interquartile and interdecile ranges of household-level average monthly energy spending vary with percentiles of the household income distribution. It shows that higher-income households tend to spend more on energy: those in the top income decile spent, on average, £143 per month, while those in the bottom decile spent £95 on average. However, variation in income accounts for only 7% of the total cross-sectional variation in energy spending.[14] There is a large degree of variation in energy spending conditional on income, across the income distribution. For instance, 10% of households in the bottom income decile spent, on average, more than £154 per month on energy and 10% of those in the top decile spent more than £226. These patterns remain similar when we describe energy spending conditional on total non-durable expenditure, rather than income (see Figure A.6 in the appendix).

Figure 2.2: *Energy spending across the income distribution*

(a) Average energy spending (b) Energy budget shares

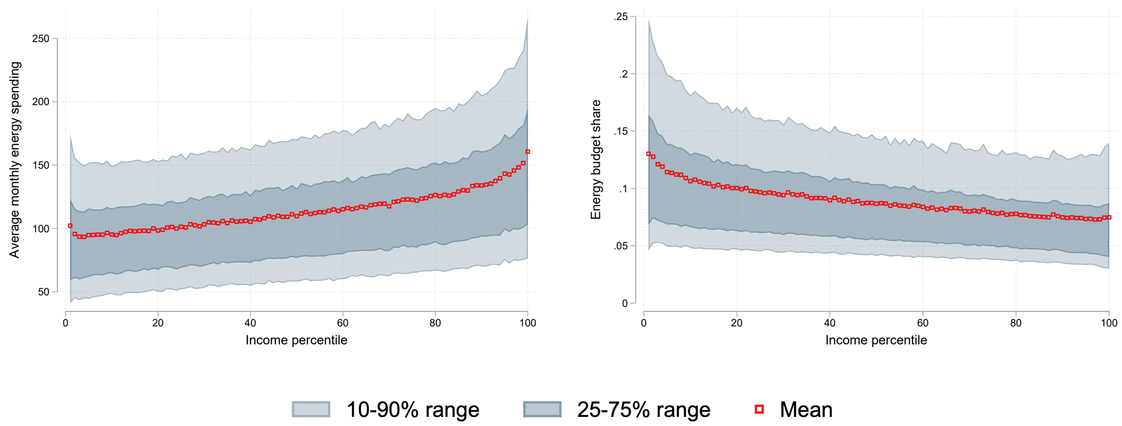

Notes: Panel (a) summarises the distribution (mean and 10th, 25th, 75th, 90th percentiles) of households' mean monthly energy spending in the pre-shock years (2019 and 2020) conditional on income (also measured in 2019 and 2020). Panel (b) shows the analogous figure for energy budget shares (energy spending over non-durable spending). Spending is expressed in 2022 prices.

Panel (b) shows how energy budget shares (i.e., energy spending as a fraction of total non-durable expenditure) vary across the income distribution. Although higher income households tend to spend more on energy, those at the bottom of the distribution tend to spend a much greater *fraction* of their budget on energy. Households in the bottom income decile allocated, on average, around 12% of their non-durable spending to energy, compared with those in the top decile, who allocated 7%. Lower-income households are therefore more exposed to energy price rises. Among lower-income households, the degree of exposure varies substantially – 10% of households in the bottom income decile spent more than 20% of their budget on energy – while some better-off households are also highly exposed to prices increases. This evidence suggests that income-based transfers alone are

[14]The R-squared of a cross-sectional regression of monthly energy spending on indicator variables for income percentiles is 0.067. We obtain a similar estimate when we conduct this exercise using the Living Costs and Food Survey. The LCFS also allows us to estimate the explanatory power of *equivalised* household income, which has a lower R-squared of 0.028. Adding to the unequivalised income specification variables that capture the age of the head of the household, the number of adults, number of children, number of rooms in the house and the region increase the overall R-squared to 0.19: there remains a substantial amount of unexplained variation in energy spending.

likely to have limited effectiveness in supporting the most exposed households when energy prices rise.

Another way to target transfers to households most exposed to price rises is to base them on historical energy usage. The effectiveness of this approach depends on the persistence of energy spending over time. If high energy use today is only weakly predicted by past consumption, transfers based on previous energy usage may leave many households exposed. We estimate a one-year autocorrelation coefficient for logged energy spending of 0.74, with an R-squared of 0.55 (see Appendix A.4 for details). This suggests a relatively high degree of persistence. However, it also points towards significant changes in some households' energy needs over time, which mean that transfers tied to historical usage may fail to support some households that are vulnerable to future price increases.

3 Evidence on Household Responses to an Energy Crisis

3.1 Energy Price Shock

UK energy prices

Global demand for energy surged following the end of COVID-19 lockdowns and as political tensions with Russia, then the world's largest exporter of natural gas, gradually worsened. This led to significant increases in wholesale gas prices in Europe in the winter of 2021 and early 2022, when Russia launched a full-scale invasion of Ukraine. These increases had a particularly large effect on energy prices in the UK, owing to its relatively strong dependence on natural gas for both residential heating and electricity generation.

A distinctive feature of the energy market in the UK is a price cap, which is administered by the energy regulator Ofgem and sets maximum values that suppliers can legally charge residential consumers for electricity and gas. Residential energy bills in the UK typically consist of unit charges for electricity and gas (charges per kilowatt-hour of use) and fixed standing charges that are independent of use. The price cap applies to both unit prices and standing charges, with different values for gas and electricity. For the majority of households the fixed fee standing charge comprises only a relatively small fraction of their overall bill – for instance, standing charges for gas and electricity accounted for 13% of the average household's bill over 2021-2023.[15] Non-linear marginal pricing, such as increasing block pricing where the marginal price increases in usage, is not a feature of the UK mar-

[15] There is evidence that when consumers face two-part energy tariffs, and there are large changes in the balance between the fixed fee and marginal price components, they respond to average rather than marginal prices changes (e.g., Labandeira et al. (2022) and Ito and Zhang (2024)). In our context the fixed fee is relatively low and stable.

ket.[16] The stated aim of the cap is to limit firm profits and to prevent those households that do not shop around for cheaper tariffs from overpaying. The cap level is based on Ofgem's estimates of supplier costs and any change is announced approximately one month before taking effect.[17]

Figure 3.1: *Energy price cap, energy price guarantee and cheapest available tariffs, 2019-2023*

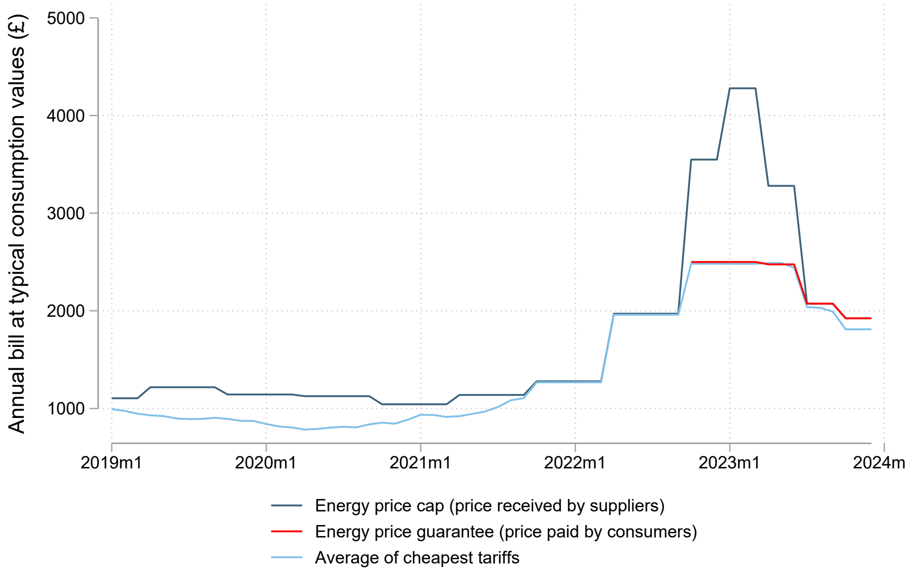

Notes: Data from Ofgem (2023). Figures are costs of an annual bill at 'typical' consumption values of gas and electricity (12,000kWh of gas and 2,900kWh of electricity) for dual fuel direct debit consumers (in nominal terms). The average of the cheapest tariffs is a simple average of direct debit tariffs from the 10 suppliers with the lowest cost tariffs (that is, only including one tariff per supplier), including fixed tariffs. Only tariffs that are generally available to consumers are included. There are small regional differences in the value of the price cap across the UK – the figure shows the average value. We account for regional variation in all our analysis.

As wholesale energy prices rose rapidly, the price cap became increasingly binding on suppliers. Figure 3.1 shows the value of a bill at capped electricity and gas prices for Ofgem's definition of a typical consumer, alongside the average annual bills this consumer would pay under the 10 cheapest tariffs on the market. In January 2020, the average of the 10 cheapest tariffs was around 70% of the cost of a bill at the default tariff cap. By January 2021, this had risen to 90%, and from October 2021 it had risen to 99%. At this point, the vast majority of consumers were paying the maximum prices specified by the cap. The fact that the cap binds in this way is useful for us, as it means we know with a high de-

[16]The majority of UK households (86% in 2012; Department for Energy and Climate Change (2013)) face a unit charge for electricity that is fixed throughout the day, while the remainder have contracts that entail a different marginal price for day-time and night-time consumption.

[17]Until October 2022, Ofgem updated the cap every 6 months to reflect changes in suppliers' costs (mostly driven by wholesale price changes). After October 2022, Ofgem updated the cap every 3 months as a response to the increasingly volatile wholesale market.

gree of certainty the tariffs consumers paid for the energy they consumed from mid-2021 onwards.[18]

Rises in the price cap led to increases in the average real cost of an energy bill at typical consumption values of 8% in October 2021 and a further 45% in April 2022.[19] The price cap rose again in October 2022, leading to a further 73% increase in the real price of energy, at which point the UK government introduced policies designed to support households. A subsequent increase in January 2023 meant that, without government intervention, households would have faced a price more than four times higher than two years earlier.

Policy intervention

In September 2022, shortly before the October price cap rise, the UK government announced the "Energy Price Guarantee" (EPG), a price subsidy for energy, and the "Energy Bill Support Scheme", a universal transfer. Anticipating that the energy price shock would likely persist, the government committed to subsidising energy prices for up to two years by pledging to cap the annual cost of the typical consumer's energy bill at £2500 over this period. It justified these measures as providing "urgent support" to "millions of families" (Department for Business Energy and Industrial Strategy, 2022a), reflecting a concern that in their absence there would be widespread hardship.[20]

The EPG capped the unit rates and standing charges on gas and electricity that suppliers could charge consumers, and paid a subsidy to suppliers to make up for the difference between the EPG rates and the energy price cap set by Ofgem. The red line in Figure 3.1 shows the cost of the typical bill under the EPG. The EPG limited the October 2022 rise in the real energy price that households faced to 22%, and protected households from a further increase in the energy price cap in January 2023. The EPG implied an average subsidy of

[18] A potential complication in using the prices set by the Ofgem cap to infer prices paid, and calculate quantities consumed, is that the cap does not apply to tariffs on fixed price contracts that consumers may have agreed before the energy price cap increases occurred. These are much less prevalent among prepay consumers, at less than 1% for both electricity and gas tariffs (Department for Business Energy and Industrial Strategy, 2022b). However, we lack data on the share of variable direct debit consumers on fixed contracts over the crisis. We therefore check the robustness of our estimates of the energy price elasticity by focusing on a subsample of households who switch suppliers between June 2021 and January 2022 (Figure 3.1 indicates that new contracts after this point are set at the cap level), and show that the elasticities for this group are very similar to the wider sample (see Table 3.1).

[19] The price cap varies slightly by payment type (prepay or direct debit) and region. Figure 3.1 shows the cap for direct debit households at prices averaged across regions, while Figure B.1 in Appendix B.1 shows the cap for prepay households. In all our analysis we account for price differences by payment type and region.

[20] Implicit in this is the idea that the government is better positioned to smooth temporary shocks than households, as many households face constraints that can lead to large consumption falls even in response to temporary income reductions or price shocks. Kaplan et al. (2014) estimate that around 35% of UK households are "hand-to-mouth," meaning the value of their liquid assets is less than half their income; the majority of these households are "wealthy hand-to-mouth" because they own illiquid assets but hold limited cash on hand.

39% applied to the marginal price of energy over the period October 2022 to March 2023.[21] Although the EPG reduced unit prices, it left fixed fee standing charges for both electricity and gas essentially unaffected. Prices fell following the end of the EPG in June 2023, with a 23% reduction in the real energy price in July and a further 6% reduction in October.

The bill support scheme provided all households with transfers of £400 paid in monthly installments over October 2022-March 2023. We return to them in Section 3.3.

Measuring the consumer price of energy

We use the regulatory cap (or EPG, when it is in place) to measure the marginal (unit) prices of electricity and gas, which we combine into a price index that measures the marginal price of energy consumption. The weights in the index are the average share of energy spending allocated to gas and to electricity (over 2019-2020). We use this price index to measure the response of energy usage to price rises.

There are two possible concerns associated with the use of this price index that arise because regulatory cap changes differentially affect the price of gas and electricity. The first is that a fixed weight index may fail to account for substitution *between* electricity and gas in response to these differential price changes, since it is consistent with a Leontief sub-utility function over electricity and gas. In Appendix B.1 we show that the aggregate quantity-shares of gas and electricity are stable over time (at around 35% for electricity). This means that using a price index that allows for substitution between gas and electricity makes very little difference to our results.[22] The second potential issue is that we use weights that are common across households, and yet there may be variation in gas and electricity shares across different types of households. To address this we also construct price indexes using household-varying weights (based on income and energy spending and measured using the LCFS), and show robustness of our price elasticity estimates to using these household-level price indexes.

For each household-month we measure the quantity (or usage) of energy by dividing a household's monthly energy spending (net of the fixed standing charge) by the price index. The resulting measure is a quantity index (rather than, say, kilowatts of usage). This presents no problems for estimating price elasticities or household-level welfare effects. However, in Section 5, when we discuss the implications of policy responses to the energy

[21] Unit prices for electricity under the EPG from October 2022 to March 2023 were on average 42% below the cap price, while unit gas prices were on average 35% lower. Weighting by average shares of spending on electricity and gas from 2019 yields the 39% average subsidy. Over the whole period of the EPG's operation (from October 2022 to June 2023) the average implied subsidy was 35%.

[22] In Appendix B.1 we shows that our fixed weight – Laspeyres – index comoves with a Fisher index. The Fisher index is a second-order approximation to an arbitrary homothetic expenditure function (Diewert, 1976) and allows flexibly for substitution effects. The two indexes co-move due the stability of the quantity shares of gas and electricity over time.

crisis for carbon emissions, we require an additional assumption about how the quantity index translates into carbon emissions.

The evolution of energy spending

In Figure 3.2 we show changes in deseasonalised log energy expenditure over the period June 2021 – June 2023. The vertical dashed lines indicate periods when the value of the energy price cap was updated and, in April 2023, when the provision of transfers to help with energy bills ended. The figure indicates that spending increased at the 8% and 45% energy price cap rises, but by less than the increase in price, implying inelastic energy demand responses. The spending changes in response to the introduction and withdrawal of the transfers indicates an economically meaningful marginal propensity to consume energy out of these funds.

Figure 3.2: *Log energy spending over the crisis*

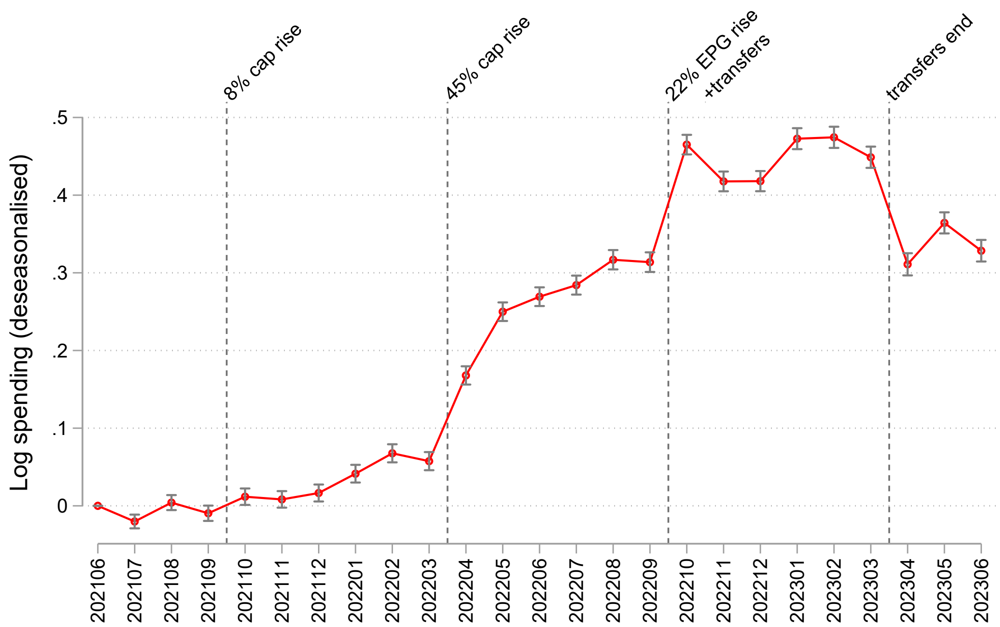

Notes: The figure shows total log deseasonalised energy spending in each month. We deseasonalise by subtracting calendar month effects estimated using the pre-crisis period (2019 and 2020). Spending and cap rises are expressed in 2022 prices; the cap rises in nominal terms are 12% (October 2021), 54% (April 2022), 27% (October 2022).

3.2 Price Elasticities

To estimate how household energy use responds to price changes we focus on the period June 2021 to September 2022. During this time, the vast majority of households faced energy prices equal to the price cap. Here we exclude the period following the cap change in

October 2022, as this price change coincides with the introduction of the transfers. We return to these below.

We estimate variants of the following equation:

$$\log o_{i\tau} = \sum_{d=DD,PP} \left(\gamma_d \log p^e_{r(i)\tau} + g_d(\text{temp}_{a(i)\tau}, \text{rain}_{a(i)\tau}) \right) + \zeta_i + \iota_{m(\tau)} + \text{adjustment}_\tau + \epsilon_{i\tau}. \quad (3.1)$$

$o_{i\tau}$ is the outcome variable, either energy spending after subtracting the fixed standing charge ($x^e_{i\tau}$) or quantity ($e_{i\tau}$) for household i in year-month τ. $p^e_{r(i)\tau}$ denotes the marginal price of a unit of energy in period τ. Note this varies across households depending on their region, which we index by r. $g_d(\text{temp}_{a(i)\tau}, \text{rain}_{a(i)\tau})$ is a flexible function of minimum and maximum monthly temperatures and rainfall in local area a;[23] ζ_i is a household fixed effect; $\iota_{m(\tau)}$ denote calendar month effects; and adjustment$_\tau$ denotes indicator variables for the months immediately before and after each cap change, which control for any anticipation or lags in adjustment. We allow the price coefficients γ_d to vary across the household's payment method – specifically, whether they pay by direct debit (DD) or prepayment (PP).

For small price changes, the coefficient capturing the effect of a change in log price on log quantity (or the coefficient minus one, when log spending is the outcome variable) approximately coincides with the percent change in quantity associated with a one percent change in price. For large price changes this is not this case. We therefore convert our estimates into the implied percentage impact on quantity of the price change in April 2022, divided by the percent (45%) real price increase (which we refer to as the elasticity). We calculate separate elasticities for households using each payment method, and then average across these using weights calculated from the nationally representative LCFS (85% on direct debit, 15% on prepay). Table 3.1 summarises our results.

The first column shows the estimated elasticity using deseasonalised log spending as the dependent variable, controlling for a flexible function of local weather conditions. We estimate that the quantity of energy demanded fell by 14%, on average, when the real price of energy rose by 45% in April 2022. This implies an own-price elasticity of energy of -0.31. The second column shows we obtain a similar estimate when using log quantity as the outcome variable, and the third column shows a comparable estimate using raw (un-deseasonalised) spending with controls for calendar month effects. The fourth column shows that the results are robust to weighting by age and region. The fifth column demonstrates robustness of the estimate to instead using a price index constructed with gas versus electricity weights that vary across households, based on quintiles of income and average pre-shock energy spending.

[23]We use weather data collected at the local area level by the Met Office, see Appendix A.2 for details.

Table 3.1: *Energy price elasticities*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
% change in e over cap change	-14.2	-13.7	-14.9	-15.8	-14.4	-12.6	-14.4
	(0.5)	(0.5)	(0.6)	(0.7)	(0.5)	(1.3)	(0.8)
Implied own-price elasticity	-0.311	-0.301	-0.328	-0.346	-0.316	-0.278	-0.316
	(0.011)	(0.011)	(0.013)	(0.015)	(0.011)	(0.029)	(0.018)
Dependent variable	$\log x^e$	$\log e$	$\log x^e$	$\log x^e$	$\log x^e$	$\log x^e$	$\log x^e$
Deseasonalised	Yes	Yes	No	Yes	Yes	Yes	No
Calendar month effects	No	No	Yes	No	No	No	Yes
Reweighting by age & region	No	No	No	Yes	No	No	No
HH-specific index weights	No	No	No	No	Yes	No	No
Restricted to new supplier	No	No	No	No	No	Yes	No
Incl. NI + year-month effects	No	No	No	No	No	No	Yes

Notes: Standard errors in parentheses and clustered at the household level. Each column shows the percentage change in energy quantity, e, over the cap change and the implied elasticity from the estimation of (3.1). Dependent variable $\log x^e$ *and* $\log e$ *refer to log of energy spending and quantity respectively. We calculate the % change in quantity for each payment type, d, as follows:* $\Delta e_d / e_d = \frac{\exp(\tilde{\gamma}_d \log p_1) - \exp(\tilde{\gamma}_d \log p_0)}{\exp(\tilde{\gamma}_d \log p_0)}$*, where* $\tilde{\gamma}_d = \hat{\gamma}_d - 1$ *if the dependent variable is* $\log x^e$ *and* $\tilde{\gamma}_d = \hat{\gamma}_d$ *if the dependent variable is* $\log e$*. The elasticity is then given by* $\epsilon_d = \frac{\Delta e_d / e_d}{(p_1 - p_0)/p_0}$*.* p_1 *denotes the price after the April 2022 increase and* p_0 *the price in March 2021. We weight the* $\Delta e_d / e_d$*s and* ϵ_d*s using the shares of each payment type in the population to get the average elasticity. All specifications include weather controls (*5^{th}*-order polynomials in local monthly minimum and maximum temperature, the squared difference between maximum and minimum temperature and local monthly rainfall) and dummies for the months immediately before and after a cap change.*

One potential concern with these estimates is any bias that may arise from households remaining on fixed-price contracts during the energy crisis and, consequently, not facing the cap rises. Due to escalating costs, several energy suppliers declared bankruptcy during the second half of 2021, resulting in their consumers being automatically transferred to the standard tariff of an alternative supplier (which entailed a price equal to the cap). Column (6) shows that restricting the sample to households who switched to a new supplier between June 2021 and January 2022 yields similar results.

Another potential concern with our estimates is that there may be time-varying factors, not fully accounted for by our seasonal and weather controls, that correlate with changes in the regulatory price cap. To address this, we estimate an alternative specification that incorporates energy spending data from Northern Ireland, which was not subject to the same regulatory price cap as the rest of the UK. In Appendix B.2, we show that energy prices increased much more gradually in Northern Ireland, diverging from those in the rest of the UK following the cap rises. This specification adds a full set of year-month dummies to equation (3.1) and uses differential geographic price variation. The final column of Table 3.1 shows that this results in a very similar estimate of the average price elasticity.

Heterogeneity in price responsiveness

Variation across households in their response to price changes is one important driver of the distribution of welfare effects due to energy price shocks. In Figure 3.3 we show how households' consumption responses to the April 2022 price increase varied across groups based on the pre-shock energy spending quintile and income quintile they belonged to. To do this we estimate the column (1) variant of equation (3.1) where we additionally interact the coefficient on log price with indicators for income quintile $\times$ pre-shock energy spending quintile (25 groups total). We calculate the elasticity for each of these 25 groups, and plot the results in Figure 3.3.[24] It shows that those households that consumed a lot of energy in the pre-crisis period have larger price elasticities, with the average elasticity varying from -0.24 for the bottom to -0.37 for the top spending quintile. It also shows that those in the highest income quintile have less elastic demand, *conditional* on pre-crisis energy consumption.

Figure 3.3: *Heterogeneity in energy price elasticities, by pre-shock energy spending and income*

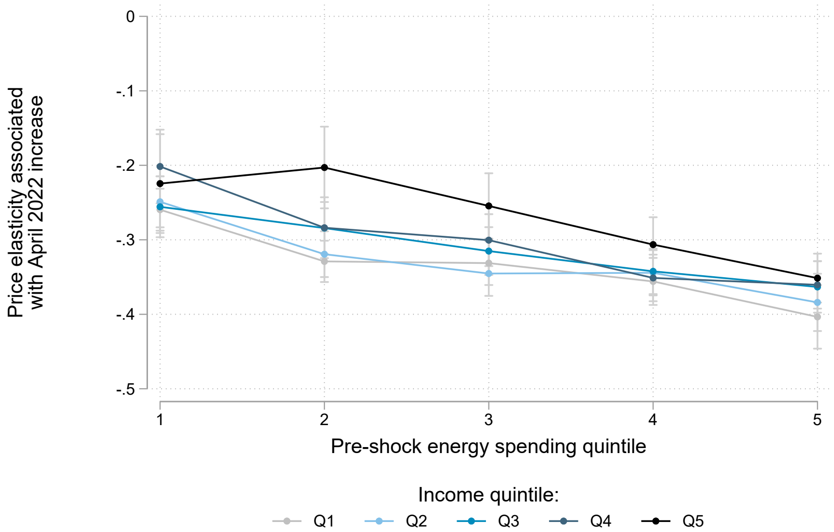

Notes: Figure shows the estimates of the price elasticity over the April 2022 price increase, across income quintile and pre-shock energy spending quintile. We allow for heterogeneity in price responsiveness across payment type and average these using weights that vary across income and pre-shock energy consumption quintile.

The results we report in Table 3.1 and Figure 3.3 do not separate out substitution effects (switching to alternatives to energy in response to relative price changes) from income effects (lowering energy consumption due to the reduced purchasing power that results from higher energy prices). The extent to which consumption responses are driven by substitution or by income effects is important for understanding the welfare effects of an energy

[24] In Appendix B.2 we show that the patterns shown in Figure 3.3 are robust to using a price index constructed with gas and electricity weights that vary across the 25 groups (see Appendix B.1 for details).

price shock and the efficiency costs of policy responses. We turn to this in Section 4, where we estimate a structural model of energy demand.

3.3 Marginal Propensity to Consume Energy Out of Transfers

A major part of the UK government's household relief package during the European Energy Crisis was the provision of £400 energy-support transfers given in monthly installments from October 2022 to March 2023. There were three main ways that households received these: (i) as vouchers (or automatic meter top-ups) for prepayment energy, (ii) as credit added to direct debit accounts and (iii) as direct cash refunds. These were all administered by households' electricity suppliers. The UK government also provided direct cash transfers to households in receipt of means-tested benefits from summer 2022 until spring 2024.[25] In this section, we estimate households' marginal propensities to consume energy (MPCE) out of transfers, and investigate whether this is different for the energy-support transfers compared with direct cash transfers.

Although the energy-support transfers were administered by energy companies and distributed to many households as either top-ups on meters or vouchers, their impact on households' budgets sets was similar to the receipt of an equivalent cash amount. This is because the vast majority of households (over 90%) spent more on energy than the value of the transfer immediately before their introduction. Despite this, there is a body of evidence documenting the existence of a "flypaper effect", where money "sticks where it hits" (Hines and Thaler, 1995).[26] A variety of underlying models may explain this empirical phenomenon, including models of mental accounting, which predict that the effects of transfers on consumption may differ according to how they are described and delivered (Shefrin and Thaler, 1988). In our setting, payments that households perceive as part of an energy "account" may be associated with a higher marginal propensity to consume energy than payments not viewed in this way.

We focus here on the impact of transfers on households that prepay for their energy, and who therefore received the transfers either in the form of vouchers or automatic top-ups on their meters.[27] Most (83%) of these households spent more on energy than the value of

[25] "Cost of living payments" were paid directly into recipients' bank accounts in six installments (with values ranging from £299 to £326) from the summer of 2022 to the spring of 2024.

[26] The "flypaper effect" was first coined by Okun, who used it to refer to the tendency of federal grants to stimulate local government spending more than an equivalent increase in local income. The term in now commonly used to describe in the non-fungibility of money in the context of household and firm decision-making (see Jacoby, 2002; Choi et al., 2009; Fafchamps et al., 2014).

[27] Vouchers and credits provided to pre-pay households could be redeemed against both electricity and gas consumption. Recipients of vouchers had the flexibility to add them to either their electricity or gas meters, or split them between the two. For households with smart pre-payment meters, payments were either split equally between their electricity and gas accounts or added by default to their electricity smart meters. These households could, however, request their supplier to transfer credit to their gas meters. Exceptions were British

the transfer in the month before their introduction, when the price was 27% lower. Even for those who received a voucher and who wished to consume energy of less value than the voucher, they could have added the balance to the meter or redeemed the voucher after the end of the transfer period (up until June 2023). It is therefore unlikely that the transfers acted as a binding constraint forcing households to consume more energy than they otherwise would.

Figure 3.4: *Energy spending over the transfer period, prepayment households*

Notes: The dashed vertical lines indicate the start and end of the period when energy-support transfers were issued. The figure shows two series for deseasonalised log energy spending for prepayment customers. The series are deseasonalised using month effects estimated in 2019-20. The solid red line shows the evolution of actual deseasonalised log spending. The dashed line shows log spending as predicted by a fixed-effects panel regression, that controls for weather effects, estimates price effects using variation outside of the transfer period only and controls for (and removes) time effects for the transfer period.

Figure 3.4 shows that gross energy spending (i.e., inclusive of the energy-support transfer value) for prepayment households increased significantly during the transfer period, followed by a decline once the transfers ceased. The dashed line shows the increase in spending that would have occurred based solely on price responses during this period. Spending is substantially higher during the transfer period than expected based purely on the price responses exhibited by households during previous price changes. Although there was some limited scope for these households to have accumulated credit on their meters (or cashed their vouchers after the transfer period), there is little evidence of a significant drop in spending relative to its predicted level in the months following the transfer period.

Gas prepay customers, or households who had different suppliers for electricity and gas, who were unable to transfer credit to their gas meters. (Note that 98% of prepay households had one energy supplier). However, as credits could be utilised until the end of June 2023, even for these consumers, the credits were likely inframarginal. Furthermore, excluding households who used British Gas, or who had two different energy suppliers, from our analysis in this section makes no material difference to our results.

We formalise this by estimating the following equation:

$$\begin{aligned}\log x^e_{i\tau} =& \gamma \log p^e_{r(i)\tau} + g(\text{temp}_{a(i)\tau}, \text{rain}_{a(i)\tau}) \\ &+ \mu^E \text{transferE}_\tau + \mu^{EP} \text{posttransferE}_\tau + \mu^G \text{transferG}_{i\tau} + \zeta_i + \epsilon_{i\tau} \end{aligned} \tag{3.2}$$

where transferE_τ is an indicator variable equal to 1 during the six month period that the transfers were distributed by energy companies (October 2022 to March 2023), $\text{posttransferE}_\tau$ is an indicator variable equal to 1 for April 2023, and $\text{transferG}_{i\tau}$ is an indicator equal to 1 if household i received a transfer directly from the government (which was badged as a general "cost of living payment") in month τ or $\tau - 1$ – accounting for the possibility that transfers received at the end of the month may affect spending in the following month. All other variables are defined as in equation (3.1). We estimate equation (3.2) using the period June 2021 (when the price cap became binding for the majority of households) to December 2023.

Using the estimated coefficients, we calculate the implied marginal propensity to consume energy (MPCE) out of transfers, which we report in Table 3.2. We estimate an MPCE out of the energy-support transfers of 0.33 (column (1)). This is substantially higher than the average budget share of energy (14%) for pre-pay households over this period. This estimate accounts for any reductions in spending in the month after the transfer period ends due to accumulated credit; column (2) shows that not adjusting for these spillover effects leads to similar estimates of the MPCE, which is consistent with Figure 3.4.[28] The MPCE out of energy-support transfers compares with an estimated MPCE of 0.04 from the cash transfer administrated directly by the government. This difference is evidence of a pronounced flypaper effect associated with the energy-support transfers paid to pre-payment households.

In Appendix B.3 we show results for households that pay by direct debit, and therefore received the transfer as credit on their account or a cash refund. We find no evidence of flypaper effects for those receiving transfers as cash. Although there is some evidence of a flypaper effect for those households that received the transfer as credit, the effect is

[28]The MPCE in the first row of column (2) is $\frac{d\mathbb{E}[x^e_{i\tau}]}{d\text{transferE}_\tau}$, which we calculate by evaluating $\hat{\mu}^E \times \exp(\log \hat{x}^e_{i\tau} + \mathbb{E}[\exp(\hat{\epsilon}_{i\tau})])$ and dividing by the 0.9 times the transfer value in each month, where $\log \hat{x}^e_{i\tau}$ is the predicted value of $\log x^e_{i\tau}$ evaluated at mean values of the covariates in the transfer months and $\hat{\epsilon}_{i\tau}$ is the regression residuals for observation over this period. We scale the transfer value by 0.9, as official statistics show 90% of prepayment vouchers were redeemed (Department for Energy Security and Net Zero (2023b)). The estimate in column (1) adjust for post-transfer spillovers. We calculate $\frac{d\mathbb{E}[x^e_{i\tau}]}{d\text{posttransferE}_\tau}$ in a similar way and subtract this value from the estimated energy-support transfer MPCE, accounting for differences in the length of the two periods. We estimate the MPCE out of government transfers by calculating $\frac{d\mathbb{E}[x^e_{i\tau}]}{d\text{transferG}_{i\tau}}$, evaluated at mean prices and covariates across the time period the transfers were made, dividing this by the average real value of the cost of living payments received, and multiplying the estimate by two (to account for the fact $\text{transferG}_{i\tau}$ is an average of spending effects across two months).

quantitatively small. Our finding that a significant flypaper effect applies to households that received energy-support transfers via vouchers, but not those who received it as cash or account credit aligns with recent experimental evidence that marginal propensities to consume vary depending on the mode of transfer (Boehm et al., 2025).

Table 3.2: *Marginal propensity to consume energy out of transfers*

	(1)	**(2)**
Transfer administered by energy companies	0.33	0.35
	(0.007)	(0.005)
Transfer administered by government, cash	0.04	0.03
	(0.002)	(0.002)
Adjust for post-transfer spillover	Yes	No

Notes: Table shows estimates of marginal propensities to consume energy out of energy-support transfers (first row) and cost of living payments (second row) for households that prepay their energy. Standard errors, clustered at the household-level, are shown in parentheses. Regressions include weather controls and household fixed effects. The dependent variable is deseasonalised log energy spending.

4 Model of Energy Demand

In this section we present an empirical model of household energy demand. Our aim in doing this is to quantify the welfare effects of a large energy price shock, and evaluate policy responses, including both those used during the European Energy Crisis and alternatives. We are interested in distributional (in addition to average) welfare consequences. Therefore, in our empirical specification, we use a flexible functional form for energy demands, designed to capture heterogeneity across households in responses.

4.1 Household Choice Model

Consider a household's decision over consumption of residential energy, e, and all other non-durables (excluding residential energy), n. Let $U(e, n; \theta)$ denote utility from choice (e, n), where utility is increasing in e and n, and θ captures any household-specific conditioning variables and parameters. Denote the marginal price (inclusive of any subsidy) for energy by p^e and the price of other non-durables by p^n. Let $\tilde{x}$ denote the household's total budget in the absence of any energy-support transfers, f denote a fixed access fee (i.e., the standing charge) for energy and $t \geq 0$ any energy-related transfer offered by the government. The household's problem is to allocate their budget, inclusive of any transfer and net of the fixed fee, $x \equiv \tilde{x} - f + t$, between the two goods. An optimising household solves the problem:

$$V(p^e, p^n, x; \theta) = \max_{e,n} U(e, n; \theta) \text{ s.t. } p^e e + p^n n \leq x. \tag{4.1}$$

We denote the solution to this problem, the optimal energy and other non-durable demands, by $e = \mathbb{e}^0(p^e, p^n, x; \theta)$ and $n = \mathbb{n}^0(p^e, p^n, x; \theta)$. The superscript $\mathbb{d} = 0$ denotes that these are privately optimal choices.

Recovering utility at suboptimal choices

We allow for the possibility that the flypaper effect associated with energy-support transfers received by prepay households reflects choices that are privately suboptimal. We denote the energy and other non-durable choice of a household that exhibits a flypaper effect by $(\mathbb{e}^1(p^e, p^n, x; \theta), \mathbb{n}^1(p^e, p^n, x; \theta))$, where the superscript $\mathbb{d} = 1$ denotes that the choices are suboptimal.

To recover the utility level associated with suboptimal choices we make the following assumption:

Assumption 1.

(a) *The flypaper effect impacts utility only through its influence on consumption bundles. Therefore the utility attained at suboptimal choices is given by:*

$$U(\mathbb{e}^1(p^e, p^n, x; \theta), \mathbb{n}^1(p^e, p^n, x; \theta); \theta)$$

(b) *In the absence of a flypaper effect – i.e., for non-prepay households and for prepay households in the absence of energy-support transfers – households choose the privately optimal consumption bundle (i.e., they solve problem (4.1)).*

This assumption enables us to sidestep fully specifying the underlying positive model that generates the departure from optimisation. The key restriction it entails is that the only impact of the flypaper effect on utility is through choices. This means that if there is a cognitive cost of making choices other than $(\mathbb{e}^1(.), \mathbb{n}^1(.))$, our efficiency cost calculations will be net of it.

Our approach is an application of Bernheim and Rangel's (2009) choice-based approach to welfare analysis; part (b) of Assumption 1 delineates a situation under which choices are optimal and part (a) entails a restriction on circumstances in which they are not, which allows us to recover welfare. Assumption 1 is similar to that used in Chetty et al. (2009) in their study of tax salience. However, whereas these authors rely on a first-order approximation to recover welfare, we instead use a fully specified choice model, which in our context, as we show below, is important for accurately recovering welfare effects.

In doing so, a challenge we face is that our empirical model entails specifying a flexible form for the expenditure function (i.e., the inverse of the indirect utility function $V(p^e, p^n, x; \theta)$)

and there is no analytical form for the corresponding direct utility function. It is therefore convenient to use the indirect utility function to measure a household's attained utility level. However, in order to do this while allowing for sub-optimal choices, we require a generalisation of the indirect utility function, which we outline in the following proposition:

Proposition 1. *Let* $\mathbb{e}^0$ *denote a household's optimal energy choice given by the solution to problem (4.1). Let* $\mathbb{e}^1 > \mathbb{e}^0$ *denote some suboptimally high energy level choice. The utility the household attains at these choices is*

$$\mathbb{V}(p^e, p^n, x, \mathbb{d}; \theta) = \begin{cases} V(p^e, p^n, x; \theta) & \text{if } \mathbb{d} = 0 \\ V(p^e(1-\phi), p^n, x - \phi p^e \mathbb{e}^1; \theta) & \quad\ \mathbb{d} = 1 \end{cases}$$

where ϕ *is such that:*

$$\mathbb{e}^1 = \mathbb{e}^0(p^e(1-\phi), p^n, x - \phi p^e \mathbb{e}^1; \theta).$$

ϕ is the compensated percent price reduction for energy that rationalises $\mathbb{e}^1$ as the optimal choice.[29]

We illustrate the intuition behind Proposition 1 in Figure 4.1. Consider a household with budget set 0AB. $(\mathbb{e}^0, \mathbb{m}^0)$ and $(\mathbb{e}^1, \mathbb{m}^1)$ denote the optimal and a suboptimal choice, respectively. The optimal choice, which is at a point of tangency between the household's indifference map and the budget constraint, yields utility level $\bar{U}$. The suboptimal choice yields utility $\bar{U}' < \bar{U}$. Consider the rotation of the budget constraint through the point $(\mathbb{e}^1, \mathbb{m}^1)$, such that the hypothetical budget constraint is tangent to the indifference map (line CD). At this budget constraint, we can use the true indirect utility function to compute the utility attained at the suboptimal choice.

While we apply Proposition 1 to the case where a transfer induces a flypaper effect, it can be applied to any phenomenon that causes a household to make a privately suboptimal choice.

4.2 Empirical Specification

We use data for our sample of households (indexed i) and year-months (indexed τ) on their energy budget shares ($w_{i\tau} \equiv \frac{e_{i\tau} p^e_{r(i)\tau}}{x_{i\tau}}$) and total net-of-fixed-fee budgets ($x_{i\tau}$), prices ($p^e_{r(i)\tau}, p^n_\tau$), and conditioning variables ($\mathbf{z}_{i\tau}$), to estimate household energy demand.[30] We

[29] As long as the underlying preference ordering is continuous, strictly monotone and convex, a unique $\phi \in [0, 1]$ is guaranteed to exist. These regularity conditions on preferences ensure, at a suboptimally high choice on the budget constraint, the indifference map crosses the budget constraint from above. Proposition 1 can straightforwardly be adjusted to account for a suboptimally low level of energy choice, by instead constructing the compensated price reduction for other non-durables.

[30] See Appendix A for the definition of non-durables. We calculate a price index for (non-energy) non-durables, p^n_τ, using the Consumer Price Index (CPI) microdata, and the methodology used for constructing the official UK CPI.

Figure 4.1: *Hypothetical budget set that rationalise suboptimal choice*

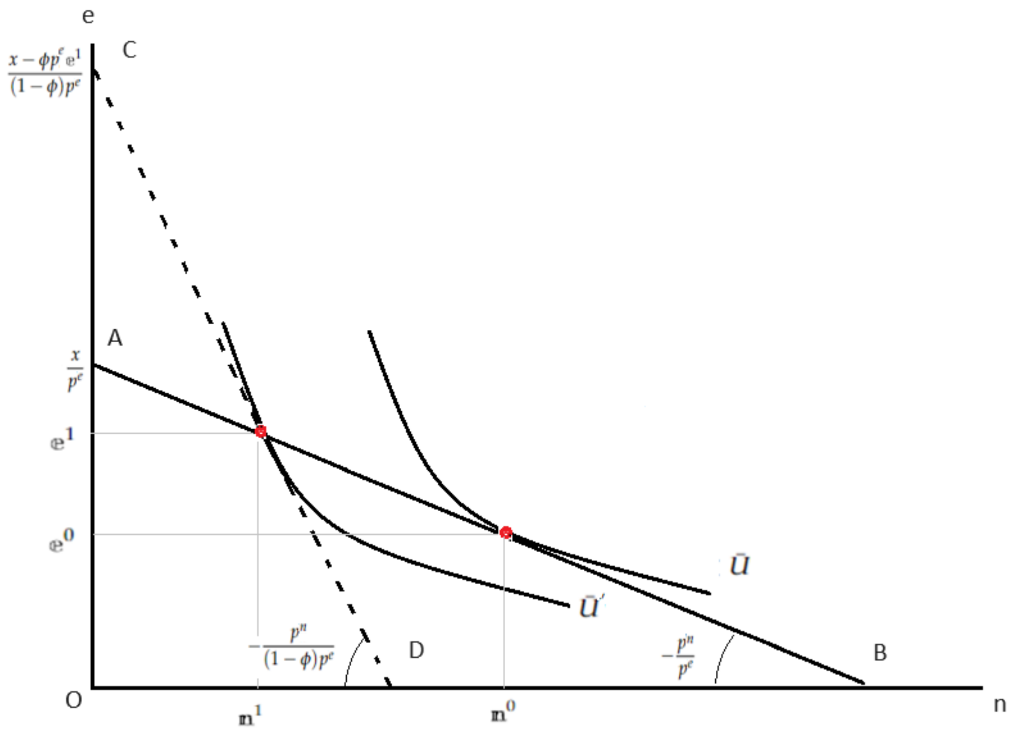

Notes: $(\mathbb{e}^1, \mathbb{m}^1)$ *and* $(\mathbb{e}^0, \mathbb{m}^0)$ *represent suboptimal and optimal choices, and* $\bar{U}'$ *and* $\bar{U}$ *the corresponding utility levels, when the consumer faces the budget set OAB. The budget set OCD represents the rotation through the point* $(\mathbb{e}^1, \mathbb{m}^1)$ *that rationalises* $(\mathbb{e}^1, \mathbb{m}^1)$ *as an optimal choice (i.e., as the solution to problem (4.1)). This entails changing the energy price from* p^e *to* $(1-\phi)p^e$ *and the budget from* x *to* $x - \phi p^e \mathbb{e}^1$*.*

also include in the model an indicator variable ($\mathbb{d}_{i\tau}$) equal to one for prepay households during the period when households received energy-support transfers. This captures any flypaper effect (i.e., responses over and above an income effect) associated with the energy-support transfer.[31]

We specify a flexible parametric form based on the Exact Affine Stone Index (EASI) demand system developed in Lewbel and Pendakur (2009). We model preference heterogeneity by incorporating a rich set of conditioning variables – including measures of past of energy usage – into the model. This approach provides a tractable way of modelling heterogenous demands that satisfy the behavioural restrictions implied by consumer theory. The model entails specifying, from a particular parametric class, a form for the expenditure function that gives rise to *implicit* uncompensated budget share demands. We specify these:

[31] The means we model the energy demand $\mathbb{e}^{\mathbb{d}}(p^e, p^n, x; \theta)$ for $\mathbb{d} = \{0, 1\}$. Under Assumption 1 we could omit "sub-optimal" observations (i.e., those for pre-pay households during the transfer period) from estimation. Doing so would mean we recover $\mathbb{e}^0(p^e, p^n, x; \theta)$, which is sufficient for assessing welfare effects (including for sub-optimal observations) under observed policy. However, as one of our counterfactual policy experiments entails varying the magnitude of a transfer that induces a flypaper effect, we prefer to explicitly model how the flypaper effect impacts demand. In addition, doing so enables us to estimate which groups – among the prepay households – exhibit the strongest flypaper effect.

$$
\begin{aligned}
\omega_{i\tau} =& (A + \sum_{l \in \mathcal{Z}_1} A_l z_{i\tau l}) + (B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l}) \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right) + (C_1 + \sum_{l \in \mathcal{Z}_2} C_{1l}) y_{i\tau} + \\
& \sum_{k>1} C_k y^k_{i\tau} + D\left(\log p^e_{r(i)\tau} - \log p^n_\tau\right) \times y_{i\tau} + (\delta + \sum_{l \in \mathcal{Z}_3} \delta_l z_{i\tau l}) \mathbb{d}_{i\tau} \qquad (4.2) \\
y_{i\tau} =& \frac{\log x_{i\tau} - (\omega_{i\tau} \log p^e_{r(i)\tau} + (1 - \omega_{i\tau}) \log p^n_\tau) + \frac{1}{2}(B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l}) \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right)^2}{1 - \frac{1}{2} D \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right)^2}, \qquad (4.3)
\end{aligned}
$$

where

$$
\Psi \equiv \left(A, \{A_l\}_{l \in \mathcal{Z}_1}, B, \{B_l\}_{l \in \mathcal{Z}_2}, C_1, \{C_{1l}\}_{l \in \mathcal{Z}_2}, \{C_k\}_{k>1}, D, \delta, \{\delta_l\}_{l \in \mathcal{Z}_2}\right)
$$

are model parameters (see Appendix C.1 for full details).

Equations (4.2) and (4.3) implicitly define the energy demand function, which we denote $\omega_{i\tau} \equiv \omega(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, \mathbb{d}_{i\tau}, \mathbf{z}_{i\tau}; \Psi)$.[32] When $\mathbb{d} = 0$, by construction, $y_{i\tau} = \log V_{i\tau}$, where $V_{i\tau}$ is the realised utility level, and $\omega(.)$ is the budget share form of the optimal demand function (both defined in problem (4.1)). When $\mathbb{d} = 1$, $\omega(.)$ deviates from the optimal choice. Budget share demand for other non-durables can be obtained from $\omega^n_{i\tau} = 1 - \omega_{i\tau}$.

By including the log difference in prices in equation (4.2), we ensure preferences satisfy adding-up, homogeneity, and Slutsky symmetry. We do not impose the two additional *inequality* constraints implied by consumer theory in estimation, which require that the expenditure function is concave in prices and monotone in decision utility, instead checking they are satisfied post-estimation.

The A parameters determine the budget share intercept, the B parameters control the demand response to a (compensated) price change, the C parameters govern the shape of the Engel curve, D allows for an interaction effect between price responses and the Engel curve, and the δ parameters capture the behavioural (i.e., flypaper) effect of the energy-support transfer.[33] We allow each to vary with a set of conditioning variables collected in the sets $\mathcal{Z}_k$ for $k = \{1, 2, 3\}$, where $\mathcal{Z}_3 \subset \mathcal{Z}_2 \subset \mathcal{Z}_1$. $\mathcal{Z}_3$, which shifts the flypaper effect, contains a set of dummy variables capturing the decile of the pre-shock energy spending distribution to which the household belongs; $\mathcal{Z}_2$, which shifts price responsiveness and the first-order Engel curve coefficient, adds to this an indicator for whether the household prepays for energy and $\mathcal{Z}_1$, which shifts the budget share constant, adds a set of dummy variables capturing the decile of the pre-shock energy spending share distribution to which

[32] This function does not have an analytical form; we solve for it numerically. A sufficient condition for the equations (4.2) and (4.3) to uniquely define $\omega_{i\tau}$ and $y_{i\tau}$ is that the expenditure function is monotone in utility. See Appendix C.1. We check that this condition is satisfied post-estimation.

[33] We set $\mathbb{d}_{i\tau} = 1$ for prepay households over October 2022-April 2023, this is the 6 month period when the energy-support transfers were given and the first month afterwards. This allows for any smoothing of the transfer into the first post month. We also include indicators for during and after the transfer period for the direct debit households to capture any smoothing behaviour.

the household belongs, as well as household region dummies, seasonal month dummies and detailed weather controls. This rich set of controls helps ensure the model captures variation in energy demands across households.

Estimation

We estimate the model using GMM (Hansen, 1982).[34] Let $\mathbf{h}_{i\tau}$ denote a vector of instruments and $\epsilon_{i\tau} \equiv w_{i\tau} - \omega(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, \mathbb{d}_{i\tau}, \mathbf{z}_{i\tau}; \Psi)$ denote the prediction error in the model's demands. We specify the population moment condition $\mathbb{E}[\epsilon_{i\tau}\mathbf{h}_{i\tau}] = 0$ and obtain parameter estimates $\widehat{\Psi}$ by selecting the value of Ψ that minimises the (quadratic form) of the sample analogue of moment conditions. We provide full details of estimation and our instrument set in Appendix C.2. We include in $\mathbf{h}_{i\tau}$ the conditioning (i.e. z) and relative price variables that enter the model. Given that residential energy prices are adjusted at pre-specified times by a regulator that bases the regulated price on international wholesale prices, it is natural to treat price as exogenous in estimation. However, we exclude log expenditure, $\log x_{i\tau}$, from $\mathbf{h}_{i\tau}$, instead including functions of the household's monthly income. This allows for the possibility that a shock to a household's energy demand is correlated with their total spending.

4.3 Model Estimates and Fit

We estimate the model using data covering June 2021 – June 2023, using the period July – December 2023 to assess the out-of-sample fit of the model. We use the sample of prepay and variable direct debit households, but in all policy simulations weight observations across the two payment types to correspond to the UK population.

In Table 4.1 we report the model parameter estimates. The first column reports estimates of the baseline parameters; we model the energy demand Engel curve as a quadratic – including higher-order terms does not meaningfully change its shape. The remaining columns show interaction effects for the budget share constant, price and first-order Engel curve term, with indicator variables for whether the household prepays and the decile of the pre-shock energy spending distribution to which they belong; we omit these for the second-order Engel curve and price-Engel curve terms as adding these parameters to the model makes little difference to the results we present in the following section. Consistent with the evidence in Section 3.3, we model a flypaper effect for the energy-support transfer administered by energy companies for those households that prepay, allowing the effect to

[34] An alternative estimator is iterated least squares (Blundell and Robin, 1999). We use this to obtain starting values for the GMM estimator. In practice both methods lead to similar parameter estimates. GMM has the advantage that it yields standard errors that are valid without any adjustment.

vary across pre-shock energy spending deciles. At the estimated parameters, over 99% of household-month observations satisfy the concavity restriction on the expenditure function and all observations satisfy the monotonicity restriction.

Table 4.1: *Parameter estimates*

		×prepay	×pre-shock energy spending decile								
			2	3	4	5	6	7	8	9	10
Constant											
(A)	1.0081	-0.0910	0.0686	0.1036	0.1334	0.1545	0.1942	0.2124	0.2449	0.2864	0.2977
	(0.0509)	(0.0113)	(0.0122)	(0.0121)	(0.0124)	(0.0129)	(0.0134)	(0.0137)	(0.0148)	(0.0167)	(0.0210)
Price											
(B)	0.2241	-0.0282	-0.0033	-0.0029	-0.0044	-0.0032	-0.0016	-0.0037	-0.0016	-0.0018	-0.0060
	(0.0085)	(0.0012)	(0.0016)	(0.0016)	(0.0016)	(0.0016)	(0.0017)	(0.0017)	(0.0017)	(0.0019)	(0.0024)
Implicit utility											
(C_1)	-0.2407	0.0115	-0.0091	-0.0133	-0.0168	-0.0192	-0.0242	-0.0261	-0.0298	-0.0344	-0.0342
	(0.0146)	(0.0016)	(0.0018)	(0.0018)	(0.0018)	(0.0018)	(0.0019)	(0.0020)	(0.0021)	(0.0023)	(0.0029)
(C_2)	0.0151	-	-	-	-	-	-	-	-	-	-
	(0.0010)										
Price× Implicit utility											
(D)	-0.0236	-	-	-	-	-	-	-	-	-	-
	(0.0012)										
Flypaper effect											
(δ)	-	0.0178	-0.0013	-0.0013	-0.0008	-0.0018	-0.0016	-0.0014	-0.0020	-0.0031	-0.0057
		(0.0009)	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0012)	(0.0011)	(0.0012)	(0.0013)	(0.0016)

Notes: Table shows (a subset) of the parameter estimates for the energy demand model given by (equations (4.2) and (4.3)). Standard errors, clustered at the household level, are shown in parenthesis. Model also includes the constant shifters: dummy variables for pre-shock energy spending share decile, region dummies, calendar month dummies, 5^{th}-order polynomials in local monthly minimum and maximum temperature, the squared difference between maximum and minimum temperature, local monthly rainfall and an indicator for Dec 2021 and Jan 2022, when a work-from-home order was in place. We estimate the model using GMM and data covering June 2021-June 2023.

In Figure 4.2 we present two graphs that validate the fit of our model. Panel (a) compares the elasticities we estimate using differences across the cap changes in Section 3.2 with corresponding elasticities based on the model. We plot the 25 elasticities that summarise the heterogeneity across pre-shock energy spending and pre-shock income quintiles. The figure shows that the two set of elasticity estimates are similar. In Appendix C.3 we plot the 25 model-based elasticities (the equivalent of Figure 3.3). Both sets exhibit the same pattern; conditional on income, elasticities are increasing in magnitude across pre-shock energy spending quintiles and conditional on spending quintile, they are declining in magnitude across income quintiles. The model allows us to decompose responses to price changes into substitution and income effects. If, in response to an energy price increase, households reduce their energy consumption primarily due to a substitution effect, the associated welfare loss is smaller than if their response is principally driven by an income effect. We find that, on average, the income effect associated with the April 2022 cap rise accounts for approximately 10% of the Marshallian response.[35]

Figure 4.2: *Model fit*

(a) In-sample; elasticities

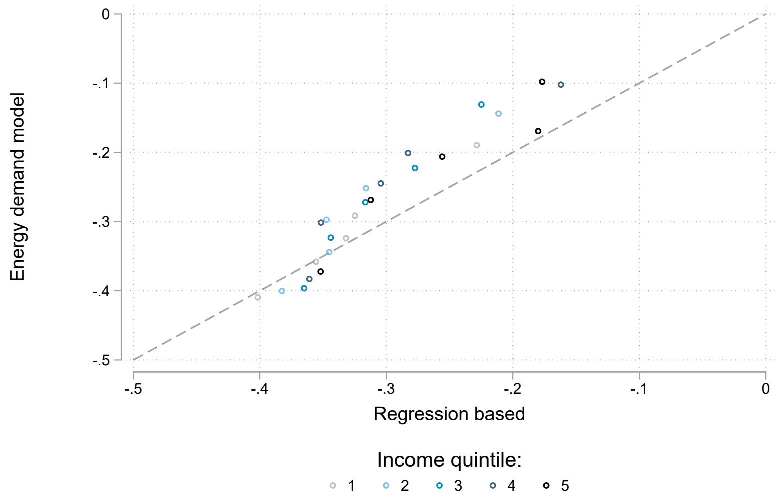

(b) Out-of-sample

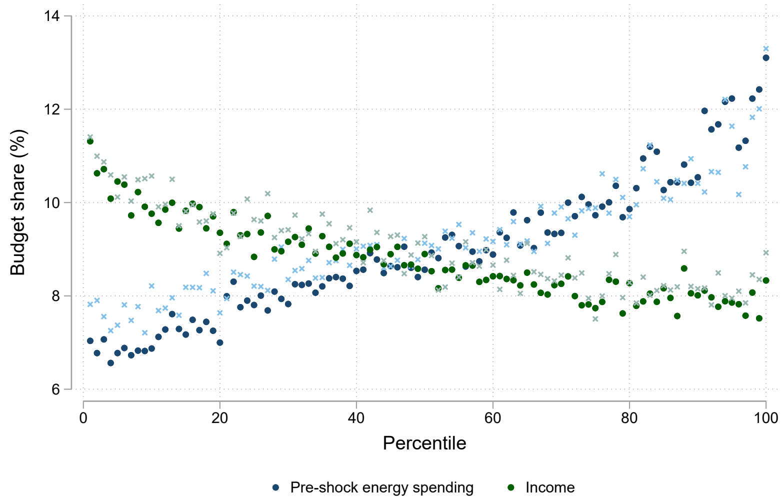

Notes: Panel (a) scatters the pre-shock energy spending-income quintile specific Marshallian price elasticities estimated using the differences design of Section 3.2, against the within-group average based on the simulated model-based Marshallian price elasticities, for the period Apr.-Sept. 2022. Panel (b) plots the average observed (crosses) and predicted (circles) budget shares across percentiles of the pre-shock energy spending (green) and the income (blue) distributions for the out of estimation sample period Oct.-Dec. 2023.

[35]In Appendix C.3 we illustrate heterogeneity in income effects, by plotting separate Engel curves for households belonging to each decile of the pre-shock energy spending distribution. At low levels of total expenditure the difference in energy budget shares between the top and bottom decile is over 15 percentage points. This gap shrinks to less than 2 percentage points at the top of the total expenditure distribution. These patterns are similar to those documented in Lewbel and Pendakur (2017), who, using cross-sectional data, model preference heterogeneity through random Barten scales.

In panel (b) of Figure 4.2 we summarise the out-of-sample fit of the model. We do this by comparing observed energy budget shares for the period October – December 2023 with their model-based predicted values. Over this three month period the energy price was 74% of its June 2023 level (the final month of our estimation sample). We plot how observed (crossed) and predicted (dots) demands vary across percentiles of the pre-shock energy spending distribution (in blue) and the income distribution (in green). The figure shows that our model successfully recovers how the demands vary across both of these dimensions.[36]

5 Welfare Effects of an Energy Crisis and Policy Responses

In this section we first outline our money-metric measure of household welfare. We then characterise the distribution of household-level welfare losses resulting from the European Energy Crisis, and how this was affected by the implemented UK policy response. Finally, we develop a social welfare framework to evaluate the effectiveness of alternative policy responses to a large energy price shock.

5.1 Money-Metric Household Welfare

Consider household i in period τ that faces the price vector $(p^e_{r(i)\tau}, p^n_\tau)$, has total budget $x_{i\tau}$, exhibits flypaper effect $\mathbb{d}_{i\tau}$, and is characterised by conditioning variables and parameters $\theta_i = (\mathbf{z}_i, \Psi)$. During the energy crisis, their attained level of utility (given by $\mathbb{V}(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, \mathbb{d}_{i\tau}; \theta_i)$ and defined in Proposition 1) is influenced by policy parameters $\mathbb{P} = (s, t, L)$. s is the subsidy rate for energy (the consumer energy price is related to the pre-subsidy price, $P^e_{r(i)\tau}$, by $p^e_{r(i)\tau} = (1-s)P^e_{r(i)\tau}$). t is the level of any transfer (with $x = \tilde{x} - f + t$, where $\tilde{x}$ is pre-transfer gross budget and f the energy fixed fee). L is an indicator variable that denotes whether the transfer induces a flypaper effect among prepay households; we refer to a transfer that does as "labelled". Note, we use this term to capture the fact that the transfers implemented by the UK government were designated for "energy support" *and* their mode of delivery was via energy suppliers.

With a slight abuse of notation we re-write attained utility directly as a function of the policy parameters $\mathbb{V}_{i\tau}(\mathbb{P})$. If the household's choice is optimal (meaning $\mathbb{d} = 0$) we recover attained utility by jointly solving equations (4.2) and (4.3). If the household's choice is influenced by a flypaper effect due to a labelled transfer, we first solve for the budget constraint

[36]The EASI demand model allows the researcher to treat the prediction errors from the GMM problem as structural preference shocks or measurement error. In the former case, in-sample predicted demands will equal observed demands by construction. Figure 4.2 shows that if we treat the errors as measurement error our model nonetheless captures how demands vary with energy need and income, even out-of-sample. In the simulations below we adopt this interpretation of the prediction errors.

pivot that supports the choice as optimal (see Proposition 1) and then jointly solve equations (4.2) and (4.3) at this budget set (and setting $\mathbb{d} = 0$) to obtain the attained level of utility (see Appendix C.4 for details). We use a cardinalisation of $\mathbb{V}_{i\tau}(\mathbb{P})$ that has the money-metric interpretation of £s at pre-crisis prices $(p^e_{r(i)0}, p^n_0)$, which we denote by $\mathbb{V}^{MM}_{i\tau}(\mathbb{P})$.[37] The money-metric – equivalent variation – loss household i suffers in month τ due to the energy price shock is then: $\mathcal{L}_{i\tau}(\mathbb{P}) \equiv (\tilde{x}_{i\tau} - f^e_{r(i)0}) - \mathbb{V}^{MM}_{i\tau}(\mathbb{P})$, where $(\tilde{x}_{i\tau} - f^e_{r(i)0})$ is their budget net of the pre-shock fixed fee. The household's loss over the energy price shock period, $\tau \in \{\underline{\tau}, \overline{\tau}\}$, is $\mathcal{L}_i(\mathbb{P}) = \sum_{\tau=\underline{\tau}}^{\overline{\tau}} \mathcal{L}_{i\tau}(\mathbb{P})$. $\mathcal{L}_i(\mathbb{P})$ can be interpreted as the household's willingness-to-pay (at pre-shock prices) to avoid the price shock under policy response $\mathbb{P}$. We denote *proportional losses* by $l^y_i(\mathbb{P}) \equiv \frac{\mathcal{L}_i(\mathbb{P})}{Y_i}$ where Y_i is household income.[38]

5.2 The European Energy Crisis and Observed Policy Response

In Figure 5.1 we summarise the distribution of monetary (panel (a)) and proportional (panel (b)) losses that households incurred during the height of the European Energy Crisis from October 2022 to March 2023. The markers show average losses and the dark and light shading show the interquartile and interdecile ranges. Panel (c) reports average and aggregate losses.

No government intervention

In the absence of government intervention, households would have been, on average, £206 per month (6.1% of income) worse off, implying an aggregate equivalent variation loss over the six month period of £35.2bn. There would also have been significant inequality in losses, with the 95th percentile of proportional losses equal to 11% of income. While proportional losses decline with income, there is substantial variation *conditional on income*. For example, within the bottom income decile, the 10th and 90th loss percentiles would have been 6 and 17%. The shock would have led to considerable hardship for many households, with the number classified as being in "energy poverty" (spending more than 10% of their after-housing-costs income on energy, see Appendix B.5) rising to over 9.5 million (compared with 4 million in 2019).

[37] Let $V(p^e_0, p^n_0, x; \theta_i) \equiv V^{\text{Pre}}_i(x)$ denote their indirect utility function evaluated at pre-shock prices and define $\chi^{\text{Pre}}_i(u)$ to be its inverse (i.e., expenditure) function. Then $\mathbb{V}^{MM}_{i\tau}(\mathbb{P}) \equiv \chi^{\text{Pre}}_i(\mathbb{V}_{i\tau}(\mathbb{P}))$.

[38] In practice we measure Y_i as the household's average monthly income over the preceding tax year. We scale it by the number of months in the shock period (i.e., six), so l^y_i is interpretable as loss as a share of income.

Figure 5.1: *Distribution of household losses by income*

(a) In £s per month (b) As % of income

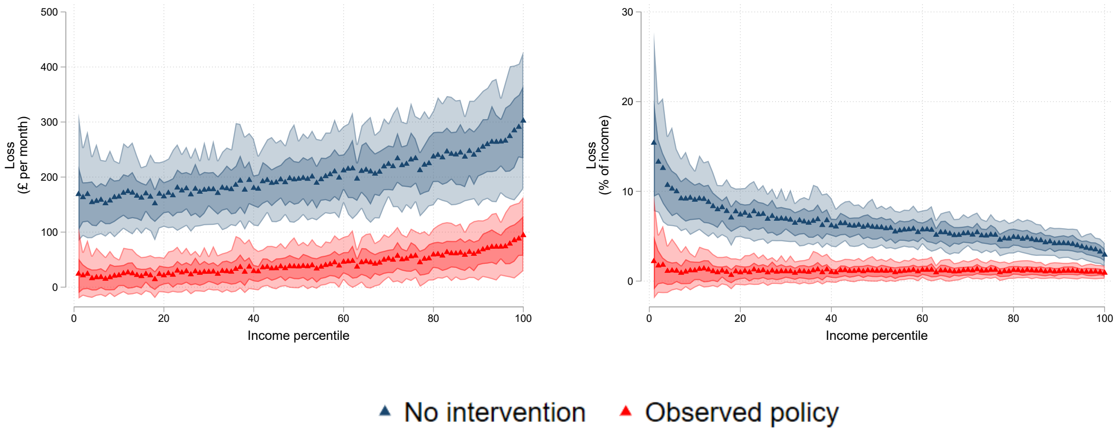

(c) Average losses

	No intervention	Observed policy	Under no behavioural response: No intervention	Under no behavioural response: Observed policy
Monetary (£pm)	206.48	42.93	326.87	80.03
	[204.33, 208.18]	[41.57, 44.32]	[319.50, 334.54]	[76.83, 83.36]
Proportional to Y (%)	6.09%	1.14%	9.67%	2.25%
	[6.03, 6.13]	[1.10, 1.18]	[9.49, 9.87]	[2.17, 2.34]
Aggregate (£bn)	35.18	7.32	55.70	13.64
	[34.82, 35.47]	[7.08, 7.55]	[54.44, 57.01]	[13.09, 14.20]
Number of households in energy poverty	9.6m	7.3m	21.6m	11.6m

Notes: For each household-month in Oct. 2022-Mar. 2023 we compute the loss from the price shock based on equivalent variation at pre-shock prices, and compute the average loss for each household. The figures summarise how the distribution of household losses varies by the percentile of the Apr. 2021-Mar. 2022 income distribution they belong to. The markers show means, the dark area shows 25th and 75th percentiles and the light area shows the 10th and 90th percentiles. "No intervention" refers to the case of no government policy intervention and "Observed policy" refers to the combination of subsidy and transfer used in practice. Panel (a) reports losses in £ per month, panel (b) reports average monthly losses scaled by average monthly income over Apr. 2021-Mar. 2022. Panel (c) reports average monthly and aggregate losses. "Under no behavioural response" refers to simulating welfare costs assuming household demands remain at their pre-shock level. Energy poverty is defined as spending more than 10% of after-housing-costs income on energy – see Appendix B.5 for further details. 95% confidence bands are reported in square brackets.

Observed policy response

Figure 5.1 shows that the UK policy response to the energy crisis significantly reduced households' losses. It lowered average losses to £43 per month, or 1.1% of income. The intervention was particularly effective at supporting the most exposed households; for example, losses for the household at the 95th percentile of the proportional loss distribution were 3%, compared with 11% of income under no intervention. In addition, the policy flattened the relationship between mean proportional losses and income. It prevented 2.3

million households from falling into energy poverty (over a third of the increase that would have resulted in the absence of government intervention).

The intervention reduced aggregate equivalent variation losses by 21%. However, the total public resources allocated to the relief packages amounted to 90% of the aggregate money-metric loss in the absence of government intervention. This gap between resource costs and the achieved reduction in losses reflects policy-induced inefficiencies. The efficiency costs totalled £3.7 bn (or 12% of total public costs of the relief package). We decompose these costs in Table 5.1.[39]

Efficiency costs of the subsidy. 69.8% of the total efficiency cost arises directly from the deadweight loss from subsidising energy – i.e., the fact that supporting households through a subsidy raises their utility by less than giving them equivalent funds as a lump-sum transfer. This effect is linked to households' willingness to substitute away from energy when its price rises. The larger these substitution effects, the higher the efficiency costs of an energy subsidy.

Efficiency costs of the transfer delivery. 12.0% of the total efficiency cost is driven by the flypaper effect for prepay households induced by the energy-support transfers. This creates two sources of inefficiency. First, there is a direct utility cost to those households from making suboptimal choices, accounting for 2.0% of total efficiency costs. Second, a bigger loss – nearly five times as large – arises from a fiscal spillover, as the transfers stimulate additional consumption of a subsidised good.[40]

Costs of extra carbon emissions. The final source of efficiency costs, accounting for 18.2% of the total, arises from the subsidy and flypaper effect acting to raise carbon emissions. To measure the social cost of these additional emissions, we calculate each household's policy-induced increase in energy consumption and multiply it by the factor α, representing the

[39] We measure efficiency costs as the reduction in aggregate money-metric utility under the observed policy, compared to a scenario where each household receives equivalent public funds via a direct transfer (see Appendix C.5). Since we base money-metric utility at *pre-shock* prices, we compute efficiency costs as a share of total public resources used, by converting these resources to their value at pre-shock prices. Alternative natural denominations for money-metric utility – and thus efficiency costs – are *no intervention post-shock* prices or *post-shock* prices. The former leads to an alternative equivalent variation measure, while the latter results in a compensating variation measure, which has the important drawback that the underlying money-metric cardinalisation of utility is policy-dependent (see Auerbach, 1985). The value we compute for total public resources as a share of non-intervention losses, as well as the extent to which these losses are reduced, remains the same under each of these three alternatives.

[40] The strength of the flypaper effect is consistent with the descriptive evidence we present in Table 3.2. The MPCE associated with the £400 energy-support transfers is 0.3 more than the MPCE from a cash transfer for the 15% of UK households (4.2m) that prepay. With a subsidy rate of 39%, this implies a fiscal spillover of $(0.39/(1\text{-}0.39))\times 0.3\times 400\times 4.2\text{m}=£322\text{m}$, which is slightly lower than the £370m reported in Table 5.1.

monetary cost of the atmospheric externality created by the carbon emissions from a unit of energy consumption.[41]

Table 5.1: *Efficiency costs*

	Total efficiency cost	Source of efficiency cost			
			Labelling:		
		Price signal	choice distortion	fiscal spillover	Carbon externality
Aggregate (£bn)	3.73	2.61	0.07	0.37	0.68
	[3.65, 3.84]	[2.53, 2.70]	[0.07, 0.08]	[0.35, 0.39]	[0.66, 0.70]
Contribution:		(69.83%)	(2.00%)	(9.96%)	(18.20%)
		[69.12, 70.86]	[1.71, 2.22]	[9.21, 10.54]	[18.05, 18.36]

Notes: We report the efficiency cost associated with the implemented UK policy of a 39% energy subsidy and £66 per month transfer, over Oct 2022.-Mar. 2023, and decompose it into 4 mutually exclusive and exhaustive sources. Row (1) reports aggregate numbers and row (2) reports the percentage contribution from each source. See Appendix C.5 for decomposition details. 95% confidence bands are reported in square brackets.

Importance of modelling behavioural response

In the final two columns of Figure 5.1(c), we present aggregate losses from the energy crisis, both without any policy intervention and under the UK's observed policy response, assuming *no behavioural responses from households*. The results show that ignoring households' behavioural responses significantly overestimates aggregate equivalent variation losses – by 58% in the absence of intervention and by 86% under the implemented policy response. In the case of a large shock and policy intervention, household substitution responses substantially mitigate welfare losses.

The quantitative significance of these responses also means that the total *fiscal externality* of the policy response – the difference between its cost under no behavioural response and its actual cost, which amounts to 21% of total policy costs – overstates the extent of policy-induced inefficiencies (equal to 12% of total policy costs). This occurs because household behavioural responses meaningfully impact welfare and reduce the true distortionary cost of intervention below the policy's fiscal externality.

Although it common to use first-order approximations of welfare costs based on pre-shock spending, and to quantify distortionary effects of policy while ignoring the impact of

[41]We then convert this cost into its value at pre-shock prices to maintain consistency with our money-metric utility cardinalisation (see Appendix C.5 for details). To measure α we use the average carbon emission generated per £ spent on residential energy in the first quarter of 2023, and use a value of £59 per tonne for the social cost of carbon (Department for Energy Security and Net Zero, 2023c). See Appendix B.4 for further details.

behavioural changes on agent's private welfare, our results highlight the significant biases this approach can produce in the context of large shocks.

5.3 Policy Design

In this section, we embed our model of household behaviour within a social welfare framework designed to quantify the policy trade-off between targeting assistance to households most affected by a price shock and minimising inefficiencies. We consider an energy price shock equivalent to that experienced during the European Energy Crisis.

Comparable policy packages

Alternative policies lead to different distributions of household-level welfare effects, different aggregate energy consumption, and hence carbon emissions, and may potentially use different amounts of public funds. When comparing alternative policies it is therefore useful to restrict attention to a subset of alternatives that are comparable in some sense. We focus on alternative policies that lead to the same level of public expenditure *inclusive of the public costs of carbon emissions*. Denote the policy response implemented by the UK during the European Energy Crisis $\mathbb{P}^O = (s^O, t^O, L^O)$, the energy consumption of household i over the energy crisis period by $e_i(\mathbb{P}^O) = \sum_{\tau=\underline{\tau}}^{\overline{\tau}} e_{i\tau}(\mathbb{P}^O)$, their energy consumption in the absence of policy intervention by $e_i(\emptyset)$, and their subsidy-exclusive energy spending over the crisis by $x_i^e(\mathbb{P}^O) = \sum_{\tau=\underline{\tau}}^{\overline{\tau}} P^e_{r(i)\tau} e_{i\tau}(\mathbb{P}^O)$. The total public resources devoted to the policy response is given by:

$$\bar{R} \equiv s^O \sum_{i=1}^{N} x_i^e(\mathbb{P}^O) + N \times 6t^O + \alpha \sum_{i=1}^{N} \Big(e_i(\mathbb{P}^O) - e_i(\emptyset) \Big), \tag{5.1}$$

where N is number of UK households and, as we focus on the 6 month period between October 2022 and March 2023 when a subsidy and transfer were in place, the monthly transfer t is pre-multiplied by six. α converts the policy-induced additional energy consumption into the social cost of the resulting carbon emissions.[42]

We consider counterfactual policies that entail expending no more (carbon externality-inclusive) public resources than $\bar{R}$, focusing on four counterfactual policy menus:

1. $(s, t, L = 0)$ – a subsidy for energy consumption and universal transfer that is not labelled.

[42]More generally, α measures the extent to which the subsidy-exclusive marginal energy price departs from the social marginal cost of energy. Borenstein and Bushnell (2022) show evidence in some locations in the US that fixed cost recovery can lead the marginal electricity price to exceed the (average) social marginal cost, which would lead to a negative value for α. As the suppliers in the UK market use two-part tariffs, which in principal facilitate fixed cost recovery without this additional distortion to marginal prices, we assume the subsidy-exclusive energy price coincides with suppliers' private marginal cost.

2. $(s, (t/Y_i), L = 0)$ – a subsidy and transfer that is proportional to inverse household income (measured based on the preceding tax year).

3. $(s, (t \times E_i), L = 0)$ – a subsidy and transfer that is proportional to a household's monthly energy usage averaged over October 2021-March 2022 (the same 6 calendar months one year earlier).

4. $(s, (t \times E_i/Y_i), L = 0)$ – a subsidy and transfer proportional to previous energy usage divided by household income.

Together, these capture the main ways in which European governments responded to the energy crisis. For instance, some policies linked transfers to income by treating them as taxable income, while others tied support to historic energy use.[43]

Social welfare

To compare the impact of counterfactual policies on the distribution of household losses, we use a social loss function that aggregates household-level losses to a single summary measure. We use a loss function that implies, if the social planner could directly allocate losses across households, it would choose to equate *proportional losses*, $l_i^y(\mathbb{P})$, across all households. In practice, the planner faces a trade-off between targeting support to get as close to equi-proportional losses as is feasible and the efficiency costs induced by policy. We therefore model the planner as choosing policy to minimise the convex sum of households' proportional losses, given by:

$$\mathcal{W}(\mathbb{P}) = \frac{1}{N} \sum_{i=1}^{N} \frac{1}{\psi} \left(\exp(\psi \times l_i^y(\mathbb{P})) - 1 \right); \quad \psi > 0. \tag{5.2}$$

$\mathcal{W}$ captures concern for both vertical equity and loss inequality, conditional on income. Vertical equity is reflected by scaling money-metric losses by the reciprocal of income; a practice standard to the optimal tax literature (e.g., Saez, 2002; Allcott et al., 2019). Aversion to loss inequality within income groups is captured by the convex transformation of losses, which allows for the possibility that the planner places more than double the weight on a household experiencing twice the loss of another household with the same income. The degree of convexity, and thus the planner's aversion to deviations from equi-proportional losses, is controlled by the parameter ψ. Note, as $\psi \to 0$, the planner becomes indifferent to

[43]For example, Belgium and Germany provided households with bill rebates, which were treated as taxable income (Germany) or were in part paid back by higher-income households through a specific income tax levy (Belgium). Germany also offered support based on past energy usage through their "Price Brake", which subsidised energy use below a personalised quota based on 80% of the previous year's consumption. For households that consume less than the quota, this policy is equivalent to a subsidy, for those that consume 80% or more it is equivalent to a transfer based on historic consumption (though with a cost for those who "bunch" at 80%). See Arregui et al. (2022) for more details.

inequality in proportional losses, and minimising equation (5.2) is equivalent to maximising the sum of money-metric utilities with Pareto weights equal to inverse income. We set ψ to rationalise the observed policy (s^O, t^O) as optimal, from among the set of policies that consist of a subsidy and a labelled universal transfer, and expend the same amount of public resources as observed policy ($\bar{R}$ in equation (5.1)). We first hold ψ at this value, before discussing the robustness of our results to varying ψ.[44]

Denote by $\xi^{\mathbb{P}}$ the level of proportional losses that, were all households to have this loss level, attains the same value of equation (5.2) as under policy $\mathbb{P}$, and notice that $\xi^{\mathbb{P}}$ is related to $\mathcal{W}(\mathbb{P})$ through a monotone transformation (i.e., $\xi = \frac{1}{\psi}\log(\psi\mathcal{W}+1)$); thus it converts the units of $\mathcal{W}$ so that they have the interpretation of the constant equi-proportional loss that leads to same level of social loss as the loss distribution induced by policy $\mathbb{P}$.[45] We therefore use $\xi^{\mathbb{P}}$ to measure social losses.

We decompose the social losses under policy $\mathbb{P}$, $\xi^{\mathbb{P}}$, into three components:

$$\xi^{\mathbb{P}} = \underbrace{\xi^{LS}}_{\text{uncompensated losses}} + \underbrace{\left(\frac{\bar{\mathcal{L}}^{\mathbb{P}}}{\bar{Y}} - \xi^{LS}\right)}_{\text{efficiency costs}} + \underbrace{\left(\xi^{\mathbb{P}} - \frac{\bar{\mathcal{L}}^{\mathbb{P}}}{\bar{Y}}\right)}_{\text{targeting costs}}. \tag{5.3}$$

ξ^{LS} denotes losses under household-specific, non-labelled, lump-sum transfers (with public resource cost $\bar{R}$) that equate proportional losses across households. This captures losses that it is not possible to compensate under the resource constraint (equation (5.1)), even under an idealised policy that induces zero efficiency costs and achieves the social loss minimising distribution of household-level losses.[46] The second term is the difference between average household-level monetary losses under policy $\mathbb{P}$, scaled by average income, and uncompensated losses. This term measures the efficiency costs of policy $\mathbb{P}$. Note that these reflect the deadweight loss associated with the allocation of public funds for household support, rather than the efficiency costs of raising these funds, which remain constant across all policies we consider.[47] The final term captures social losses that arise when a policy fails to

[44]In Appendix C.5 we plot the associated welfare weights. They imply for households with losses of £43pm (the average under observed policy) the planner places 2.7 times more weight per £ of loss on a household at the 25th income percentile compared to one at the 75th percentile. They also imply, for a household with median income, the planner places 2.1 times more weight per £ of loss at the 75th monetary loss percentile compared to one at the 25th percentile.

[45]ξ is defined by $\frac{1}{N}\sum_i^N \frac{1}{\psi}(\exp(\psi\xi)-1) = \mathcal{W}$. Hence, $\xi = \frac{1}{\psi}\log(\psi\mathcal{W}+1)$. If policy $\mathbb{P}^*$ minimises $\mathcal{W}$ it also minimises ξ.

[46]"Zero efficiency cost" is only true approximately. The lump-sum transfers raise energy demands, slightly, through an income effect, compared with no intervention, inducing an increase in carbon emissions. We take this into account when computing the lump-sum transfer scheme (so they observe the resource constraint). The increase in carbon emissions is very small, accounting for only 0.36% of the total public resource cost of the transfer scheme.

[47]This separation of the incidence and efficiency costs of how funds are spent from how they are raised echos that in the marginal value of public funds framework (see Hahn et al. (2024)). The UK government funded its energy support package through a combination of increased borrowing and special levies on oil and gas producers, which raised £6bn in 2022-23 (Office for Budget Responsibility, 2023).

achieve equi-proportional losses. We refer to these as targeting costs. They are equal to the difference between the welfare-equivalent constant loss level ($\xi^{\mathbb{P}}$) and the average monetary loss level scaled by average income.

The efficiency-targeting trade-off

Figure 5.2(a) illustrates the efficiency-targeting trade-off for different policy menus. For each menu, we vary the fraction of public resources $\bar{R}$ allocated to the subsidy, with the remainder used to fund transfers. The red line represents the trade-off for the policy menu $(s, t, L = 1)$.[48] Moving leftwards along the line corresponds to a higher subsidy, which improves targeting by reducing the proportional losses of those most vulnerable to price increases but comes with increased efficiency cost. The cross denotes the observed UK policy, which, by construction, is the social-loss minimising policy (based on our calibration of ψ). The shaded gray area delineates the combinations of targeting and efficiency costs that result in lower social losses than those under the observed policy. The social value of losses under this policy is equivalent to all households experiencing a loss of 1.96% of income. Panel (b) decomposes this into the contribution from uncompensated losses (0.58 percentage points), efficiency costs (0.57 percentage points) and targeting costs (0.81 percentage points).

The dark blue line (in panel (a)) illustrates the welfare gains that arise if the transfer is unlabelled and therefore does not generate a flypaper effect. For all values of s below the maximum level (53%), the policy menu $(s, t, L = 0)$ dominates $(s, t, L = 1)$. This is because, for any policy in the latter menu, there exists one in the former with the same level of targeting costs but lower efficiency costs. By avoiding the flypaper effect, the $(s, t, L = 0)$ menu avoids distortions in the choices of prepay households and the associated fiscal spillover. The government could have achieved the same level of targeting costs as observed policy, but at 18.7% lower efficiency costs.

The lighter blue line show the effect of specifying transfers to be inversely proportional to income. This policy menu, $(s, (t/Y_i), L = 0)$, significantly improves on those based on a universal transfer. The optimal policy under this menu results in total social losses of 1.51%, which is around four-fifths of those under the menu $(s, t, L = 0)$. However, the optimal subsidy level remains high at 40%. Moreover, a policy of solely income-based transfers – represented by the rightmost point of the light blue line – leads to larger social losses

[48]The δ parameters in the demand equation (4.2) capture the magnitude of the flypaper effect when the transfer is £66 per month. As we vary the magnitude of the labelled transfer, we scale these parameters (for instance if the transfer is £33, we half them). This means we hold the excess marginal propensity to consume energy for those with a flypaper effect constant.

compared to the observed UK policy. This is because, under such a policy, many households still face large proportional losses, resulting in high targeting costs.

Figure 5.2: *Counterfactual policy responses*

(a) Efficiency-targeting trade-off

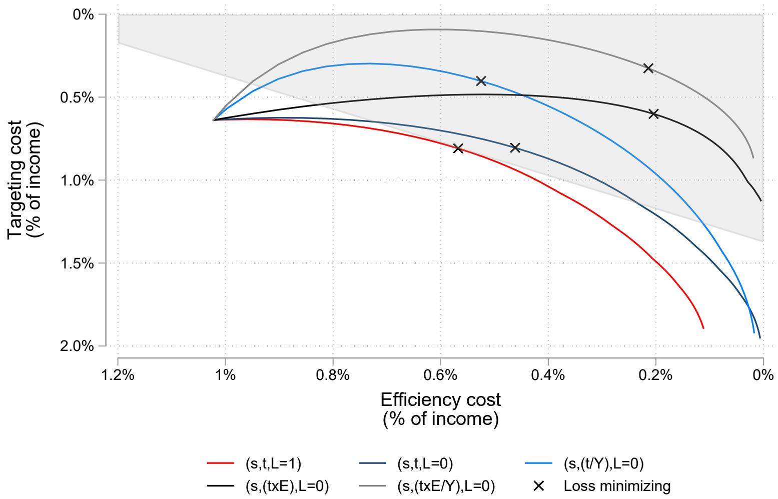

(b) Social loss minimising policy

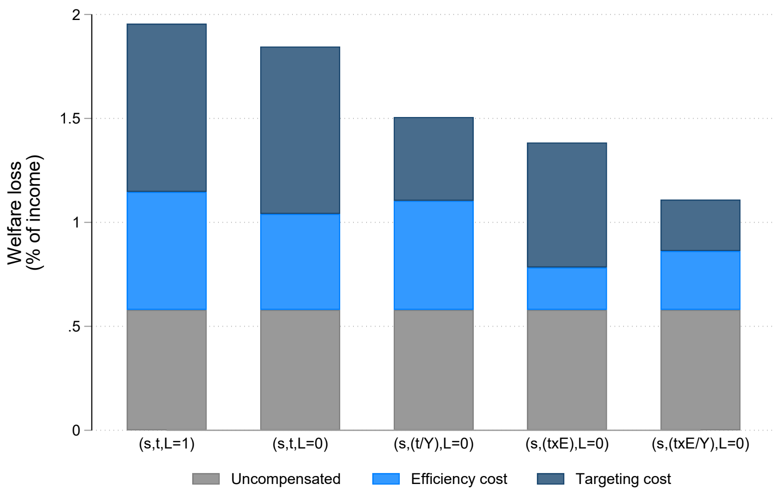

Notes: Panel (a) reports the targeting and efficiency costs associated with policy menus described in the text as $s \geq 0$ *and* $t \geq 0$ *are varied holding public resources expended fixed. The cross denotes the point corresponding to social loss minimising policy. Panel (b) decomposes social losses under loss minimising policy for each policy menu into uncompensated losses (those that would remain if public funds were allocated lump-sum and equated proportional losses), efficiency costs and targeting costs. See Appendix C.5 for decomposition details.*

Another strategy is to base transfers on past energy use. The black and gray lines represent the policy menus when transfers are solely a function of past energy use and a function of past energy in proportion to income, respectively. Both approaches perform substantially better than a policy than entails a universal transfer. Using information on both past energy use and income together yields the best results. In this case, the optimal subsidy rate is 30%, and the welfare loss is limited to 1.11% of income. Of this, 0.53% reflects uncompensated losses – those that cannot be addressed due to the public resource constraint. This policy closes 60% of the gap in losses between observed UK policy and those under an (idealised) system of personalised lump-sum transfers that equate proportional losses.

An advantage of structuring transfers based on past energy usage – especially when adjusted for income – is that it targets households exposed to high proportional losses, reducing reliance on the energy price subsidy. Unlike subsidies, transfers do not distort price signals and they contribute much less strongly to increased carbon emissions, thus avoiding significant efficiency costs. However, even with transfers linked to past energy use, our findings indicate that optimal policy still entails a substantial energy price subsidy. This is because past usage is an imperfect predictor of energy spending 12 months later. Relying solely on these transfers would leave some households – particularly those with significant increases in energy needs – facing substantial losses. Nevertheless, the efficiency costs of an optimally designed policy using transfers based on past energy use are significantly lower than those of policies that do not use information on previous consumption patterns. Conversely, unlike solely income-based transfers, policy that relies only on transfers based on past usage with zero subsidy – represented by the rightmost points of the black and gray lines – outperforms the observed UK policy.

Robustness to varying social preferences

Figure 5.2 illustrates how social losses vary across different policies while holding the total public resource cost constant at that incurred under observed policy. An alternative approach is to quantify how the public resource cost changes when limiting social losses to the level that occurred under observed policy, across different policy menus. In Appendix C.5, we demonstrate that this alternative criterion yields the same social ranking over policy options, and that the UK could have achieved the same level of social losses at 17.5% lower public resource cost under policy menu $(s, (t \times E_i/Y_i), L = 0)$.

The precise quantification of social losses in Figure 5.2 depends on the social preference parameter ψ. In Appendix C.5, we show how this quantification varies with ψ. We also show that as long as the planner has some degree of aversion to large losses – placing more than double the weight on a household experiencing twice the loss of another household with the same income – all our main qualitative takeaways hold, including the social preference ranking over the policy menus, optimal policy entailing a positive subsidy, and the relatively poor performance of income-based transfers alone.

6 Conclusion

Sudden increases in the cost of living can lead to significant hardship for many households, with negative implications for the wider economy and even political stability. Understanding not only the average impact of such shocks but also their *variation* across households is crucial for designing effective relief packages, which governments are often compelled to implement rapidly and with limited information.

Our findings point towards broader lessons for policymakers. Exposure to shocks varies widely and is often only weakly correlated with income. As a result, policies that rely solely on income-based transfers are likely to leave many households facing significant welfare losses. Governments often use price subsidies instead, with examples including "gas tax holidays" in the US and wheat and rice subsidies in developing countries. These measures target support to heavy consumers of affected goods. However, our results suggest that even for an inelastically demanded good such as residential energy, substitution responses are significant enough to mean subsidies incur substantial efficiency costs.

We demonstrate that policymakers may be able to reduce their reliance on subsidies by implementing more targeted transfers that exploit information on households' circumstances. Energy suppliers record information on energy usage, which makes it feasible to base transfers on historic usage. However, in other contexts, information on consumption histories of affected goods may not be readily available. In addition, such history-dependent policies, if used repeatedly, may incentivise higher consumption during normal times. Therefore, there is value in making use of other variables that are observable to the government, correlate with exposure, and on which rapidly deployed transfers may be conditioned. Nonetheless, the substantial variability in household welfare losses suggests subsidies may often remain a valuable component of relief package design.

The use of energy subsidies appears directly at odds with many governments' aspirations to achieve *Net Zero*. Their use stems from the fact that income-based transfers only imperfectly compensate households for energy price increases after they occur. However, reducing reliance on fossil fuels purchased on international commodity markets offers the dual benefit of lowering carbon emissions and reducing exposure to wholesale price volatility along with the associated costs of relief packages discussed in this paper. Energy taxes will likely be needed to incentivise reductions in energy use and investments in green technology. Yet, to the extent that they raise energy prices, policymakers must carefully consider their differential incidence both across income and within groups. Designing policies that ensure the costs of the green transition are shared equitably across households will be key to ensuring its political feasibility.

References

Allcott, H., B. B. Lockwood, and D. Taubinsky (2019). Regressive Sin Taxes, with an Application to the Optimal Soda Tax. *Quarterly Journal of Economics 134*(3), 1557–1626.

Arregui, N., O. Celasun, D. Iakova, A. Mineshima, V. Mylona, F. Toscani, Y. C. Wong, and J. Zhou (2022). Targeted, Implementable, and Practical Energy Relief Measures for Households in Europe. *IMF Working Papers 2022*(262), 1.

Auclert, A., H. Monnery, M. Rognlie, and L. Straub (2023). Managing an Energy Shock: Fiscal and Monetary Policy. *NBER Working Paper 31543*.

Auerbach, A. J. (1985). The theory of excess burden and optimal taxation. *Handbook of public economics Vol. 1*, 61–127.

Bachmann, R., D. Baqaee, C. Bayer, M. Kuhn, A. Löschel, B. Moll, A. Peichl, K. Pittel, and M. Schularick (2022). What if? The Economic Effects for Germany of a Stop of Energy Imports from Russia. *ECONtribute Policy Brief No. 028*.

Batlle, C., P. Heller, C. Knittel, and T. Schittekatte (2024). US federal resource allocations are inconsistent with concentrations of energy poverty. *Science Advances 10*(41).

Beatty, T. K. M., L. Blow, T. F. Crossley, and C. O'Dea (2014). Cash by any other name? Evidence on labeling from the UK Winter Fuel Payment. *Journal of Public Economics 118*, 86–96.

Benhassine, N., F. Devoto, E. Duflo, P. Dupas, and V. Pouliquen (2015). Turning a Shove into a Nudge? A "Labeled Cash Transfer" for Education. *American Economic Journal: Economic Policy 7*(3), 86–125.

Bernheim, B. D. and A. Rangel (2009). Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics. *Quarterly Journal of Economics 124*(1), 51–104.

Blundell, R. and J. M. Robin (1999). Estimation in Large and Disaggregated Demand Systems: An Estimator for Conditionally Linear Systems. *Journal of Applied Econometrics 14*(3), 209–232.

Boardman, B. (1991). *Fuel Poverty: From Cold Homes to Affordable Warmth*. Belhaven Press.

Boehm, J., E. Fize, and X. Jaravel (2025). Five Facts about MPCs: Evidence from a Randomized Experiment. *American Economic Review 115*(1), 1–42.

Borenstein, S. (2012). The Redistributional Impact of Nonlinear Electricity Pricing. *American Economic Journal: Economic Policy 4*(3), 56–90.

Borenstein, S. and J. B. Bushnell (2022). Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency. *American Economic Journal. Economic Policy 14*(4).

Borusyak, K. and X. Jaravel (2024). Are trade wars class wars? The importance of trade-induced horizontal inequality. *Journal of International Economics 159*.

Chen, T., P. Levell, and M. O'Connell (2024). Cheapflation and the rise of inflation inequality. *CEPR Discussion Series DP 19388*.

Chetty, R., A. Looney, and K. Kroft (2009). Salience and Taxation: Theory and Evidence. *American Economic Review 99*(4), 1145–1177.

Choi, J. J., D. Laibson, and B. C. Madrian (2009). Mental Accounting in Portfolio Choice: Evidence from a Flypaper Effect. *American Economic Review 99*(5), 2085–2095.

Christensen, L. R., D. W. Jorgenson, and L. J. Lau (1973). Transcendental Logarithmic Production Frontiers. *Review of Economic and Statistics 55*(1), 28–45.

Consumer Council for Northern Ireland (2025). Home Energy Index | Consumer Council. https://www.consumercouncil.org.uk/research/home-energy-index.

Costa, F. and F. Gerard (2021). Hysteresis and the Welfare Effect of Corrective Policies: Theory and Evidence from an Energy-Saving Program. *Journal of Political Economy 129*(6), 1705–1743.

Department for Business Energy and Industrial Strategy (2022a). Government announces Energy Price Guarantee for families and businesses while urgently taking action to reform broken energy market. https://www.gov.uk/government/news/government-announces-energy-price-guarantee-for-families-and-businesses-while-urgently-taking-action-to-reform-broken-energy-market.

Department for Business Energy and Industrial Strategy (2022b). Quarterly Energy Prices. Technical report.

Department for Business Energy and Industrial Strategy (2022c). Subnational Electricity and Gas Consumption Statistics. Technical report.

Department for Energy and Climate Change (2013). Special feature - Domestic energy bills in 2012. Technical report, Department for Energy and Climate Change.

Department for Energy Security and Net Zero (2023a). 2023 Government Greenhouse Gas Conversion Factors for Company Reporting. Technical report.

Department for Energy Security and Net Zero (2023b). Energy Bills Support Scheme GB: Payments made by electricity suppliers to customers. Technical report.

Department for Energy Security and Net Zero (2023c). Traded carbon values used for modelling purposes, 2023. Technical report.

Department for Energy Security and Net Zero (2024a). Annual fuel poverty statistics in England, 2024. https://assets.publishing.service.gov.uk/media/65ccecba1d939500129466a9/annual-fuel-poverty-statistics-report-2024.pdf.

Department for Energy Security and Net Zero (2024b). UK Energy Trends: June 2024. Technical report.

Dergiades, T. and L. Tsoulfidis (2008). Estimating residential demand for electricity in the United States, 1965–2006. *Energy Economics 30*(5), 2722–2730.

Deryugina, T., A. MacKay, and J. Reif (2020). The Long-Run Dynamics of Electricity Demand: Evidence from Municipal Aggregation. *American Economic Journal: Applied Economics 12*(1), 86–114.

Diewert, W. (1976). Exact and superlative index numbers. *Journal of Econometrics 4*(2), 115–145.

European Central Bank (2024). Fiscal policy measures in response to the energy and inflation shock and climate change.

Fafchamps, M., D. McKenzie, S. Quinn, and C. Woodruff (2014). Microenterprise growth and the flypaper effect: Evidence from a randomized experiment in Ghana. *Journal of Development Economics 106*, 211–226.

Fetzer, T., L. Gazze, and M. Bishop (2023). Distributional and climate implications of policy responses to the energy crisis: Lessons from the UK. *CEPR Working Paper*.

Fetzer, T., C. Shaw, and J. Edenhofer (2024). Informational Boundaries of the State. *SSRN Electronic Journal*.

Finkelstein, A. and N. Hendren (2020). Welfare Analysis Meets Causal Inference. *Journal of Economic Perspectives 34*(4), 146–167.

Fleurbaey, M. and F. Maniquet (2018). Optimal Income Taxation Theory and Principles of Fairness. *Journal of Economic Literature 56*(3), 1029–1079.

Gadenne, L., S. Norris, M. Singhal, and S. Sukhtankar (2024). In-Kind Transfers as Insurance. *American Economic Review 114*(9), 2861–2897.

Hahn, R. W., N. Hendren, R. D. Metcalfe, and B. Sprung-Keyser (2024). A Welfare Analysis of Policies Impacting Climate Change. *NBER Working Paper 32728*.

Hahn, R. W. and R. D. Metcalfe (2021). Efficiency and Equity Impacts of Energy Subsidies. *American Economic Review 111*(5), 1658–1688.

Hansen, L. P. (1982). Large Sample Properties of Generalized Mathod of Moments Estimators. *Econometrica* (4).

Hastings, J. and J. M. Shapiro (2018). How Are SNAP Benefits Spent? Evidence from a Retail Panel. *American Economic Review 108*(12), 3493–3540.

Hines, J. R. and R. H. Thaler (1995). Anomalies: The Flypaper Effect. *Journal of Economic Perspectives 9*(4), 217–226.

Ito, K. (2014). Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing. *American Economic Review 104*(2), 537–563.

Ito, K. and S. Zhang (2024). Do Consumers Distinguish Fixed Cost from Variable Cost? “Schmeduling" in Two-Part Tariffs in Energy. *American Economic Journal: Economic Policy forthcoming*.

Jacoby, H. G. (2002). Is There an Intrahousehold ‘Flypaper Effect’? Evidence from a School Feeding Programme. *Economic Journal 112*(476), 196–221.

Jaravel, X. (2021). Inflation Inequality: Measurement, Causes, and Policy Implications. *Annual Review of Economics 13*(1), 599–629.

Jaravel, X. (2024). Distributional Consumer Price Indices. *mimeo*.

Kaplan, G., G. L. Violante, and J. Weidner (2014). The Wealthy Hand-to-Mouth. *NBER Working Paper w20073*.

Kilian, L. (2008). The Economic Effects of Energy Price Shocks. *Journal of Economic Literature 46*(4), 871–909.

Klick, J. and A. Stockburger (2024). Examining U.S. inflation across households grouped by equivalized income. Technical Report Monthly Labor Review, Bureau of Labor Statistics.

Labandeira, X., J. M. Labeaga, and X. López-Otero (2017). A meta-analysis on the price elasticity of energy demand. *Energy Policy 102*, 549–568.

Labandeira, X., J. M. Labeaga, and J. J. Teixidó (2022). Major Reforms in Electricity Pricing: Evidence from a Quasi-Experiment. *Economic Journal 132*(644), 1517–1541.

Lewbel, A. and K. Pendakur (2009). Tricks with Hicks: The EASI Demand System. *American Economic Review 99*(3), 827–863.

Lewbel, A. and K. Pendakur (2017). Unobserved Preference Heterogeneity in Demand Using Generalized Random Coefficients. *Journal of Political Economy 125*(4), 947–1243.

Lieber, E. M. J. and L. M. Lockwood (2019). Targeting with In-Kind Transfers: Evidence from Medicaid Home Care. *American Economic Review 109*(4), 1461–1485.

Moll, B., M. Schularick, and G. Zachmann (2023). The Power of Substitution: The Great German Gas Debate in Retrospect. *Brookings Papers on Economic Activity 2023*(2), 395–481.

Office for Budget Responsibility (2023). The cost of the Government’s energy support policies. https://obr.uk/box/the-cost-of-the-governments-energy-support-policies/.

Office for Budget Responsibility (2024). Economic and fiscal outlook - March 2024. https://obr.uk/docs/dlm_uploads/E03057758_OBR_EFO-March-2024_Web-AccessibleFinal.pdf.

Office for National Statistics (2021). Estimates of the population for the UK, England, Wales, Scotland, and Northern Ireland. Technical report.

Office for National Statistics (2024). Living Costs and Food Survey. [data series]. 4th Release. Technical Report SN: 2000028, UK Data Service.

Ofgem (2023). Retail Market Indicators. https://www.ofgem.gov.uk/retail-market-indicators#:~:text=The%20cheapest%20tariff%20in%20the,was%20unchanged%20at%20%C2%A318.

Olivi, A., V. Sterk, and D. Xhani (2024). Optimal Monetary Policy during a Cost-of-Living Crisis. *mimeo*.

Pieroni, V. (2023). Energy shortages and aggregate demand: Output loss and unequal burden from HANK. *European Economic Review 154*, 104428.

Piketty, T. and E. Saez (2013). Chapter 7 - Optimal Labor Income Taxation. In A. J. Auerbach, R. Chetty, M. Feldstein, and E. Saez (Eds.), *Handbook of Public Economics*, Volume 5 of *Handbook of Public Economics, Vol. 5*, pp. 391–474. Elsevier.

Reiss, P. C. and M. W. White (2005). Household Electricity Demand, Revisited. *Review of Economic Studies 72*(3), 853–883.

Reiss, P. C. and M. W. White (2008). What changes energy consumption? Prices and public pressures. *RAND Journal of Economics 39*(3), 636–663.

Saez, E. (2002). Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses. *Quarterly Journal of Economics* (117).

Shaffer, B. (2020). Misunderstanding Nonlinear Prices: Evidence from a Natural Experiment on Residential Electricity Demand. *American Economic Journal: Economic Policy 12*(3), 433–461.

Shefrin, H. M. and R. H. Thaler (1988). The Behavioral Life-Cycle Hypothesis. *Economic Inquiry 26*(4), 609–643.

Online Appendix

A Supplementary Details on the Data

A.1 Bank Account Data

A.1.1 Sample Construction

Our main dataset is the *ExactOne Transactional Dataset* collected by ClearScore. ExactOne contains data extracted from individuals' bank accounts and credit card statements. For each transaction, we see the date, amount, merchant, and a description. In many (but not all) cases we also see whether a payment was a regular direct debit payment, a card payment or a standing order. ExactOne includes a set of 150 categories that group similar transactions e.g., 'Energy (Gas, Elec, Other)', 'Food, Groceries, Household', 'Entertainment, TV, Media'. Our focus is on ClearScore users who are responsible for paying the households' energy bills. We start by identifying consumers with at least one account that records spending on energy.

In a first step, we separate energy transactions into those that are "likely prepayment", direct debit, and other payment types (primarily people who pay their bills either monthly or quarterly on receipt of a bill, known as standard credit). We define "likely prepayment" transactions as those payments that are multiples of £5 and less than £100, that were not coded in the data as being direct debit payments (or have a transaction description describing them as direct debit), and that are not single monthly payments regularly paid out on the same date each month.

To validate this method, we check the share of transactions that are classified as likely prepayment for two suppliers: Boost and Utilita. Boost is a solely prepay supplier and we categorise 93% of their transactions as likely prepayment. Utilita is a company that specialises in prepayment (but does offer some credit tariffs), and we class 85% of their transactions as likely prepay. For the other suppliers, who all have both direct debit and prepay customers, the share of transactions we class as likely prepayment is between 5% and 20%.

Direct debit transactions are identified using the bank code variable in the data and/or the transaction description. Most people that pay for their energy by direct debit do so on a monthly basis and 80% of direct debit households have a single monthly payment that covers both gas and electricity. The remaining have separate direct debits for gas and electricity (whether from a single or multiple suppliers).

We define periods of continuous payment as those that have a direct debit payment at least once every 60 days. Within these periods, we define "variable" direct debits as those for whom the payment amount changes on at least half of the months, i.e., at least once every 2 months on average. This could either be because households have a smart meter (which sends usage information directly to the supplier) and bills that adjust automatically, or they send in regular meter readings such that suppliers can ensure their monthly bills accurately reflect their energy usage.

We also identify receipt of the Energy Bill Support Scheme cash refunds in this step, which are defined as receiving a payment of £66 or £67 from an energy supplier in any of the months between October 2022 and March 2023.

We keep periods where households continuously pay for energy for a spell of at least 6 months with the same supplier, using the same payment method. This requires them to purchase energy (by prepayment or direct debit) at least once every two months, or for households that use another payment type, at least once every four months (this covers those paying quarterly). We classify all months in a spell with a particular supplier as prepayment if a majority of months within that spell have over 50% of energy transactions classified as likely prepayment. Otherwise spells are classified as either direct debit or standard credit according to households' observed payment methods.[49]

We then collapse the data to the household-year-month level, excluding around 1% of households who either have more than two energy suppliers or a supplier that exclusively sells heating oil. For those households that use two payment types with different suppliers within a month (e.g., prepay for one and direct debit for another, or two separate direct debits) we keep information on both payment methods.

We define prepay household-year-months as those that prepay for *all* their energy in that month, and variable direct debit household-year-months as those that are on variable billing for *all* their payments in that month. We sum energy spending across payment modes for each household-year-month.

The energy spending of households on variable billing exhibits much more seasonality than those on fixed billing. Figure A.1 shows the log deviations in monthly energy spending across the calendar year in the pre-crisis period (2019 and 2020). For prepayment households and those on variable direct debits, there is pronounced seasonality, with lower spending in the summer than winter. For all other payment types – the vast majority of which are fixed direct debits – spending is approximately stable across the calendar year.

[49]The spell is classified as direct debit if at least one payment in each month of the spell is classified as direct debit. Otherwise it is classified as standard credit. If some months in a spell have direct debit payments, and others do not, the spell is excluded from our sample.

This gives us confidence that energy spending for the variable billing sample is a good measure of energy usage.

Figure A.1: *Seasonality of energy spending by payment type*

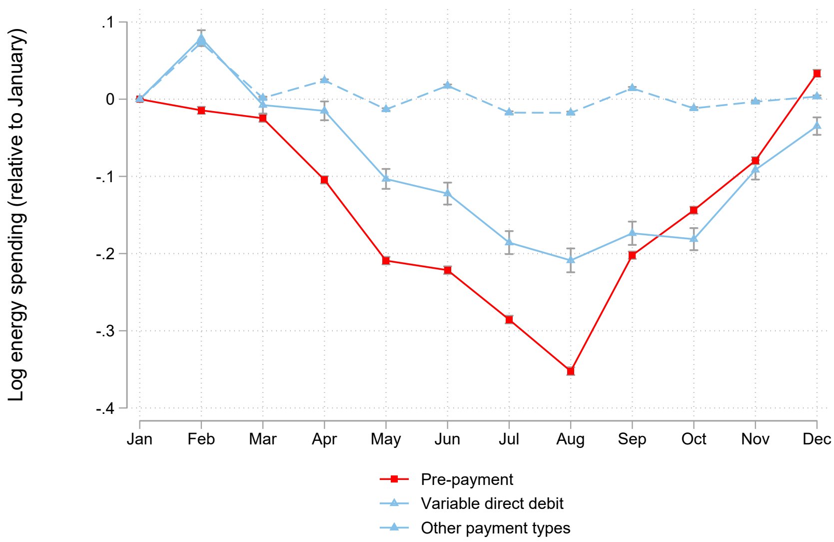

Notes: We estimate deviations in log energy spending in each calendar month, relative to January, over 2019 and 2020. The red line shows this for the prepayment sample, the solid blue line for the variable direct debit sample, and the blue dashed line for all other payment types. We control for household fixed effects. 95% confidence intervals are shown.

For each household-year-month, we also construct measures of non-durable spending and income using the ExactOne data on all transactions associated with their linked accounts. Non-durable spending covers expenditures on: energy, groceries, vehicle fuel (gasoline), discretionary leisure spending (e.g., going out, entertainment), personal services (e.g., hair cuts), phone and TV subscriptions, other bills, transport. We measure income as the sum of inflows into the account, excluding transfers in from other accounts. ExactOne provides information on a banded (gross) salary measure, which allows us to validate our income measure. Although we would not expect our income measure and self-reported salary band to match perfectly (the former is after-tax and includes other transfers and benefits, and potentially the earnings of multiple household members), there is nonetheless a high degree of correlation between the two. ExactOne also provides information on users' sex, age, and geographic area down to postcode sector (each of which covers approximately 3000 households).

In a final step, we drop outliers (1st and 99th percentiles) at the household-year-month level of energy budget share, non-durable spending, and log income. We also require that households are present in the pre-crisis period (2019 and 2020), which allows us to construct our measure of pre-crisis energy spending, and that they are present for at least 6 months between June 2021 and December 2023 (i.e., covering our main estimation period).

Figure A.2: *Correlation of self-reported salary band and mean monthly income*

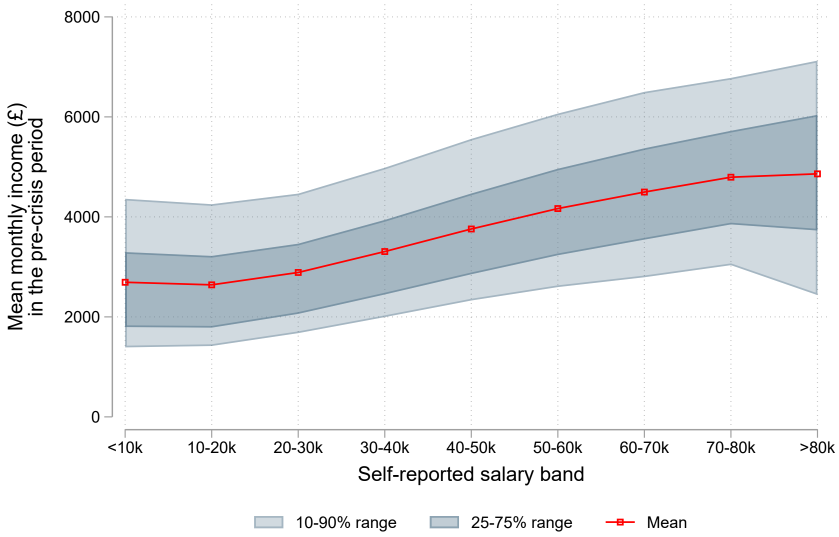

Notes: The figure shows the mean, interquartile and interdecile ranges of mean monthly income (computed over 2019 and 2020), conditional on self-reported salary band of the household member using the ClearScore app.

A.1.2 Representativeness

Table A.1 shows the composition of payment types in our main sample. 25% of our sample are classified as prepayment households – this is higher than the 15% in the nationally representative Living Costs and Food Survey. We reweight all our results to match the share of prepay households in the population.

Table A.1: *Summary statistics of ExactOne sample*

	Analysis sample?	Households No.	Households %	Mean energy spend (£)
Pre-payment	✓	70,270	24.7	107.8
Variable direct debit	✓	17,872	6.3	120.5
Fixed direct debit & other payment modes	✗	196,205	69.0	119.7
Total		284,347	100.0	

Notes: The table shows the number and share of households in our full ExactOne sample over 2019-23, by payment type (constructed as described in Appendix A.1.1). The final column shows the mean energy spend across households in each payment type in 2019. Households can move between payment types over time. Those that pay via pre-payment and variable direct debit constitute our analysis sample.

Figure A.3 compares the age and regional distributions of the ExactOne sample (by payment type) with the distributions in the UK population (Office for National Statistics, 2021). The ExactOne data overrepresents younger individuals and those living in the North West, and underrepresents those in London and the South East. We use the population data to construct weights based on region and 5-year age bands and show robustness of our results to reweighting to match the UK population.

Figure A.3: *Age and geographic sample composition*

(a) Region

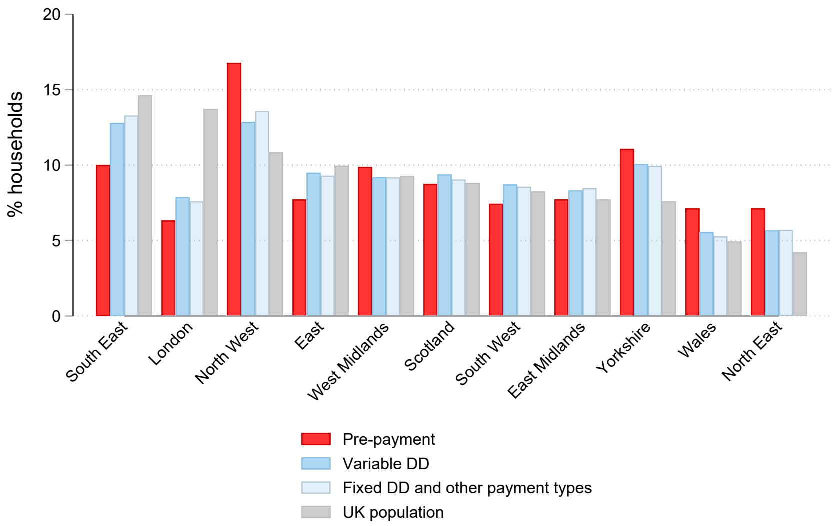

(b) Age

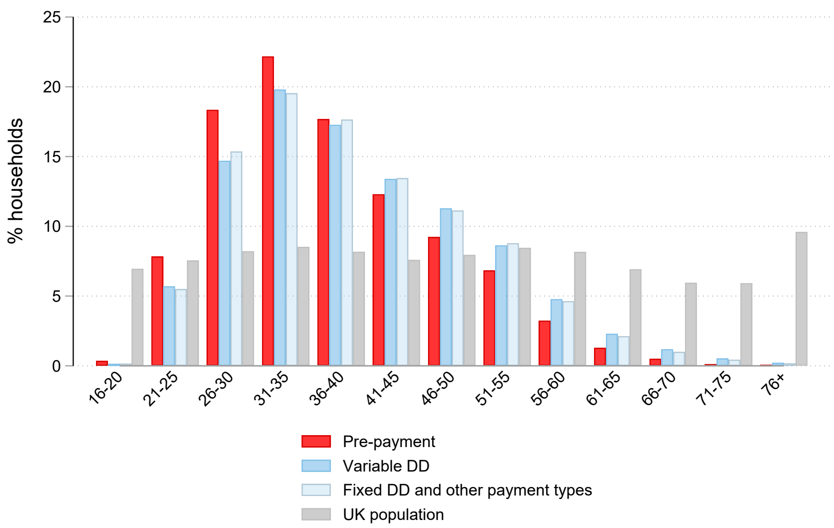

Notes: The top panel shows the fraction of households using each payment type in the ExactOne sample and in the UK population in each region. The bottom panel shows the fraction of households using each payment type in the ExactOne sample and in the UK population in 5-year age bands (measured in 2021).

Figure A.4 shows that the distributions of non-durable and energy spending are similar in the ExactOne data and the LCFS. The main difference between the ExactOne and survey data is that the former are considerably less noisy and it does not exhibit "bunching" at rounded amounts of monthly energy spending.

Figure A.4: *Non-durable and energy spending in ExactOne and Living Costs and Food Survey*

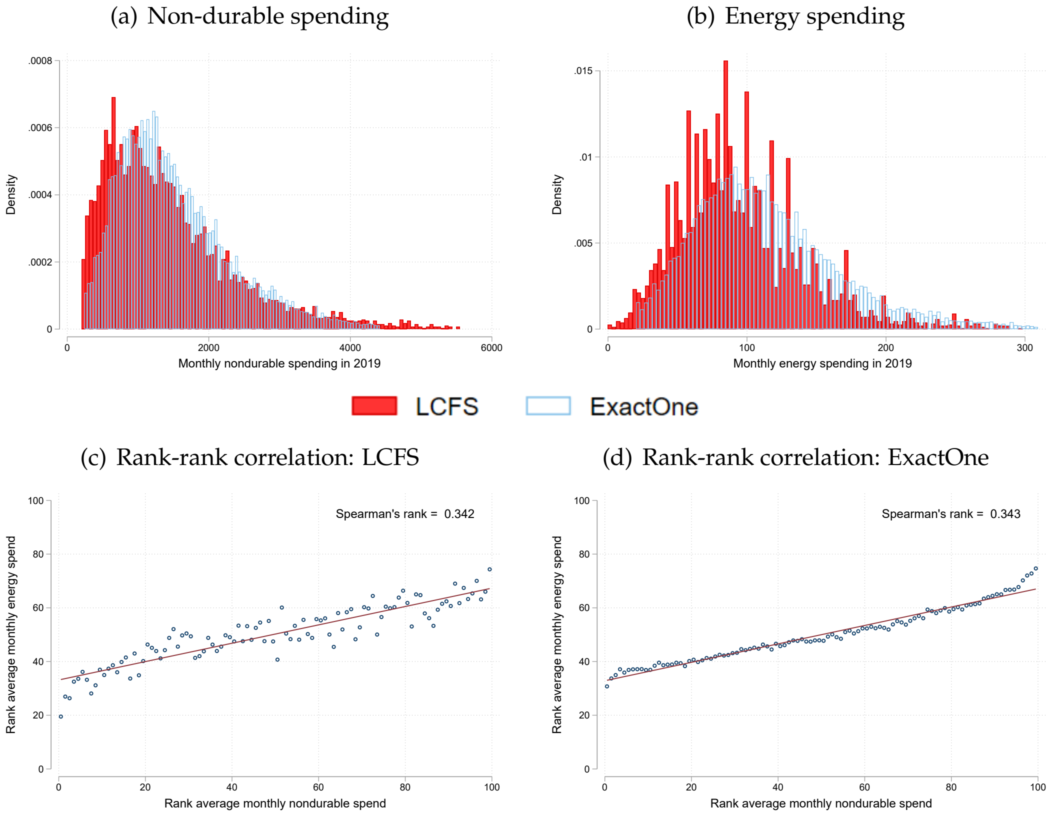

Notes: The top two panels compare the distribution of monthly non-durable spending (top left), and energy spending (top right) in the ExactOne data with that in the Living Costs and Food Survey (LCFS). Distributions are drawn for 2019; for the ExactOne data we take the within-household average of each variable across the months the household is present in the sample. The bottom two panels compare the rank-rank correlations between average monthly non-durable and energy spending in the LCFS (left) and ExactOne data (right). ExactOne data are reweighted by age and region to match the UK population. Spending is expressed in 2022 prices.

Figure A.1 shows that variable billing households have significant seasonality in their energy, giving us confidence that their spending reflects energy usage. It is important, however, that they are otherwise representative of the broader population. Within the variable billing sample, we distinguish between those households on prepayment and those using variable direct debit. Pre-payment households have lower incomes and spend less on energy, on average, than direct debit households – see Figure A.5. They are also younger and more likely to live in northern regions of the UK – see Figure A.3. Variable direct debit households, however, look much more similar in observable terms, to those on fixed direct debits and other payment types. Figure A.5 shows that they have almost identical distributions of income and energy spending, and Figure A.3 shows they are broadly comparable across age and region. This gives us confidence that these households are broadly representative of the more commonly used, fixed billing types. Providing we weight the prepayment and direct debit samples to match the population – which we do everywhere – we are confident that our sample is representative of the wider UK population. This is confirmed in Figure 2.1, which shows that we match the National Accounts measure of energy spending over 2019-2023.

Figure A.5: *Income and energy spend, by payment type*

(a) Income

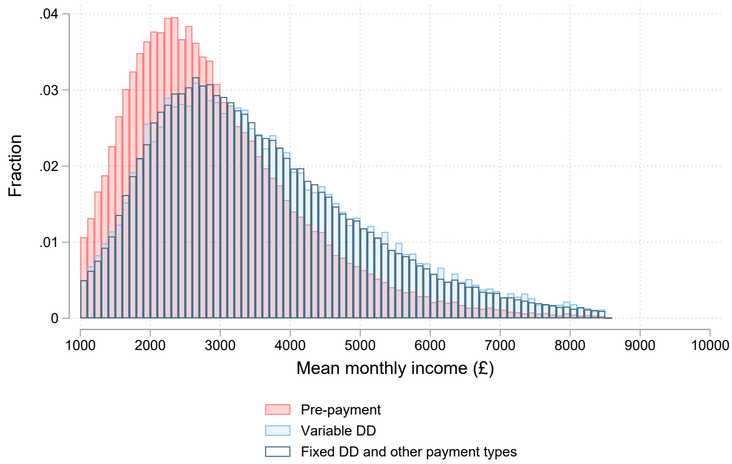

(b) Energy spend

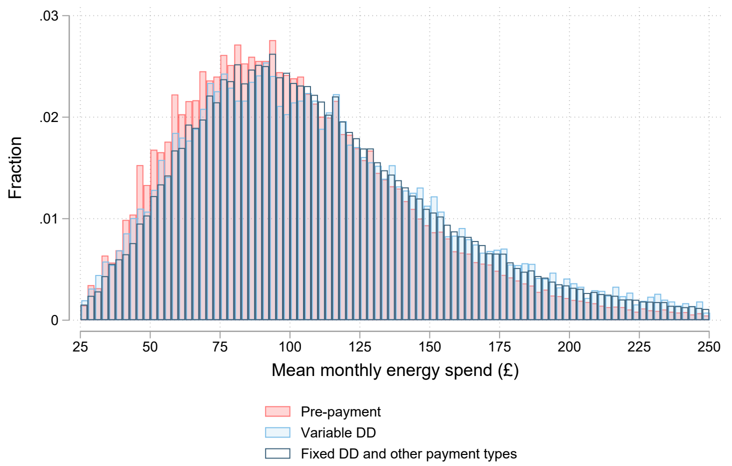

Notes: The top (bottom) panel shows the distributions of mean monthly income (energy spending) calculated over the pre-crisis period (2019-20), for the different payment types in the ExactOne sample. Number of households and observations for each payment type are shown in Table A.1.

A.1.3 Estimation Periods

Table A.2 details the estimation periods used in different parts of the analysis, and associated number of observations and households. We conduct our analysis from June 2021 onwards, since this is when the energy price cap starts to bind for almost all households. For estimation of price elasticities in Section 3.2 we use the period from June 2021 to September 2022 to avoid the confounding effects of transfers, which were introduced in October 2022. To estimate the marginal propensity to consume energy out of transfers we use June 2021 – December 2023 (i.e., the entire period after the price cap became binding for the vast majority of households). For demand estimation, we use the period June 2021 to June 2023, to estimate the model, and the last three months of 2023 to conduct an out-of-sample validation.

Table A.2: *Estimation periods used in the different analyses*

	No. households	No. obs
Analysis sample over full crisis period		
June 2021 – December 2023	74,499	1,262,188
Price elasticities (Section 3.2)		
June 2021 – September 2022	68,683	757,651
Marginal propensity to consume energy (Section 3.3)		
June 2021 – December 2023 (pre-pay only)	60,772	1,098,620
Demand model (Section 4.3)		
June 2021 – June 2023 (in-sample)	72,642	1,086,764
October 2023 – December 2023 (out-of-sample)	24,017	59,235

Notes: The table shows the number of households and number of observations in each of the estimation periods used in our analysis. The first panel pertains to estimation of the price elasticities in Section 3.2, the second panel to estimation of the marginal propensity to consume energy out of transfers in Section 3.3, and the third panel to demand estimation.

A.2 Weather Data

We use data on minimum and maximum monthly temperatures and total monthly rainfall provided by the UK Met Office to control for seasonal and local weather changes that may affect energy demand. This data is collected from 37 weather stations situated across the UK. We interpolate using inverse distance weighting to obtain temperatures and rainfall at the level of Lower Super Output Areas (which have a mean population of 1500), and then merge this information into our ExactOne dataset based on the household's residential location.

A.3 Pre-Crisis Energy Spending

Figure A.6(a) repeats the Figure 2.2(a) from the main paper, with the exception that it is constructed only using data from 2019 (i.e., excluding the COVID-19 pandemic). The two pictures are very similar. Panel (b) summarises the joint distribution of energy spending total non-durable spending. The relationship is similar to that between energy spending and income.

Figure A.6: *Energy spending across the income and total expenditure distributions*

(a) Average energy spending vs income, 2019 only (b) Average energy spending vs total expenditure

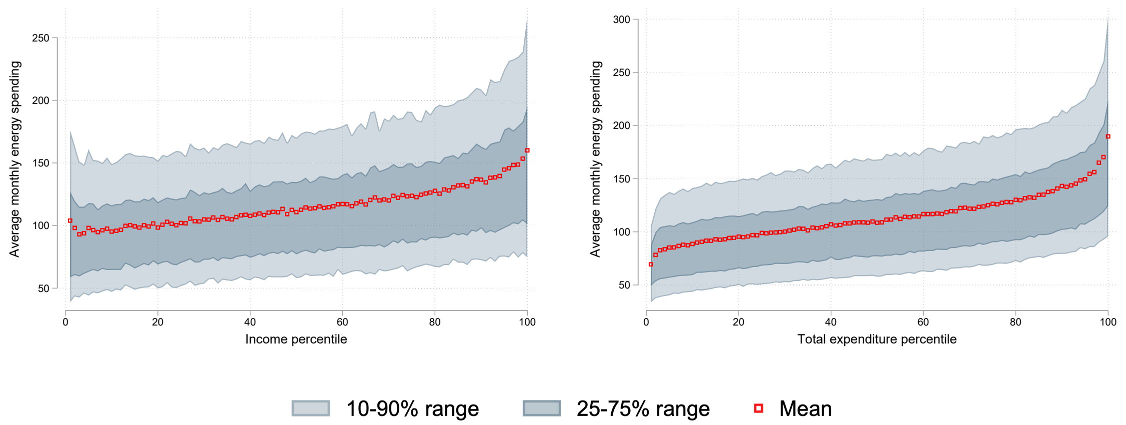

Notes: Panel (a) summarises the distribution (mean and 10th, 25th, 75th, 90th percentiles) of households' mean monthly energy spending conditional on income distribution. The figure is constructing using data from 2019. Panel (b) summarises the distribution (mean and 10th, 25th, 75th, 90th percentiles) of households' mean monthly energy spending in each of the pre-shock years conditional on mean monthly total non-durable spending. Spending is expressed in 2022 prices. The figure is constructing using data from 2019 and 2020. Spending is expressed in 2022 prices.

A.4 Persistence in Energy Spending

We estimate an autocorrelation regression to determine the extent to which energy spending is explained by its yearly lag. For each household, we calculate their energy spending in each year-six month period, where the six month periods are defined as winter (October – March) and summer (April – September). Let x^e_{iyp} denote the average energy spending by household i in year y in period $p = \{$ summer, winter$\}$. We estimate:

$$\log x^e_{ipy} = \rho_0 + \rho_1 x^e_{ipy-1} + \tau_y + \epsilon_{ipy} \tag{A.1}$$

where τ_y are year effects. Table A.3 shows the estimated $\hat{\rho}_1$ and partial R-squared over two time periods. Column (1) focuses just on the pre-crisis period (2019 and 2020); we omit the year effects in this specification. Column (2) shows the results over the whole time period (2019-23), controlling for time effects. We estimate an autocorrelation coefficient of approximately 0.7 in both specifications, with lagged energy spending explaining around half the variation in current energy spending.

Table A.3: *Persistence of energy spending*

	(1)	**(2)**
Autocorrelation coefficient, $\hat{\rho}_1$	0.741	0.708
	(0.003)	(0.001)
(Partial) r-squared	0.549	0.499
Time period	2019-20	2019-23
Year effects	No	Yes

Notes: Table shows estimates of equation (A.1) over the pre-crisis and whole time period. We restrict the sample to households that are present for the full six months within each time period.

B Supplementary Details on Prices, Elasticities and MPCEs

B.1 Measuring the Energy Price and Quantities

Figure B.1 shows the values of the energy price cap and guarantee for prepay consumers compared with those on direct debits (averaging across regional variation). There are small differences between the value of the caps, but the prices essentially co-move.

The price cap is set regionally, but the differences in prices across the regions are small. For example, the average nominal increase in the value of the cap in April 2022 is 54% – the lowest price increase was 50% in the South West and the highest was 57% in London. We use the value of the price cap in the household's region, based on their postcode sector, which we observe in the ExactOne data.

Price index. We construct our consumer energy price index using a fixed-weight Laspeyres index. Let w^g denote the mean share of spending on gas in 2019. The price of energy faced by household i in region $r(i)$ in period t is:

$$p_{r(i)t} = w^g p^g_{r(i)t} + (1 - w^g) p^{el}_{r(i)t} \tag{B.1}$$

where $p^g_{r(i)t}$ is the unit price of gas and $p^{el}_{r(i)t}$ the unit price of electricity in region r. A potential concern is that over the crisis electricity prices rose by more than gas, potentially leading people to switch from electricity to gas. However, Figure B.2 plots the share of total energy from electricity, which is stable over the period before and during the crisis.

Figure B.1: *Energy price cap, energy price guarantee and cheapest available tariffs, 2019-2023*

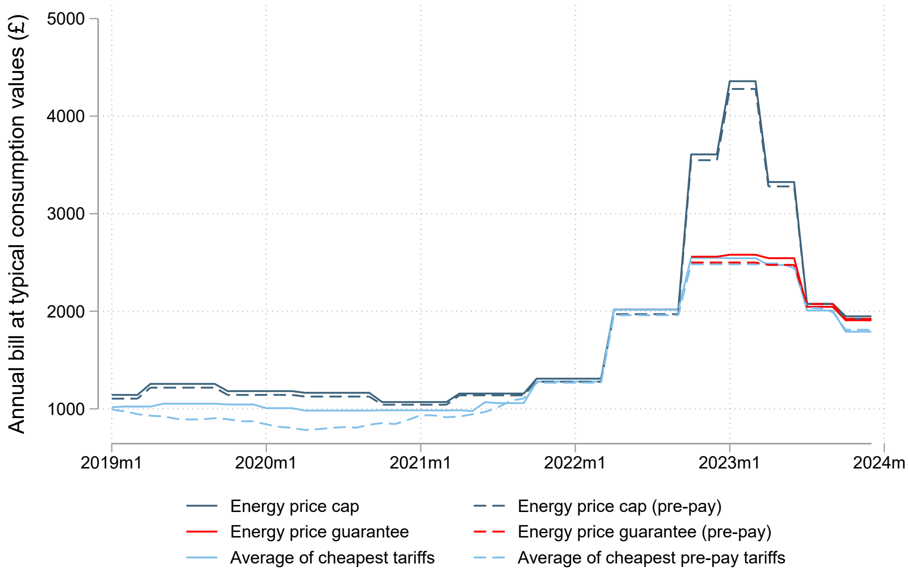

Notes: Data from Ofgem (2023). Figures are costs of an annual bill at 'typical' consumption values of gas and electricity (12,000kWh of gas and 2,900kWh of electricity) for dual fuel direct debit consumers. The average of the cheapest tariffs is a simple average of direct debit tariffs from the 10 suppliers with the lowest cost tariffs (that is, only including one tariff per supplier), including fixed tariffs. Only tariffs that are generally available to consumers are included. The costs of the energy price cap are an average across regions in Great Britain. There are small regional differences in the value of the cap – we use the regional value of the cap when constructing the price index faced by households.

Figure B.2: *Quantity share of total energy from electricity*

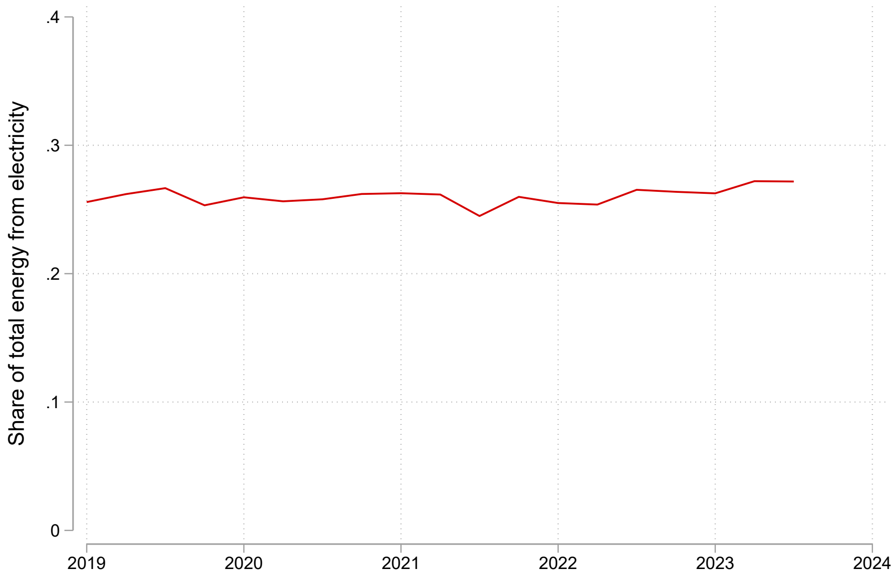

Notes: Data from Department for Energy Security and Net Zero (2024b). The figure shows the share of total energy (gas + electricity) that comes from electricity, measured in tonnes of oil equivalent. Data is at the quarterly level and seasonally adjusted.

This means our fixed weight Laspeyres index (using the annual average 2019 share of energy spending on gas) is very similar to alternative indexes that accommodate substitution responses. Figure B.3 shows that the index we use for our main results is very similar to both a monthly chained Laspeyres index using seasonally adjusted gas and electricity quantities as weights, or a Fisher index (a geometric average of a Laspeyres index using January 2019 weights and a Paasche index using current quantities as weights). Note the Fisher index allows for substitution between electricity and gas; it is second-order approximation to an arbitrary homothetic sub cost-of-living index (Diewert, 1976).

Figure B.3: *Alternative price indexes*

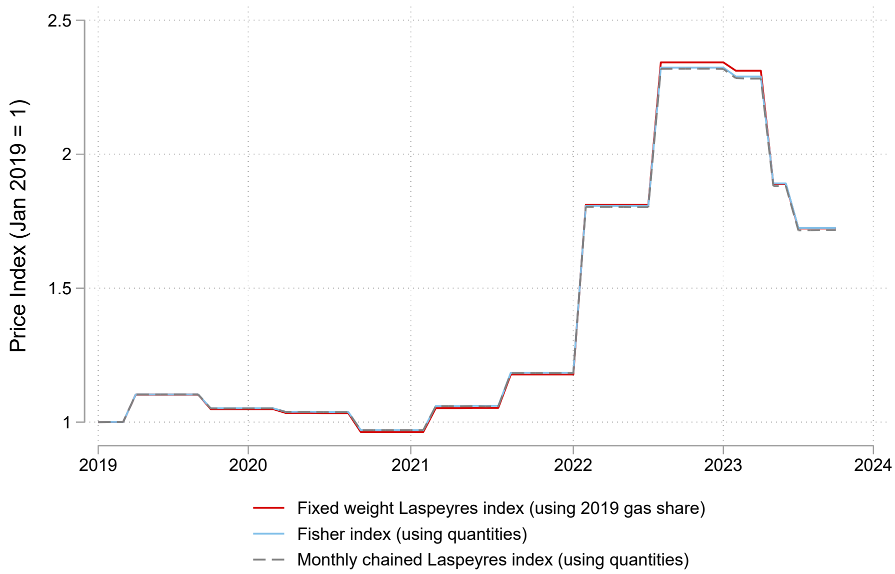

Notes: Price indexes calculated using data on residential gas and consumption (quantities) from Department for Energy Security and Net Zero (2024b) or the share of energy spending on gas and electricity from the Living Costs and Food Survey (2019). The "fixed weight Laspeyres index (using 2019 gas share)" is the index used for our main results.

Household-varying price indexes. To account for variation across households in their use of electricity and gas, we also calculate versions of the Laspreyres index (B.1) with budget shares that vary by income and energy spending quintile. We calculate these using the 2019 version of the Living Costs and Food Survey. Figure B.4 shows these shares: they are similar (at around 0.45) for all quintiles, except for households in the bottom quintile of energy spending, who devote a smaller fraction of their spending to gas.

Variable energy spending. Let $\tilde{x}^e_{i\tau}$ denote the energy spending by household i in year-month τ. We obtain variable energy spending, $x^e_{i\tau}$ by subtracting the fixed costs associated with purchases gas and electricity: $x^e_{i\tau} = \tilde{x}_{i\tau} - F^g_{i\tau} - F^e_{i\tau}$, where $F^g_{i\tau}$ and $F^e_{i\tau}$ denote the standing charges associated with gas and electricity use.

Figure B.4: *Share of energy spending on gas, by income quintile and energy spending quintile*

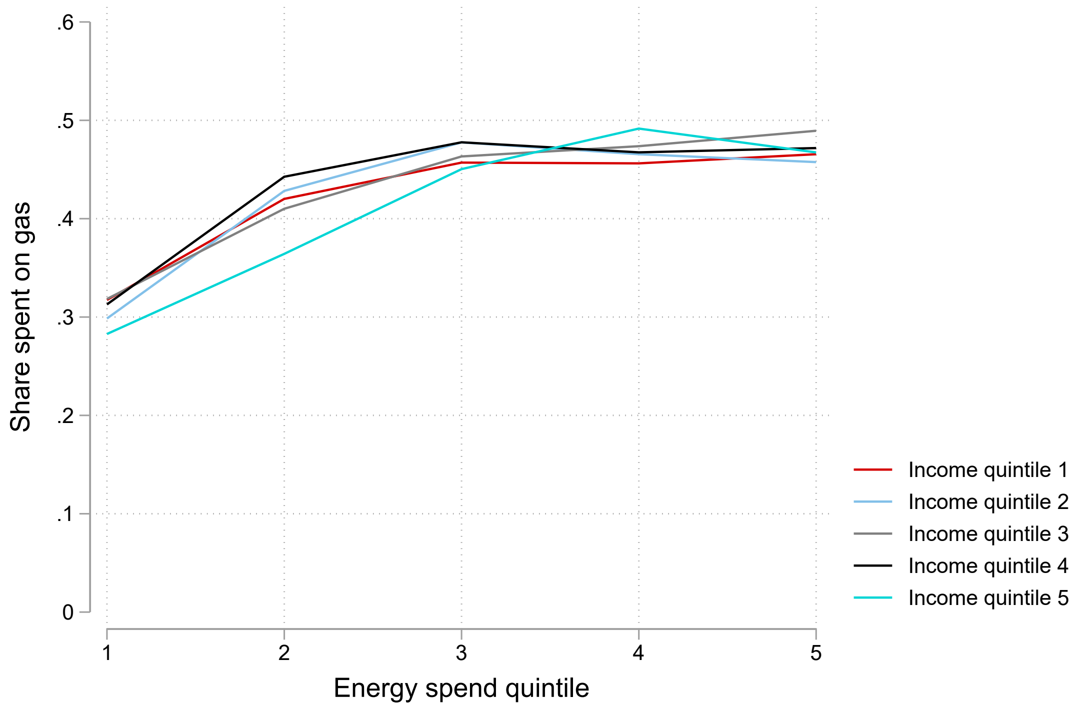

Notes: Data from the 2019 Living Costs and Food Survey.

Implied spending from energy-support transfers. The UK government paid rebates on household energy bills of £400 in monthly installments of either £66 or £67 from October 2022 to March 2023 (under the "Energy Bill Support Scheme"). Direct debit consumers either received these as credit added to their account or as a refund in each billing period. We can identify those that received refunds from the bank account data as those who received a credit equal to £66 or £67 in the each of the months transfers were paid – approximately 20% of direct debit consumers received a refund. Standard credit and prepayment consumers that used smart meters had credit allocated directly to their meters. Pre-payment consumers with traditional meters were sent vouchers that they could use to top up their meters (which could be redeemed up to 30 June 2023). The vast majority of these were redeemed in the period they were issued – aggregate statistics suggest that over 90% of pre-payment vouchers were claimed (Department for Energy Security and Net Zero (2023b)). Letting $R_{i\tau}$ denote the value of the transfer issued in each month (either £66 or £67), we construct a transfer-inclusive measure of energy spending, $\tilde{\tilde{x}}^e_{i\tau}$ as follows:

$$\tilde{\tilde{x}}^e_{i\tau} = \begin{cases} x^e_{i\tau} & \text{if consumer received a direct cash refund} \\ x^e_{i\tau} + 0.9 \times R_{i\tau} & \text{if the consumer was a prepayment household} \\ x^e_{i\tau} + R_{i\tau} & \text{otherwise} \end{cases}$$

where we scale the transfer value by 0.9 for prepayment households to reflect the fact that redemption by these consumers was 90%.

Quantity measurement. Quantity, q^e_{it} is then calculated as follows:

$$q^e_{i\tau} = \frac{\tilde{\tilde{x}}^e_{i\tau}}{p_{r(i)\tau}}$$

where $p_{r(i)\tau}$ denotes the price index described above.

B.2 Price Elasticities

In Figure B.5 we replicate Figure 3.3 from the main paper, with the exception that we allow the gas and electricity weights in the price index for energy to vary across the pre-shock income and energy spending quintiles (computed using the LCFS). It shows that the main patterns of variation in elasticities across the 25 groups are unaffected.

Figure B.5: *Heterogeneity in elasticities, estimated using pre-shock energy spending and income specific price index weights*

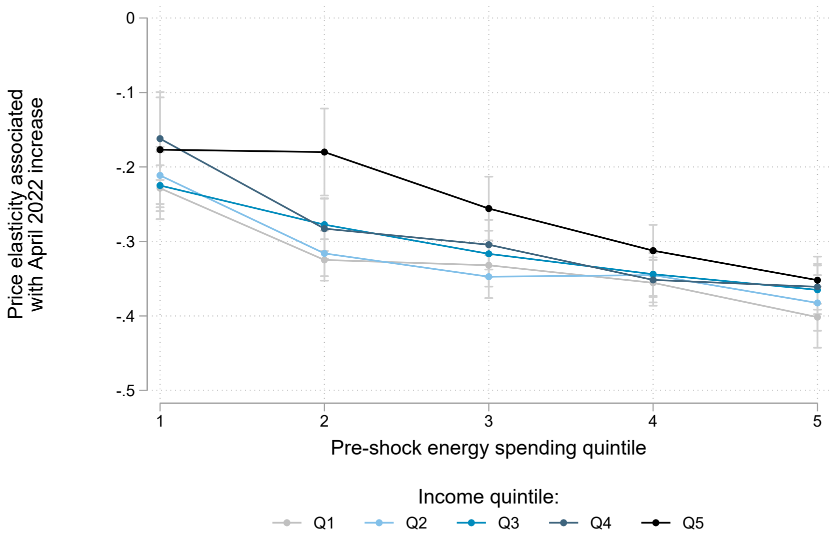

Notes: Figure shows the estimates of the average price elasticity over the April 2022 price increase, for each income quintile and pre-shock energy consumption quintile. We allow for heterogeneity in price responsiveness across payment type and average these using weights that vary across income and pre-shock energy consumption quintile. We use the price index constructed with weights for gas and electricity that are pre-shock income and energy spending group specific (see Appendix B.1 for details).

A potential concern with our main specification is that our seasonal and weather variables do not fully control for time-varying factors that correlate with changes in the regulatory price cap and also influence energy demand. To address this, we estimate an alternative specification that incorporates energy spending data from Northern Ireland (NI), which was not subject to the same regulatory price cap as the UK. We exclude approximately half of Northern Irish households that we observe buying heating oil (which is common source of residential in NI, but not in the rest of the UK) focusing on those that use electricity and gas and are therefore more comparable to the rest of our sample. Figure B.6(a) illustrates the differential price variation in the two jurisdictions over this period: the energy price in NI rose much more gradually over this period, with no discrete jumps in October 2021 and April 2022. Figure B.6(b) shows that over the first half of 2021, (un-deseasonalised) spending in the two countries broadly evolved in parallel. In April 2022, the two series diverge, with a big jump in spending in the UK as the cap increased, and no corresponding change in NI (un-deseasonalised spending in NI actually fell, reflecting the seasonal change in Spring). Table 3.1 shows that when we use this differential geographic price variation, incorporating a full set of year-month dummies in our estimating equation, we estimate a very similar elasticity to our main specification.

Figure B.6: *Energy prices and spending in Northern Ireland relative to the rest of the UK*

(a) Prices (b) Spending

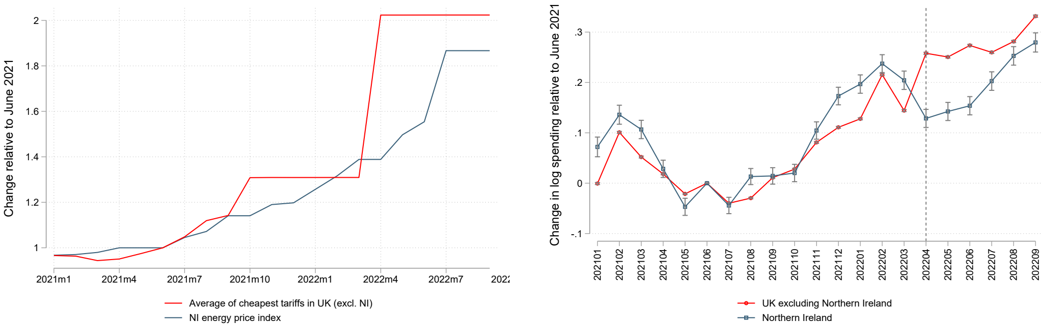

Notes: Panel (a) compares the average price of energy in Northern Ireland with the rest of the UK between January 2021 and September 2022. The price series for Northern Ireland is a composite index for electricity and gas calculated by the Consumer Council for Northern Ireland (2025). Panel (b) shows the change in average un-deseasonalised spending across households in Northern Ireland (excluding those that we observe buying heating oil) compared with the rest of the UK. Here we use information on all payment types, and do not restrict to the variable payment sample.

B.3 Marginal Propensity to Consume Energy Out of Transfers: Direct Debit Households

In Section 3.3 we present evidence that households that prepay for their energy have a substantially higher marginal propensity to consume energy (MPCE) out of energy-support transfers (distributed to them as meter top-ups or vouchers from electricity suppliers) than out of alternative cash transfers. Here we conduct a similar analysis for direct debit households, and show that there while there is some evidence of their MPCEs varying by transfer type, the effects are much weaker than for households who use prepayment.

Households that pay their energy bills via direct debit either received their £400 energy-support transfers as a credit on their bills or via a direct cash refund from their energy company. We start by considering how spending was affected for households receiving their transfers as cash. The log spending of these individuals is shown as the solid line in Panel (a) of Figure B.7. This figure, and the subsequent regression analysis, includes households with (relatively) fixed direct debits alongside the variable direct debit sample we use in our main analysis. This is because, once we split the variable sample into those that receive transfers in cash or credit, the sample sizes in the two groups are smaller and the estimates noisier. The figure shows that direct debits for those receiving the transfer as cash rose steadily over the crisis period.

The dashed line in Panel (a) of Figure B.7 shows predicted spending over October 2022 to March 2023 based on price responses outside of this period. We make these predictions in the same way as for prepayment households (shown in Figure 3.4): by regressing de-seasonalised spending on log prices, weather controls, and a dummy variable for the transfer period, and then plotting results setting the transfer dummy to zero. To account for the possibility that direct debit payments depend on past and future prices, we also include lags and leads of log prices alongside current prices. The figure shows that actual spending and predicted spending, based on prices responses, are very similar, indicating little evidence of flypaper effects for this group.

To further quantify this, we estimate a variant of equation (3.2) for households that received the energy-support transfer as cash (as above including lags and leads of prices as additional controls). We allow for the possibility that higher consumption during the transfer period was paid for in the following quarter by accounting for spillover effects with a post-transfer indicator (for the months of April to June 2023). We then calculate MPCEs analogously to how we described constructing them for prepayment households. As before, we also show how the MPCE out of transfers compares with those for the cost of living payments. Column (1) of Table B.1 shows that the implied MPCE out of energy-support transfers received in cash is small and close to zero. There is no evidence of a flypaper effect for this group of households.

We study the impact of transfers for households who received them as account credit with their energy company by comparing them to those who received them as cash, shown in Figure B.7(b). For those households that received the energy-support transfers as credit, we add the value of the transfer to their observed monthly spending, which gives a gross measure comparable to those who received the transfer as cash. Prior to the introduction of the transfers, the levels of spending for the two groups evolved in a similar way. Following the introduction of the transfers in October 2022 (concurrent with the price rise), gross spending for those receiving transfers as credit noticeably increased, relative to those receiving transfers as cash. Once the transfer period ended in March 2023, gross spending by the credit households fell sharply. By the end of 2023, spending for these two groups had reverted to the pre-transfer period difference. This pattern suggests that those receiving the transfers as credit initially accumulated this as surplus on their accounts, which they subsequently drew down after the transfer period ended.

Figure B.7: *Energy spending over the transfer period, direct debit households*

(a) Direct debit consumers who received transfers as cash, actual spending vs predicted based only on price responses

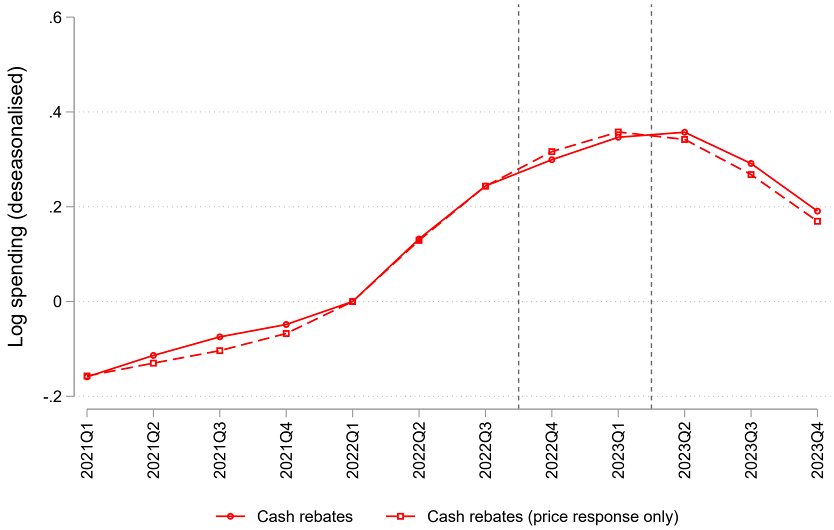

(b) Direct debit consumers who received transfers as cash vs credit

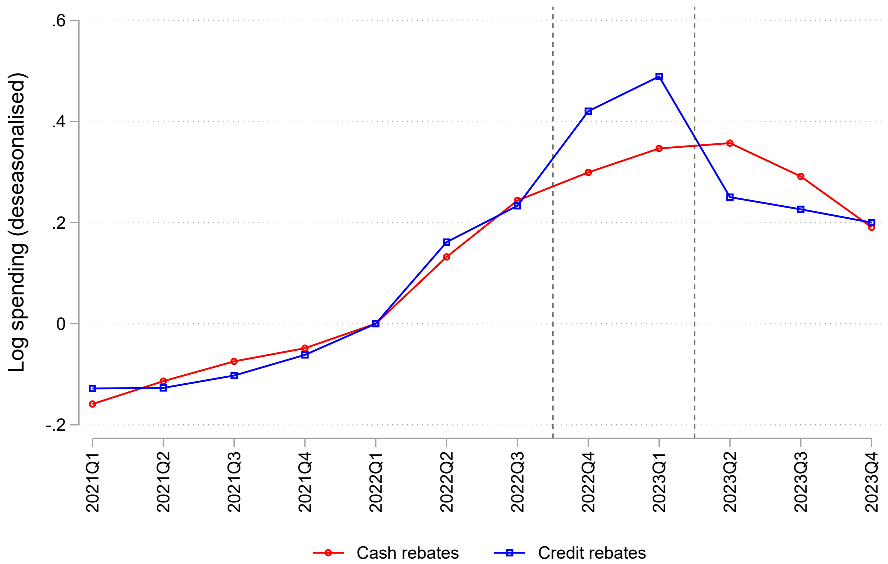

Notes: The dashed vertical lines indicate the start and end of period when energy-support transfers were distributed via electricity suppliers. The solid red line in Panel (a) shows the evolution of actual deseasonalised log spending. The dashed line shows log spending as predicted by a fixed-effects panel regression, that controls for weather effects, estimates price effects using variation outside of the transfer period only and controls for (and removes) time effects for the transfer period. Panel (b) shows log deseasonalised energy spending (inclusive of transfers) for direct debit consumers who either received energy-support transfer as cash or as account credit.

To estimate whether this intertemporal smoothing fully accounted for any increase in spending over the transfer period, we estimate a differences-in-differences specification, with credit receipt as our treatment group and cash receipt as our comparator group:

$$x^e_{i\tau} = \sum_\tau \text{time}_\tau \times \text{credit}_i + \sum_\tau \text{time}_\tau + \zeta_i + \nu_{i\tau}, \tag{B.2}$$

where time_τ are year-month effects, and credit_i is an indicator variable for whether the transfers were paid to the consumer as credit. The interaction terms $\sum_\tau \text{time}_\tau \times \text{credit}_i$ tells us how much greater energy spending was for those receiving payments as credit than for those who received it as cash. Summing the value of these coefficients and dividing by the value of transfers received yields an estimate of the MPCE for those receiving transfers as credit. The second column of Table B.1 reports the results. There is a small difference in cumulative spending between these two groups. For those who receive transfers via credit we estimate that the MPCE from energy-support transfers is 0.09 higher than for those receiving transfers as cash. Although this is indicative of a modest flypaper effect, it is much smaller than the effect we estimate for households that pay via prepayment.

Table B.1: *MPCE estimates for transfers (households paying by direct debit)*

	(1)	**(2)**
Transfer administered by energy companies	-0.03	0.09
	(0.007)	(0.005)
Transfer administered by government, cash	-0.01	
	(0.003)	
Sample	Cash rebates DD	All DD

Notes: Column (1) shows implied marginal propensities to consume energy (MPCE) out of energy-support transfers transfers and cost of living payments for households paying for their energy via direct debit, and who receive the energy-support transfer in cash. Column (2) shows the differences in the MPCE out of energy-support transfer, for direct debit paying households, between those that receive them as credits and those that receive them as cash. Standard errors are shown in parentheses, clustered at the household-level. Regressions include weather controls and household fixed effects. The regression in column (1) includes lags and leads of log real energy prices. The dependent variable in both columns is deseasonalised log energy spending.

B.4 Carbon Emissions

We calculate carbon emissions per £ spent on energy in the first quarter of 2023 as follows. We take total residential consumption of gas and electricity in that period from Table 4.1 and 5.2 of Department for Energy Security and Net Zero (2024b) and multiply this by the unit prices for that period to arrive at total electricity and gas spending. We then calculate the fraction of each £ spent over this period that is electricity and gas (54% is gas and 46% is electricity). We also obtain the number of kWh of electricity and gas per £ spent. Multiplying the latter by the carbon intensity of grid average electricity consumption and gas consumption then gives $CO2$ emissions per £ of electricity and gas. These are taken from Tables 9 and 10 of Department for Energy Security and Net Zero (2023a). Summing these two gives us $CO2$ emissions per £ of energy spending, yielding an estimate of 1.24kg of $CO2$

for each £. We then assume a social cost of carbon of £59 per tonne (taken from Department for Energy Security and Net Zero (2023c)) to obtain the cost of the carbon externality.

B.5 Energy Poverty

We use an affordability measure of energy poverty (Boardman, 1991), which classifies a household as "energy poor" if they spend more than a set percentage of their annual income on energy. In the UK, this threshold is set at 10% of after-housing-costs income.[50] In the US, the accepted threshold for high energy burden is 6%, on the basis that housing costs should not exceed 30% of income and utility costs should not exceed 20% of housing costs (Batlle et al., 2024).

We do not observe (reliably for our whole sample) housing costs in the ExactOne data, so we use the Living Costs and Food Survey to impute after-housing-costs income in our sample. Using the LCFS, we construct the ratio of after-housing-costs income, y^{AHC} to non-durable spending, x, which we denote by $\tilde{y}^{AHC} = y^{AHC}/x$. We regress this on indicator variables for income, non-durable spending, and energy spending deciles and calculate the predicted values, $\tilde{y}_d^{AHC}$, where d indexes the triad of income, non-durable spending, and energy spending decile. We match these ratios into the ExactOne data and multiply by the household's non-durable spending to get their predicted after-housing-cost income:

$$y_{it}^{AHC} = \tilde{y}_{d(i)}^{AHC} x_{it}.$$

Figure B.8(a) shows the correlation between this imputed AHC income and the income measure in the ExactOne data. We define an energy poverty indicator equal to 1 if energy spending, e_{it} as a share of after-housing-cost income, y_{it}^{AHC}, is greater than 10%.

In 2019, we estimate that 3.9 million (14%) households are in energy poverty, which aligns with the official statistics provided by Department for Energy Security and Net Zero (2024a). Figure B.8(b) shows the share of households in energy poverty across the income distribution in 2019. More than 40% of households in the bottom income decile are energy poor, compared with less than 4% of those in the top decile.

[50] There was a recent change to the official measure of energy poverty in England, which is now based on the energy efficiency of the home and the household's income. However, the UK government continue to produce the 10% affordability measure, which is also still used in Scotland, Wales and and Northern Ireland for official energy poverty statistics (Department for Energy Security and Net Zero, 2024a).

Figure B.8: *After-housing costs income and energy poverty*

(a) Imputed AHC income versus observed income (b) % households in energy poverty

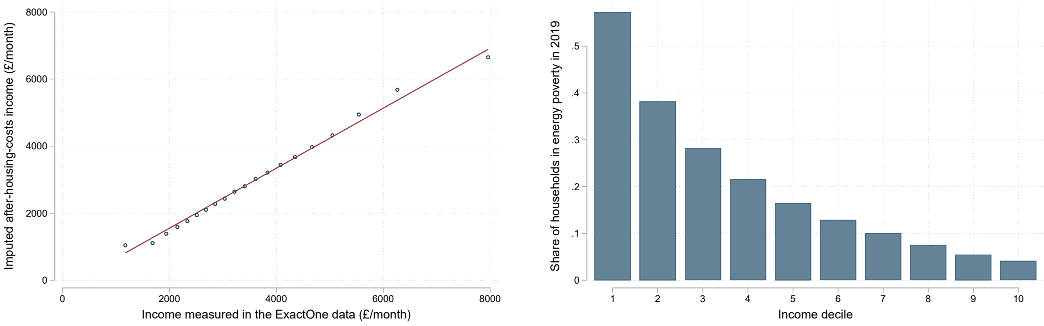

Notes: The left-hand panel shows the correlation between our imputed after-housing-costs (AHC) income measure, constructed as described in the text, and the household's average income in the ExactOne data in 2019. The right-hand panel shows the share of households in energy poverty (i.e., those for whom energy spending is more than 10% of AHC income) across income deciles (based on the ExactOne measure) in 2019.

C Supplementary Details on the Welfare Framework and Results

C.1 Empirical Demand Model

Let i index households and τ year-months. Let $(p^e_{r(i)\tau}, p^n_\tau)$ denote energy and other non-durable prices, $x_{i\tau}$ denote the household's total net-of-fixed-fee budget and $\theta_{i\tau} = (\mathbf{z}_{i\tau}, \Psi)$ collect conditioning variables $\mathbf{z}_{i\tau}$ and parameters Ψ. $\mathbb{d}_{i\tau}$ is an indicator variable that denotes whether a household is influenced by a flypaper effect; this equals 1 for prepay households during the period when the energy-support transfer is distributed.

Empirical counterpart of problem (4.1)

Consider the case when $\mathbb{d} = 0$. This corresponds to the situation where the household makes privately optimal choices (and to observed choice behaviour outside of prepay households during the energy-support transfer period). We denote the corresponding indirect utility function $u_{i\tau} = V(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, \theta_{i\tau})$. Let $x_{i\tau} = \chi(p^e_{r(i)\tau}, p^n_\tau, u_{i\tau}, \theta_{i\tau}) \equiv (V)^{-1}(p^e_{r(i)\tau}, p^n_\tau, ., \theta_{i\tau})$ denote the expenditure function. We specify this as:

$$\begin{aligned}\log \chi = \log u_{i\tau} &+ (A + \sum_{l \in \mathcal{Z}_1} A_l z_{i\tau l}) \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right) + \frac{1}{2}(B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l}) \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right)^2 \\ &+ \left((C_1 + \sum_{l \in \mathcal{Z}_2} C_{1l}) \log u_{i\tau} + \sum_{k>1} C_k (\log u_{i\tau})^k\right) \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right) \\ &+ \frac{1}{2} D \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right)^2 \times \log u_{i\tau} \qquad \text{(C.1)}\end{aligned}$$

We write the log expenditure function directly as a function of the log price difference between the two goods, which, in a two-good demand system, ensures the associated demands satisfy adding up (budget shares sum to one), homogeneity in prices and budgets and Slutsky symmetry. If $D = 0$ the expenditure function belongs to the translog class (Christensen et al., 1973). When $D \neq 0$, the first-order Engel curve and relative price terms interact, and the expenditure function belongs to the EASI class (Lewbel and Pendakur, 2009).

By Shephard's Lemma, the corresponding Hicksian energy budget share demand is:

$$\begin{aligned}\omega_{i\tau} =& (A + \sum_{l\in\mathcal{Z}_1} A_l z_{i\tau l}) + (B + \sum_{l\in\mathcal{Z}_2} B_l z_{i\tau l}) \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right) + (C_1 + \sum_{l\in\mathcal{Z}_2} C_{1l}) \log u_{i\tau} + \\ & \sum_{k>1} C_k (\log u_{i\tau})^k + D\left(\log p^e_{r(i)\tau} - \log p^n_\tau\right) \times \log u_{i\tau}\end{aligned}$$

Substituting this into equation (C.1) and rearranging yields:

$$\begin{aligned}\omega_{i\tau} =& (A + \sum_{l\in\mathcal{Z}_1} A_l z_{i\tau l}) + (B + \sum_{l\in\mathcal{Z}_2} B_l z_{i\tau l}) \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right) + (C_1 + \sum_{l\in\mathcal{Z}_2} C_{1l}) y_{i\tau} + \\ & \sum_{k>1} C_k y^k_{i\tau} + D\left(\log p^e_{r(i)\tau} - \log p^n_\tau\right) \times y_{i\tau} \end{aligned} \tag{C.2}$$

$$y_{i\tau} = \frac{\log x_{i\tau} - (\omega_{i\tau} \log p^e_{r(i)\tau} + (1-\omega_{i\tau}) \log p^n_\tau) + \frac{1}{2}(B + \sum_{l\in\mathcal{Z}_2} B_l z_{i\tau l}) \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right)^2}{1 - \frac{1}{2} D \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right)^2}, \tag{C.3}$$

where $y_{i\tau} = \log u_{i\tau}$. Lewbel and Pendakur (2009) refer to these as the *implicit* Marshallian budget share demand and indirect utility functions. Together they define the (privately optimal) Marshallian budget share energy demand $\omega_{i\tau} \equiv \omega^0(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, \mathbf{z}_{i\tau}; \Psi)$.

Empirical demands

Our demand specification (equations (4.2) and (4.3)) accommodates choices when $\mathbb{d} = 1$, corresponding to a flypaper effect. We do this by augmenting equation (C.2) with $\big(\delta + \sum_{l\in\mathcal{Z}_3} \delta_l z_{i\tau l}\big)\mathbb{d}_{i\tau}$. For completeness we repeat equations (4.2) and (4.3) here:

$$\begin{aligned}\omega_{i\tau} =& (A + \sum_{l\in\mathcal{Z}_1} A_l z_{i\tau l}) + (B + \sum_{l\in\mathcal{Z}_2} B_l z_{i\tau l}) \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right) + (C_1 + \sum_{l\in\mathcal{Z}_2} C_{1l}) y_{i\tau} + \\ & \sum_{k>1} C_k y^k_{i\tau} + D\left(\log p^e_{r(i)\tau} - \log p^n_\tau\right) \times y_{i\tau} + (\delta + \sum_{l\in\mathcal{Z}_3} \delta_l z_{i\tau l})\mathbb{d}_{i\tau}\end{aligned}$$

$$y_{i\tau} = \frac{\log x_{i\tau} - (\omega_{i\tau} \log p^e_{r(i)\tau} + (1-\omega_{i\tau}) \log p^n_\tau) + \frac{1}{2}(B + \sum_{l\in\mathcal{Z}_2} B_l z_{i\tau l}) \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right)^2}{1 - \frac{1}{2} D \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right)^2},$$

Together these equations define the budget share energy demand $\omega_{i\tau} \equiv \omega(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, \mathbb{d}_{i\tau}, \mathbf{z}_{i\tau}; \Psi)$, which we estimate. While this function does not have an analytical formula it can be obtained numerically from recursively solving equations (4.2) and (4.3). A sufficient condition for these equations to uniquely define $\omega_{i\tau}$ and $y_{i\tau}$ is that the Jacobian of the system of equations is non-singular, which is satisfied if equation (C.5) below is satisfied.[51]

Inequality constraints

Two additional regularity conditions, implied by consumer theory, that we do not impose during estimation are that the function $\chi(p^e_{r(i)\tau}, p^n_\tau, u_{i\tau}, \theta_{i\tau})$ is concave in prices and increasing in utility. This implies the inequality restrictions:

$$\omega_{i\tau}^2 - \omega_{i\tau} + B + \sum_{l \in \mathcal{Z}_2} B_l z_l + D y_{i\tau} < 0 \tag{C.4}$$

$$\left(\log p^e_{r(i)\tau} - \log p^n_\tau\right)\left((C_1 + \sum_{l \in \mathcal{Z}_2} C_{1l}) + \sum_{k>1} C_k k y_{i\tau}^{k-1} + \frac{1}{2} D \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right)\right) > -1, \tag{C.5}$$

which we check are satisfied post estimation.

C.2 Estimation

GMM estimator. We estimate the model using GMM. Let $w_{i\tau}$ denote the observed budget share of household i and period τ. We specify the population moment condition:

$$\mathbb{E}(g_{i\tau}(\Psi)) = 0 \tag{C.6}$$

where $g_{i\tau}(\Psi) \equiv (w_{i\tau} - \omega(\mathbf{p}_\tau, x_{i\tau}, \mathbb{d}_{i\tau}, \mathbf{z}_{i\tau}; \Psi))\mathbf{h}_{i\tau}$ and $\mathbf{h}_{i\tau}$ is a vector of instruments. In practice, we specify the instrument vector such that $|\mathbf{h}_{i\tau}| = |\Psi|$. We obtain the estimate $\widehat{\Psi}$ from:

$$\widehat{\Psi} = \arg\min_{\Psi} \bar{g}(\Psi)' W \bar{g}(\Psi)$$

where $\bar{g}(\Psi)$ is the sample mean of $g_{i\tau}(\Psi)$ and W is the identity matrix.

The vector of instruments we use is:

$$\mathbf{h}_{i\tau} = \begin{pmatrix} 1 \\ \{\mathbf{z}_{i\tau}\}_{\mathcal{Z}_1} \\ \log p^e_{r(i)\tau} - \log p^n_\tau \\ \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right) \times \{\mathbf{z}_{i\tau}\}_{\mathcal{Z}_2} \\ \{\widetilde{\log \mathrm{inc}}_{i\tau}\}_r \\ \widetilde{\log \mathrm{inc}}_{i\tau} \times \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right) \\ \mathbb{d}_{i\tau} \times \{\mathbf{z}_{i\tau}\}_{\mathcal{Z}_3} \end{pmatrix}$$

[51]More specifically, the Jacobian on the system of equations (4.2) and (4.3) is non-zero as long as $\left(\log p^e_{r(i)\tau} - \log p^n_\tau\right)\left((C_1 + \sum_{l \in \mathcal{Z}_2} C_{1l}) + \sum_{k>1} C_k k y_{i\tau}^{k-1} + \frac{1}{2} D \left(\log p^e_{r(i)\tau} - \log p^n_\tau\right)\right) \neq -1$.

Note this excludes the household's total budget ($x_{i\tau}$) and includes functions of the log of their total monthly income ($\text{inc}_{i\tau}$) deflated with a Stone prices index: $\widetilde{\log \text{inc}}_{i\tau} \equiv \log \text{inc}_{i\tau} - \left(\bar{w} \log p^e_{r(i)\tau} + (1-\bar{w}) \log p^n_{\tau}\right)$.

Starting values. To obtain starting values for the GMM estimation procedure, we use the iterated least squares procedure of Blundell and Robin (1999). This entails fixing an initial guess of $y_{i\tau}$, estimating the energy budget share demand (equation (4.2)) using linear methods (IV), updating $y_{i\tau}$ (equation (4.3)) and continuing the procedure until convergence. At each iteration, when estimating the budget share equation, we instrument $\log y_{i\tau}$ using a generated instrument defined by equation (4.3), but with $\text{inc}_{i\tau}$ in place of $x_{i\tau}$.[52] The procedure yields estimates that, under the orthogonality condition (C.6), are consistent but asymptotically inefficient.

Standard errors. We compute standard errors allowing for arbitrary dependence between observations within i. Specifically, let T_i denote the total number of periods household i is observed in the data, denote the average moment condition, evaluated at estimated parameters, for household i by $g_i(\widehat{\Upsilon}) = \frac{1}{T_i}\sum_{\tau=1}^{T_i} g_{i\tau}(\widehat{\Upsilon})$, and let $V = \frac{1}{N}\sum_i g_i(\widehat{\Upsilon})g_i(\widehat{\Upsilon})'$ denote the variance of the average moment conditions. Let $D = \overline{\bigtriangledown g_{i\tau}}$ be the sample average of the gradient of the moments evaluated at estimated parameters. The asymptotic variance matrix of the estimates is then:

$$\text{Var}(\widehat{\Upsilon}) = (D'WD)^{-1}\left(D'WVWD\right)(D'WD)^{-1}.$$

When we report statistics based on the model estimates (for instance, total efficiency costs of policy), we report Monte Carlo confidence intervals. To compute these, we take 100 draws from the asymptotic variance matrix of the estimates, calculate the statistic of interest, and use the resulting distribution to construct confidence intervals.

C.3 Estimates

In Figure 4.2(a) of the main paper we scatter the elasticities shown in Figure 3.3 with their counterparts computed based on the demand model estimates. Figure C.1 plots how the model-based elasticities vary across the 25 groups.

[52]In practice, when iterating between equations (4.2) and (4.3) we do note update the parameters values used to construct the generated instrument. Rather, we initially iterate between the two equations using $\log \text{inc}_{i\tau} - (\bar{w}\log p^e_{r(i)\tau} + (1-\bar{w})\log p^n_{\tau})$, where $\bar{w}$ is the sample average budget share, as the instrument. Using the converged parameters we construct the generated instrument. We then undertake the iterative estimation procedure a second time with this new instrument.

Figure C.1: *Heterogeneity in elasticities, model-based estimates*

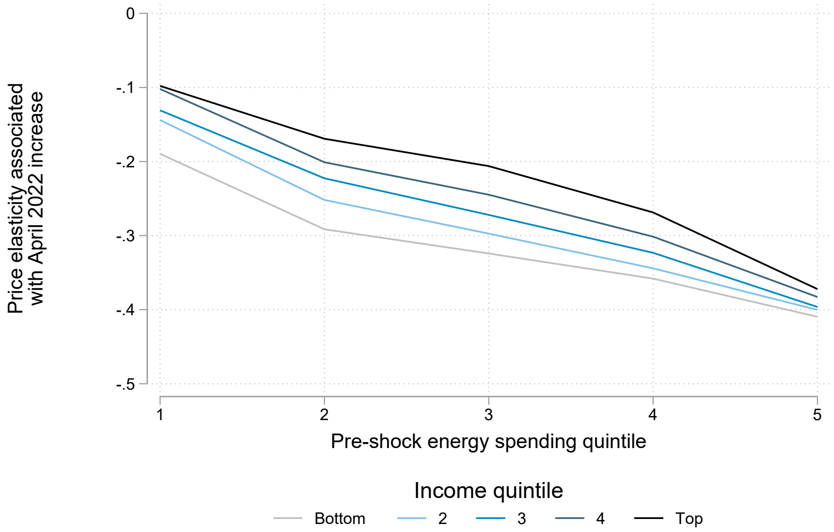

Notes: Figure shows the estimates of the average price elasticity over the April 2022 price increase, for each income quintile and pre-shock energy consumption quintile based on our energy demand model.

In Figure C.2 we summarise heterogeneity in Engel curves. Each line show the relationship between the energy budget share (ω) and the log of total net-of-fixed-fee expenditure ($\log x$). We draw a separate line for households belonging to each decile of the pre-shock energy spending distribution. We hold prices and temperature controls at their mean level, and controls for prepay and decile of the pre-shock energy spending share distribution at their pre-shock energy spending decile specific mean. All means are computed over April-September 2022. The graphs shows significant heterogeneity in Engel curves. The extent of this heterogeneity is similar to that found in Lewbel and Pendakur (2017), who estimate energy demand in Canada using cross-sectional household survey data, and model preference heterogeneity via random Barten scales.

Figure C.2: *Heterogeneity in Engel curves*

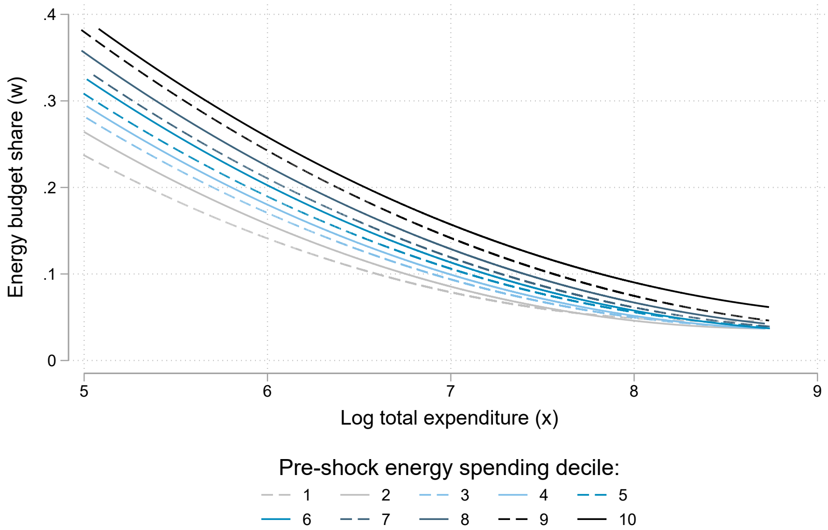

Notes: Figure shows Engel curves estimates.

C.4 Computing Household-Level Welfare

Optimal choice. When a household's choice is not subject to a flypaper effect (meaning $\mathbb{d} = 0$), we can obtain their attained level of utility when faced with the budget set $(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau})$ by solving equations (4.2) and equation (4.3) to obtain $y_{i\tau}$. Attained utility is then given by $\exp(y_{i\tau})$.

Suboptimal choice. When the household's choice is subject to a flypaper effect, we solve for the pivot in their budget constraint (parameterised by ϕ) that rationalises the choice as optimal (see Proposition 1). In practice we do this as follows. Let $\mathbb{e}^1_{i\tau} = \mathbb{e}(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, 1, \theta_{i\tau})$ denote the household's quantity choice, (i.e., $\mathbb{e}^1_{i\tau} = \omega(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, 1, \theta_{i\tau})\frac{x_{i\tau}}{p^e_{r(i)\tau}}$. We solve for the budget rotation $\phi_{i\tau}$ using the iterative algorithm:

$$\phi^{(l)}_{i\tau} = \phi^{(l-1)}_{i\tau} + \log \mathbb{e}^1_{i\tau} - \log \mathbb{e}(p^e_{r(i)\tau}(1 - \phi^{(i-1)}_{i\tau}), p^n_\tau, x_{i\tau} - \phi^{(l-1)}_{i\tau} p^e_{r(i)\tau}\mathbb{e}^1_{i\tau}, 0, \theta_{i\tau}),$$

until $||\phi^{(i)}_{i\tau} - \phi^{(i-1)}_{i\tau}|| < 10e^{-3}$. Once we have obtained $\phi_{i\tau}$ we solve equation (4.2) and equation (4.3), with $\mathbb{d}$ set to 0, at the hypothetical budget set $(p^e_{r(i)\tau}(1 - \phi_{i\tau}), p^n_\tau, x_{i\tau} - \phi_{i\tau}\mathbb{e}^1_{i\tau})$ to obtain $y_{i\tau}$ and hence attained utility as $\exp(y_{i\tau})$.

Money-metric utility. The preceding two paragraphs outline how we obtain the function that maps choices into the level of utility they attain

$$\mathbb{V}(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, \mathbb{d}, \theta_{i\tau}) = \begin{cases} V(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, \theta_{i\tau}) & \text{if} \quad \mathbb{d} = 0 \\ V(p^e_{r(i)\tau}(1-\phi_{i\tau}), p^n_\tau, x_{i\tau} - \phi_{i\tau} p^e_{r(i)\tau} \mathbb{e}^1_{i\tau}; \theta_{i\tau}) & \qquad\;\; \mathbb{d} = 1. \end{cases}$$

We use the expenditure function (specified in equation (C.1)) to cardinalise utility in money-metric terms. Let $(p^e_{r(i)0}, p^n_0)$ denote pre-shock prices. We use the money-metric cardinalisation of utility given by:

$$\mathbb{V}^{MM}(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, \mathbb{d}_{i\tau}, \theta_{i\tau}) = \chi(p^e_{r(i)0}, p^n_0, \mathbb{V}(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, \mathbb{d}_{i\tau}, \theta_{i\tau}), \theta_{i\tau})$$

Note that if the household faces pre-shock marginal prices $(p^e_{r(i)0}, p^n_0)$ and fixed fee $f^e_{r(i)0}$ (and, since there is no transfer, $\mathbb{d} = 0$) then money-metric utility is:

$$\begin{aligned} \mathbb{V}^{MM}(p^e_{r(i)0}, p^n_0, \tilde{x}_{i\tau} - f^e_{r(i)0}, 0, \theta_{i\tau}) &= \chi(p^e_{r(i)0}, p^n_0, \mathbb{V}(p^e_{r(i)0}, p^n_0, \tilde{x}_{i\tau} - f^e_{r(i)0}, 0, \theta_{i\tau}), \theta_{i\tau}) \\ &= \chi(p^e_{r(i)0}, p^n_0, V(p^e_{r(i)0}, p^n_{r(i)0}, \tilde{x}_{i\tau} - f^e_{r(i)0}, \theta_{i\tau}), \theta_{i\tau}) \\ &= \tilde{x}_{i\tau} - f^e_{r(i)0} \end{aligned}$$

The money-metric utility function has the interpretation of the minimum expenditure the household requires at pre-shock prices $(p^e_{r(i)0}, p^n_0)$ to reach the level of utility attained at the post-shock budget set (including any flypaper effect induced by the transfer). The policy menus we consider, $\mathbb{P}$, influence the energy price through a subsidy s, the consumer's budget by a transfer t, and may entail a flypaper effect (for prepay households) if the transfer is labelled, captured by $L \in \{0, 1\}$. For notational parsimony, in the paper we denote money-metric utility $\mathbb{V}^{MM}_{i\tau}(\mathbb{P}) \equiv \mathbb{V}^{MM}(p^e_{r(i)\tau}, p^n_\tau, x_{i\tau}, \mathbb{d}_{i\tau}, \theta_{i\tau})$. The reduction in money-metric utility for household i in period τ due to the shock is:

$$\mathcal{L}_{i\tau}(\mathbb{P}) \equiv (\tilde{x}_{i\tau} - f^e_{r(i)0}) - \mathbb{V}^{MM}_{i\tau}(\mathbb{P}),$$

and the household's money-metric utility loss over the whole shock period, $\tau \in \{\underline{\tau}, \overline{\tau}\}$, is $\mathcal{L}_i(\mathbb{P}) = \sum_{\tau=\underline{\tau}}^{\overline{\tau}} \mathcal{L}_{i\tau}(\mathbb{P})$. In some places we scale this loss by household income (average household monthly over April 2021-March 2022, scaled by the number of months in the shock period), to obtain $l^y_i = \frac{L_i}{Y_i}$.

C.5 Computing Efficiency Costs and Social Welfare

Efficiency costs. The observed government policy response ($\mathbb{P}^O = (s^O, t^O, L^O)$) expended public funds inclusive of carbon emissions given by $\bar{R}$ defined by equation (5.1), which we repeat here:

$$\bar{R} \equiv s^O \sum_{i=1}^{N} x_i^e(\mathbb{P}^O) + N \times 6t^O + \alpha \sum_{i=1}^{N} \Big(e_i(\mathbb{P}^O) - e_i(\emptyset) \Big),$$

where $e_i(\mathbb{P}^O) = \sum_{\tau=\underline{\tau}}^{\overline{\tau}} e_{i\tau}(\mathbb{P}^O)$ is household i's total energy consumption over the crisis, $e_i(\emptyset)$ is their consumption in the absence of any government policy response, $x_i^e(\mathbb{P}^O) = \sum_{\tau=\underline{\tau}}^{\overline{\tau}} P^e_{r(i)\tau} e_{i\tau}(\mathbb{P}^O)$ is their subsidy-exclusive energy spending over the crisis, and α converts energy consumption into the social cost of the associated carbon emissions. We can re-write equation (5.1):

$$\bar{R} = \sum_i \sum_{\tau=\underline{\tau}}^{\overline{\tau}} (R_{i\tau} + \bar{\alpha}),$$

where $R_{i\tau} = s^O x_i^e(\mathbb{P}^O) + t^O$ is the public funds provided to household i in period τ and $\bar{\alpha} = \frac{\alpha}{6N} \sum_{i=1}^{N} \big(e_i(\mathbb{P}^O) - e_i(\emptyset) \big)$ is the average monthly social costs of carbon emissions.

To compute the efficiency cost associated with the implemented policy and to decompose it into its sources we compute the following for each (i, τ):

1. $\mathbb{V}^{MM}_{i\tau}(s^O, t^O, L^O)$: money-metric utility at observed policy
2. $\mathbb{V}^{MM}_{i\tau}(0, R_{i\tau} + \bar{\alpha}, 0)$: money-metric utility under an individual specific transfer equal to $R_{i\tau} + \bar{\alpha}$, the monetary value of funds given to the household and the average value of the carbon externality.
3. $\mathbb{V}^{MM}_{i\tau}(0, R_{i\tau}, 0)$: money-metric utility under a transfer equal $R_{i\tau}$, the monetary value of funds given to the household
4. $\mathbb{V}^{MM}_{i\tau}(s^O, t^O, 0)$: money-metric utility under the observed subsidy rate and transfer, but under no transfer labelling
5. $\mathbb{V}^{MM}_{i\tau}(s^O, \tilde{t}, 0)$ where $\tilde{t}$ is such that $s^O \sum_{i=1}^{N} x_i^e(s^O, \tilde{t}, 0) + N \times 6\tilde{t} = \bar{R}$: money-metric-utility under the observed subsidy rate and a transfer (with no transfer labelling) set to expend the same amount of government revenue as observed policy.

We measure aggregate efficiency costs of the implemented policy as

$$\text{efficiency cost} = \sum_i \sum_{\tau=\underline{\tau}}^{\overline{\tau}} \Big(\mathbb{V}^{MM}_{i\tau}(0, R_{i\tau} + \bar{\alpha}, 0) - \mathbb{V}^{MM}_{i\tau}(s^O, t^O, L^O) \Big)$$

and decompose it according to:

$$\begin{aligned}
\text{efficiency cost} = \quad & \textstyle\sum_i \sum_{\tau=\underline{\tau}}^{\overline{\tau}} \big(\mathbb{V}^{MM}_{i\tau}(0, R_{i\tau} + \bar{\alpha}, 0) - \mathbb{V}^{MM}_{i\tau}(0, R_{i\tau}, 0) \big) && \text{(carbon emissions)} \\
& \textstyle\sum_i \sum_{\tau=\underline{\tau}}^{\overline{\tau}} \big(\mathbb{V}^{MM}_{i\tau}(0, R_{i\tau}, 0) - \mathbb{V}^{MM}_{i\tau}(s^O, \tilde{t}, 0) \big) + && \text{(price signal)} \\
& \textstyle\sum_i \sum_{\tau=\underline{\tau}}^{\overline{\tau}} \big(\mathbb{V}^{MM}_{i\tau}(s^O, \tilde{t}, 0) - \mathbb{V}^{MM}_{i\tau}(s^O, t^O, 0) \big) + && \text{(labelling: fiscal spillover)} \\
& \textstyle\sum_i \sum_{\tau=\underline{\tau}}^{\overline{\tau}} \big(\mathbb{V}^{MM}_{i\tau}(s^O, t^O, 0) + \mathbb{V}^{MM}_{i\tau}(s^O, t^O, L^O) \big) && \text{(labelling: choice distortion).}
\end{aligned}$$

Counterfactual policies. We consider a series of counterfactual policies that all expend $\bar{R}$ public resources (inclusive of the social costs of carbon emissions). These counterfactual policies are:

1. All values with $s \geq 0$ and $t \geq 0$ of $(s, t, L = 1)$. Note $(s = 0.39, t = 66, L = 1) = (s^O, t^O, L = 1)$ corresponds to the policy implemented in practice

2. All values with $s \geq 0$ and $t \geq 0$ of $(s, t, L = 0)$.

3. All values with $s \geq 0$ and $t \geq 0$ of $(s, t/Y_i, L = 0)$, where Y_i is households i's average monthly income over April 2021-March 2022.

4. All values with $s \geq 0$ and $t \geq 0$ of $(s, t \times E_i, L = 0)$. E_i is a measure of past energy use. We obtain this by computing the Marshallian budget share energy demand for each household i in each period τ in the 6 month period October 2021-March 2022 at observed variables. We then convert this to E_i by multiplying the budget share demand by the household's total period τ budget, dividing it by the period-marginal energy price, and then averaging over the 6 months for each household.

5. All values with $s \geq 0$ and $t \geq 0$ of $(s, t \times E_i/Y_i, L = 0)$, where Y_i and E_i are computed as in the preceding two cases.

We use the social loss function $\mathcal{W}(\mathbb{P}) = \frac{1}{N}\sum_{i=1}^{N} \frac{1}{\psi}\left(\exp(\psi \times l_i^y(\mathbb{P})) - 1\right); \quad \psi > 0$. We fix ψ at the value such that:

$$(s^O, t^O) = \arg\min_{(s,t)} \mathcal{W}(s, t, L = 1),$$

that is at the value that rationalises the observed choice of subsidy and transfer from the menu $(s, t, L = 1)$.

Define the *average social welfare weight* of a households with income Y and equivalent variation losses L as the relative weight the planner places on each £ of loss they suffer. This is given by their contribution to the social loss function scaled by their monetary loss:

$$w(L, Y) = \frac{1}{\psi \mathcal{L}}\left(\exp\left(\psi \frac{\mathcal{L}}{Y}\right) - 1\right)$$

In Figure C.3 we plot combinations of $(Y, \mathcal{L})$ that are assigned the same w (i.e., the level sets of the welfare weights). Panel (a) shows these for our calibrated value of ψ. Panel (b) shows them for the limiting case when $\psi \to 0$; in this case the social loss function reflects only vertical equity concerns, and minimising it is equivalent to maximising money-metric utility scaled by inverse income. The comparison between the level sets in the panels shows how the social loss function's convexity alters the standard vertical welfare weights.

Figure C.3: *Level set of average welfare weights*

(a) Calibrated ψ (b) ψ=0

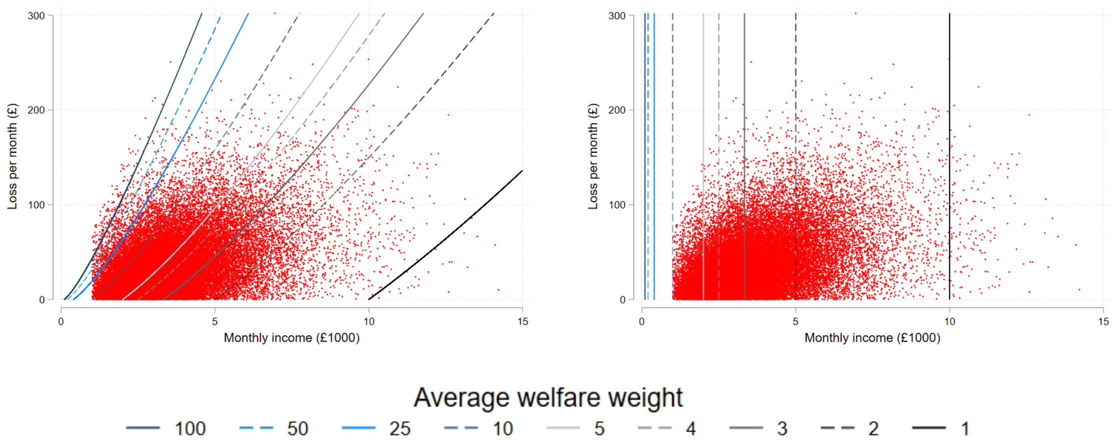

Notes: Figure shows combinations of income and monetary welfare loss that lead to the same average welfare weight. Panel (a) is based on our calibrated value of ψ. Panel (b) shows the case where $\psi = 0$, meaning the planner is indifferent to loss inequality conditional on income. The red dots scatter the household-level joint distribution of monetary losses and income under the observed policy response. We restrict the graph to positive losses; under the convex social loss function, households with negative losses (gains) are assigned average welfare weights that are very close to zero.

For each policy $\mathbb{P}$ we measure aggregate social losses as:

$$\xi^{\mathbb{P}} = \frac{1}{\psi} \log(\psi \mathcal{W}(\mathbb{P}) + 1),$$

that is the constant level of proportional loss that would result in the same level of the social loss function as the distribution of losses under policy $\mathbb{P}$.

We decompose $\xi^{\mathbb{P}}$ into uncompensated losses, efficiency costs, and cost from failing to achieve equal proportional losses (targeting costs). The decomposition is:

$$\xi^{\mathbb{P}} = \underbrace{\xi^{LS}}_{\text{uncompensated losses}} + \underbrace{\left(\frac{\bar{\mathcal{L}}^{\mathbb{P}}}{\bar{Y}} - \xi^{LS}\right)}_{\text{efficiency costs}} + \underbrace{\left(\xi^{\mathbb{P}} - \frac{\bar{\mathcal{L}}^{\mathbb{P}}}{\bar{Y}}\right)}_{\text{targeting costs}},$$

where ξ^{LS} are losses under the individualised lump-sum subsidy scheme that expends $\bar{R}$ public resources and equates proportional losses, $\bar{\mathcal{L}}^{\mathbb{P}}$ is average monetary losses under the policy $\mathbb{P}$ and $\bar{Y}$ is average income.

The dual of the planner's problem. The results reported in the paper (Figure 5.2) compare the value of social losses attained under alternative policies holding fixed their public resource cost (at $\bar{R}$, the costs under observed policy). In other words, they entail minimising equation (5.2) subject to the constraint (5.1). An alternative approach to comparing alternative policies is to minimise costs subject to social losses equalling their level under observed policy. This alternative framing results in similar conclusions. To illustrate this point, in

Table C.1, for each counterfactual policy menu, under both the social loss minimising subsidy level, and when the subsidy is set to zero, we solve for the level of transfers the result in the same social loss level us under observed policy. Columns (2)-(5) report the *reduction* in public resource costs under each policy (negative values entail higher costs than under observed policy). So, for instance, the policy of a subsidy and transfer tied to past energy consumption scaled by income could have attained the same social loss level as observed policy at 17.5% (or £6bn) lower public cost. The social ranking of the policies in Table C.1 are the same as those implied by Figure 5.2.

Table C.1: *Public resource costs of reaching* $\mathcal{W}(\mathbb{P}^O)$ *under alternative policies*

	(1)	(2)	(3)	(4)	(5)
		% public resource cost savings under:			
	$(s^O, t^O, L=1)$	$(s, t, L=0)$	$(s, t/Y_i, L=0)$	$(s, t \times E_i, L=0)$	$(s, t \times E_i/Y_i, L=0)$
s=s*	£34.43bn	+1.5%	+8.4%	+7.9%	+17.5%
s= 0		-12.3%	-20.1%	+4.9%	+13.0%

Notes: In row (1) the subsidy is set at the level that minimises social losses subject to using $\bar{R}$ *public resources. In row (2) the subsidy is set to 0. In each case we solve for the transfer value that equates social losses to those realised under observed policy. Column (1) reports the resource costs of observed policy* ($\bar{R}$)*. Columns (2)-(5) report the resource cost savings under each alternative policy.*

Robustness of results to social preference specification. In Figure C.4(a) we show how the optimal subsidy rate, under each policy menu, varies with the convexity parameter, ψ, in the social loss function. The vertical line denotes our baseline value. Under each policy menu, higher convexity leads to a higher optimal subsidy rate. In panel (b) we show how social losses *under optimal policy* vary with ψ. To make results comparable across different social preferences, we express social losses as a % of those under incurred under the menu $(s, t, L = 1)$. For almost all values of ψ, optimal policy under the menu $(s, t, L = 1)$ is dominated by that under $(s, t, L = 0)$, which, in turn is dominated by that under $(s, \frac{t}{Y}, L = 0)$ and $(s, t \times E, L = 0)$, with $(s, t \times \frac{E}{Y}, L = 0)$ performing best. The exception is when the convexity in social losses is very low, in which case the planner places little weight on reducing inequality in losses conditional on income, and the menu $(s, t \times E, L = 0)$ becomes the poorest performing.

In Figure C.4(c) we show how social losses based on pure transfer schemes (i.e., with $s = 0$) vary with ψ. For the policy menus, $(s = 0, t, L = 0)$, $(s = 0, \frac{t}{Y}, L = 0)$, $(s = 0, t \times E, L = 0)$, $(s = 0, t \times \frac{E}{Y}, L = 0)$, we plot social losses as a % of those under incurred under the menu $(s, t, L = 1)$ – note this entails the optimal rather than zero subsidy. It shows that a pure transfer scheme based on past energy use over income, $(s = 0, t \times \frac{E}{Y}, L = 0)$, always outperforms the policy menu $(s, t, L = 1)$. Conversely, except when the convexity in social losses is very low, a pure transfer scheme based on income performs worse than the menu $(s, t, L = 1)$.

Figure C.4: *Variation in optimal policy with ψ*

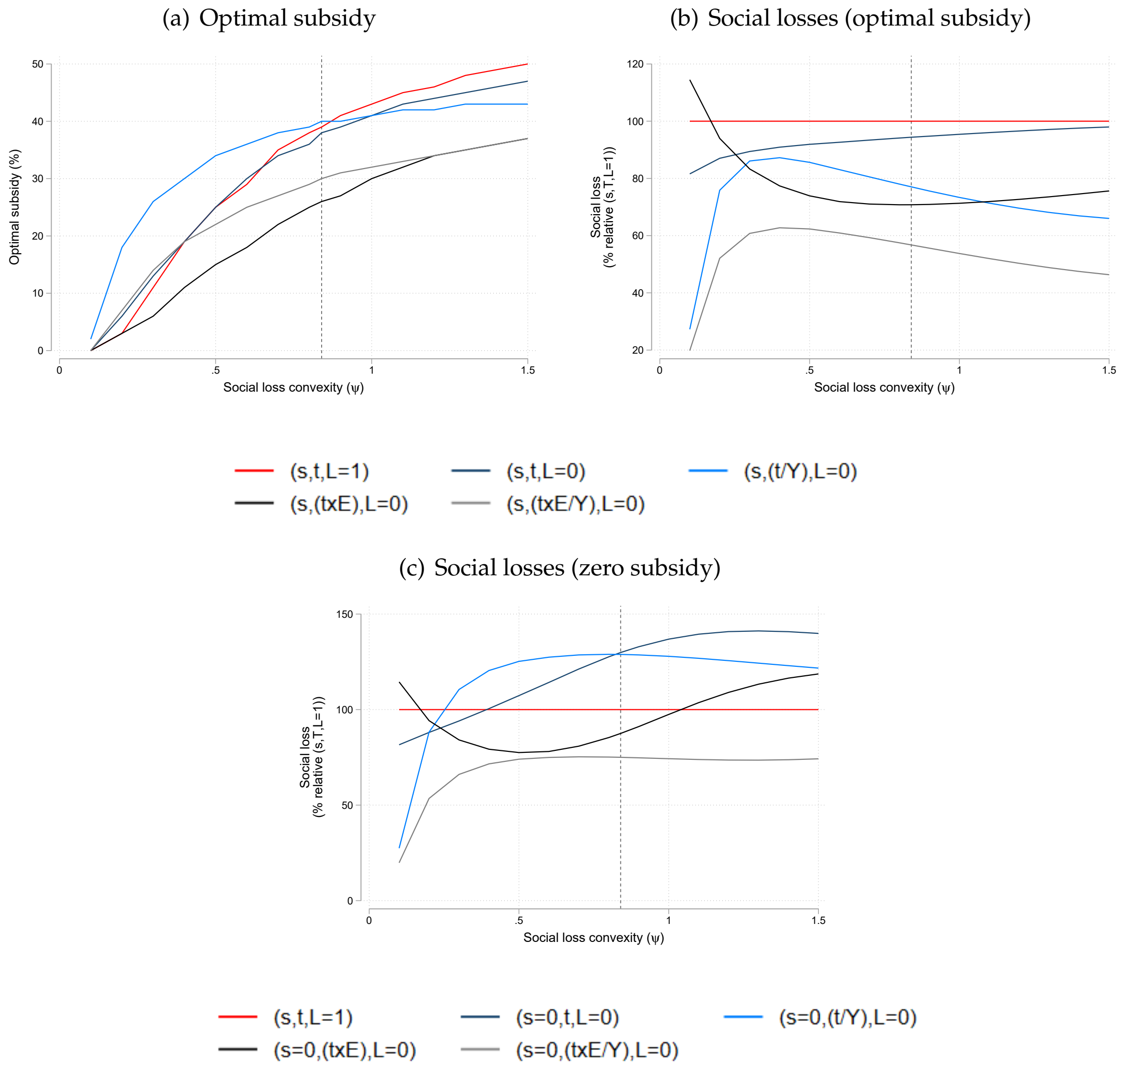

Notes: Panel (a) shows the optimal subsidy rate under each policy menu for different values of ψ. Panel (b) shows how the associated welfare losses vary, expressed as a % of those under the menu $(s, t, L = 1)$. Panel (c) shows how welfare losses under the policy menus $(s, t, L = 0)$, $(s, (t/Y), L = 0)$, $(s, (t \times E), L = 0)$, and $(s, (t \times E/Y), L = 0)$ at $s = 0$ vary with ψ, expressed as a % of those under the menu $(s, t, L = 1)$ at the optimal subsidy. The vertical line indicates our baseline value of ψ.