

Peter Levell
Martin O'Connell
Kate Smith

Working paper

25/55

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

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

Peter Levell, Martin O'Connell and Kate Smith*

November 2025

Abstract

Aggregate price shocks generate unequal welfare losses both across and within income groups. Policymakers face a trade-off: price subsidies target those most affected but create inefficiencies, while transfers are less distortionary but harder to target. We develop and implement a framework to quantify this trade-off using rich panel data on households' energy spending and incomes, alongside price and policy variation from the 2022–23 European Energy Crisis. Absent policy intervention, average household welfare losses would have equalled 6% of income, with some households facing much larger losses. The combination of an energy-price subsidy and universal transfers reduced both the mean and dispersion of losses, but incurred efficiency costs equal to 12% of the total relief package revenue costs. Optimal policy entails a strictly positive price subsidy; its level is lower when transfers can be targeted using income and past energy usage.

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 capacity to smooth temporary shocks often prompt substantial fiscal interventions. Recent examples include price subsidies and transfers implemented by more than 25 countries in response to the European Energy Crisis, with fiscal costs approaching 2% of European GDP in 2022 (European Central Bank, 2024). Yet despite the recurrence of energy shocks and the widespread use of such measures, evidence on the distribution of households' responses and its implications for optimal policy design remains limited.

This paper provides new evidence on the distribution of welfare losses from a large energy price shock and develops a framework for evaluating policy responses, drawing on price and policy variation during the 2022–23 European Energy Crisis. A central element of our analysis is the challenge of efficiently targeting support to the most affected households. We provide a new quantification of the incidence of large energy price shocks—incorporating heterogeneous behavioural responses—and show that incidence is highly dispersed and only weakly correlated with income. Consequently, conventional redistribution via the tax-and-transfer system is ill-suited to compensating households. We show that optimal policy features a strictly positive energy-price subsidy, despite its sizeable efficiency costs and even when the government can deploy a wide range of transfer schemes.

We study the UK during the European Energy Crisis. Over this period, retail prices for residential gas and electricity were constrained by a binding regulatory cap that adjusted in discrete steps based on international wholesale costs. This generates variation that we use to estimate households' demand responses to energy prices. During the crisis, the UK government also implemented an energy-price subsidy alongside transfers intended to help with both energy bills and broader cost-of-living pressures; we use this policy variation to estimate how transfer design affects energy consumption. Our analysis uses bank account and credit card data covering income and spending for a panel of a quarter of a million UK households from 2019 to 2023. We document households' exposure to energy price shocks and their responses to price increases and accompanying policy interventions. We then estimate a flexible model of energy demand and embed it in a social-welfare framework to quantify welfare effects and evaluate alternative policies.

A key driver of households' exposure to energy price shocks is their level of energy expenditure. Our data records household-level energy spending and income at high frequency,

¹See Klick and Stockburger (2024), 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.

overcoming a common limitation of energy-usage datasets that lack income information (Borenstein, 2012). We show that, although higher-income households spend more on energy in levels (but less as a share of total spending), income explains only 7% of the cross-sectional variation in energy spending. The panel dimension of our data also lets us quantify persistence: we show that lagged household energy spending explains about 55% of the variation in current spending. Thus, while historic usage can be a useful criterion for targeting energy support, households' energy needs still fluctuate substantially over time.

Households' willingness to substitute away from energy use when prices rise mitigates the welfare impact of shocks but increases the efficiency costs of price subsidies.² We estimate demand responsiveness by exploiting large, periodic jumps in regulated energy prices generated by cap adjustments, combining these discrete changes with high-frequency spending data and granular controls for seasonality and weather. In April 2022, a cap adjustment raised the real residential energy price by 45%, prompting a near-immediate average reduction of 15% in household energy consumption—an implied own-price elasticity of -0.33. Households with higher baseline (pre-crisis) energy spending reduce consumption more in response to price increases, while higher-income households are slightly less price-responsive than lower-income households. These findings contribute new evidence on both average and heterogeneous demand responses—by income and prior usage—in residential energy demand to a large, salient price change.³

Another way to support households is through transfers; these avoid the efficiency costs from distorting price signals but can be difficult to target towards those most affected by a price shock. Transfers can also introduce other inefficiencies, such as labelling-induced over-consumption. We exploit variation from the UK's energy-support transfers—administered by energy suppliers—together with “cost-of-living” payments paid directly to low-income households to estimate the marginal propensity to consume energy (MPCE) out of each transfer type. For households with prepayment meters, the average MPCE out of energy-support transfers is 34%, substantially higher than the 3% MPCE from cost-of-living payments. This finding contributes to evidence on the “flypaper effect” (Hines and Thaler, 1995), whereby agents do not treat money as fungible.⁴ If this reflects distortions to household decision-

²Even modest substitution elasticities can materially dampen welfare losses; see Moll et al. (2023) on the macroeconomic effects of the cut-off from Russian gas in Germany.

³The literature reports a wide range of elasticities (see 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 (Dergiades and Tsoufidis, 2008) or non-linear contracts (Reiss and White, 2005)), and limited household responses to marginal price incentives embedded in complex non-linear contracts (e.g., Ito, 2014; Shaffer, 2020). Our evidence of rapid consumption responses aligns with findings from the 2000–01 California energy crisis (Reiss and White, 2008).

⁴Examples include the allocation of the “Winter Fuel Payment” to energy spending (Beatty et al., 2014), education-support transfers (Benhassine et al., 2015), and households' spending response to SNAP vouchers (Hastings and Shapiro, 2018).

making—such as mental accounting (Shefrin and Thaler, 1988)—then observed choices may diverge from underlying welfare.

Informed by these findings, we estimate a model of household energy demand to evaluate the welfare effects of the crisis—fully accounting for households’ behavioural responses—and to compare counterfactual policy designs. We employ a flexible empirical specification based on the EASI demand system (Lewbel and Pendakur, 2009), exploiting our longitudinal data to capture rich preference heterogeneity. To assess welfare when households may make privately suboptimal choices, we extend the revealed-preference approach of Bernheim and Rangel (2009) (applied by Chetty et al. (2009) in their study of tax salience), to a rich empirical demand model with non-marginal price and policy changes. The model performs well out-of-sample, predicting both the distribution of energy demands and its variation with income.

We use the model to provide the first comprehensive analysis of the incidence of the 2022–2023 European Energy Crisis that accounts for households’ behavioural adjustments.⁵ Absent government intervention, the crisis would have generated aggregate equivalent-variation losses of £35.2 billion over six months (£207 per household-month); this corresponds to 6% of after-tax income on average, rising to 11% at the 95th percentile. Lower-income households were disproportionately affected: at the 10th income percentile, average proportional losses are 5 percentage points higher than at the 90th percentile. Dispersion within income groups is even larger: among the poorest 10% of households, the 90–10 loss gap is 11 percentage points. First-order welfare approximations based on pre-shock spending—which hold usage fixed and thus ignore behavioural response—can perform poorly for large shocks; in our setting, ignoring households’ behavioural responses overstates the welfare losses by 59%.

The UK household relief package reduced aggregate losses by £27.8 billion and compressed the upper tail, lowering the 95th percentile of losses from 11% to 3% of income. This came with efficiency costs of £3.8 billion. Efficiency losses arose from two sources. First, subsidised prices induced substitution toward energy use, generating excess consumption and higher social costs from carbon emissions. Second, the flypaper effect imposed a direct utility loss by steering transfers toward energy use rather than households’ preferred allocations. These two mechanisms interact: flypaper-induced overconsumption of subsidised energy creates a fiscal spillover more than five times larger than the direct distortion of choices. Implementing transfers in a way that eliminates the flypaper effect would have reduced efficiency costs by about 20%.

⁵A handful of papers study its macroeconomic effects (e.g., Auclert et al., 2023; Pieroni, 2023), or distributional effects using only pre-crisis spending data (e.g., Bachmann et al., 2022; Fetzner et al., 2023). More broadly, much of the literature on energy price shocks focuses on macroeconomic and average effects (Kilian, 2008).

This raises the question: could the relief package have been designed to target assistance more effectively to those most affected? To answer this, we develop a social-welfare framework that characterises the trade-off between targeting and efficiency. We assume that the policymaker’s objective is to bring households as close as possible to their pre-shock circumstances. Our specification of social preferences accommodates both vertical equity—a given financial loss is more burdensome for lower-income households—and aversion to within-income loss inequality—conditional on income, hardship increases non-linearly with loss size. This formulation nests a standard social-welfare function and extends it by incorporating a proportional version of the equal-sacrifice principle (see Fleurbaey and Maniquet, 2018). The ideal policy would allocate public funds as personalised lump-sum transfers to equalise losses as a share of income. However, *feasible* policies must balance targeting gains against efficiency costs.

We compare policy menus that pair a subsidy with transfers that are either universal or proportional to income, prior energy usage, or both. This spans the main approaches European governments adopted during the crisis.⁶ We show that income-based transfers alone would have led to larger welfare losses than the UK’s implemented package because many households would still have experienced large proportional losses. Setting transfers as a function of *both* income and prior energy use better targets those most exposed to large proportional losses. Nonetheless, fluctuations in energy demand over time are substantial enough to justify a strictly positive subsidy even in this case. The optimal policy package—combining a subsidy with transfers based on income and past energy usage—closes over 60% of the gap in social losses between first-best personalised-transfer benchmark and the implemented UK response. Sensitivity analysis shows that, provided the policymaker exhibits some aversion to large losses, our qualitative findings remain: optimal policy includes a positive subsidy, and income-based transfers alone perform poorly.

Much of the normative public-finance literature focuses on the trade-off between efficiency and *vertical* equity when redistributing toward lower-income households (Piketty and Saez, 2013), including in the analysis of energy policy (e.g., Bento et al., 2009; Borenstein, 2012). In that context, Hahn and Metcalfe (2021) show that using energy subsidies to support low-income households can be welfare-reducing. A growing literature, however, highlights how heterogeneous exposure to idiosyncratic shocks (e.g., Lieber and Lockwood, 2019) or localised (place-specific) risk (e.g., Gadenne et al., 2024) can justify in-kind instruments to better target the most affected households. We contribute to this literature by showing that

⁶Basing transfers on prior use is feasible because energy suppliers retain consumption records. For example, Germany and Austria tied support to past consumption during the European Energy Crisis; Brazil did so in an earlier crisis (Costa and Gerard, 2021). The role of information in shaping policy is explored by Fetzer et al. (2024), who argue that incorporating additional data—such as property characteristics or prior energy consumption—could have reduced losses for most UK households.

aggregate shocks with heterogeneous incidence within income groups generate a targeting–efficiency trade-off that can justify price subsidies. Our quantification focuses on energy—a recurring source of aggregate shocks—but because similar patterns of heterogeneous exposure arise under other aggregate shocks (e.g., trade shocks; Borusyak and Jaravel (2024)), these qualitative policy trade-offs apply more broadly.

The rest of this paper is structured as follows. The next section summarises the institutional setting and data; Section 3 describes households’ exposure to energy price shocks. Section 4 presents our results on consumption responses to price increases and evidence of a flypaper effect from transfers. Section 5 outlines our model of household energy demand. Section 6 presents the welfare effects of a large energy price shock and evaluates alternative policy responses. A final section concludes.

2 Institutional Setting and Data

2.1 UK Residential Energy Market

Residential energy bills in the UK typically comprise (i) unit charges for electricity and gas (per kWh), which generally do not vary with total usage or time of day; and (ii) fixed standing charges that are independent of usage and constitute a relatively small share of households’ bills.⁷ A distinctive feature of the UK (excluding Northern Ireland) retail energy market is a price cap, administered by the energy regulator Ofgem, which sets maximum unit and standing charges for residential electricity and gas.⁸ The cap aims to limit excess profits and protect non-switching households from overpaying. Ofgem sets the cap level based on estimates of supplier costs and announces any changes roughly one month before they take effect. Until October 2022 the cap was updated every six months; since then it has been updated quarterly in response to greater wholesale price volatility.

Households pay for energy in different ways. We use the UK’s national budget survey, the Living Costs and Food Survey (LCFS; Office for National Statistics, 2024) to calculate the share of household electricity and gas spending by payment method. 79% of spending is paid for by direct debit with payments automatically collected each month. 14% is paid for using prepayment meters, with credit that households top up in advance (online or in shops) before the energy can be consumed. The remaining 7% is paid for using standard credit, settling bills after receiving an invoice for actual use.

⁷Increasing-block pricing—where the marginal price increases with usage—is not a feature of the UK market. The majority of UK households (86% in 2012; Department for Energy and Climate Change, 2013) face an electricity unit charge that is fixed throughout the day. Over 2021–2023, the standing charges account for about 13% of the average household bill.

⁸Northern Ireland has a separate regulator and sets policy independently. Unless stated otherwise, we focus on Great Britain and, for brevity, refer to it as “the UK”.

Direct debit customers use two common payment arrangements: (i) *smoothed* direct debit, under which the customer pays a fixed monthly instalment based on expected annual usage; suppliers periodically review this amount (e.g., quarterly, semi-annually, or after major price changes) and adjust based on actual consumption; and (ii) *variable* direct debit, under which monthly payments track actual usage recorded via smart meters or submitted meter readings.

2.2 Energy Price Shock

Global demand for energy surged after COVID-19 lockdowns ended, while political tensions with Russia—then the world’s largest exporter of natural gas—intensified, leading to a sharp rise in European wholesale gas prices. Prices spiked further in early 2022, coinciding with Russia’s full-scale invasion of Ukraine. These increases had an especially large impact on UK retail energy prices given the country’s heavy reliance on natural gas for both residential heating and electricity generation.


As wholesale energy prices rose, the Ofgem price cap became increasingly binding on suppliers. Figure 2.1 plots the cap over time (evaluated at Ofgem’s “typical” consumption values) alongside the average annual bills—at the same consumption values—for the ten cheapest pricing plans in the market. In January 2020, the average of the ten cheapest plans was about 70% of the cap; by January 2021 it was 90%, and from October 2021 onward it was 99%. As legacy contracts expired, consumers were moved (actively or passively) onto plans that tracked the cap.

We use the Ofgem cap to measure the unit prices households faced. A potential source of mismeasurement is that some consumers remained on legacy fixed-price plans whose unit price did not immediately track the cap. By the peak of the crisis in early 2023, fewer than 10% of consumers were on fixed-price contracts (Department for Energy Security and Net Zero, 2023c). Fixed contracts are rarer among households on variable billing: (i) for variable-direct-debit customers, suppliers typically offer “trackers” rather than fixed-price deals—standing charges are fixed, but unit prices move with the cap; and (ii) at the outset of the crisis, fewer than 1% of prepay households were on fixed-price contracts (Department for Business Energy and Industrial Strategy, 2022b). Hence, from mid-2021 the cap closely approximates the unit price faced by households on variable billing.

Cap adjustments translated into large movements in the real price of energy: evaluated at Ofgem’s typical consumption values, the real price rose by 8% in October 2021 and a further 45% in April 2022. The cap led to a further 73% rise in October 2022, at which point the UK government introduced a package of household relief policies. A further cap increase in

January 2023 implied that, absent policy intervention, households would have faced a real energy price around than three and a half times higher than two years earlier.

Figure 2.1: Energy price cap, energy price guarantee and cheapest available pricing plans



Notes: The figure reports the annual bill at Ofgem’s “typical” consumption values (12,000kWh of gas and 2,900kWh of electricity) for dual-fuel direct-debit consumers, in nominal terms (Ofgem, 2023). The “average of the cheapest direct-debit pricing plans” is a simple mean of the lowest-cost direct-debit plans by each of the 10 cheapest suppliers (one plan per supplier), and includes fixed-price plans; only plans generally available to consumers are included. Small regional differences in the cap exist across Great Britain; the figure shows the average value. There are also very small differences in the cap value applied to prepayment plans (see Appendix Figure A.1). We account for payment-type and regional variation throughout our analysis.

2.3 Policy Intervention

In September 2022, shortly before the October cap rise, the UK government introduced an energy-price subsidy—the Energy Price Guarantee (EPG)—and a universal transfer—the Energy Bill Support Scheme (EBSS)—to provide “urgent support” to “millions of families” (Department for Business Energy and Industrial Strategy, 2022a). Two considerations motivate such interventions: exposure to energy price shocks varies substantially across households, and the government is better placed than many households to smooth temporary shocks.⁹

The EPG imposed a ceiling on unit rates and standing charges for gas and electricity at levels below Ofgem’s cap and compensated suppliers for the difference. Rates were set so that a “typical” household’s annual bill would not exceed £2,500. In practice, the EPG operated from October 2022 through June 2023; as wholesale prices fell, the Ofgem cap

⁹Many households face constraints that generate large consumption falls even in response to temporary income 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; most are “wealthy hand-to-mouth,” owning illiquid assets but holding limited cash on hand.

dropped below the EPG level, rendering the guarantee non-binding. The red line in Figure 2.1 shows the cap evaluated at typical consumption values. The EPG limited the October 2022 increase in the real energy price faced by households to 22% and prevented a further rise in January 2023. On average, it implied a subsidy of 39% on the marginal price of energy over October 2022–March 2023.¹⁰ Although, even with the EPG in place, unit prices in October 2022 were higher than their pre-EPG levels, standing charges were essentially unchanged. After the EPG ended in June 2023, prices declined, with a 23% fall in the real energy price in July and a further 6% in October.

The EBSS provided all households with a transfer of £400 paid in six monthly instalments from October 2022 to March 2023. Households received support as: (i) vouchers or automatic meter top-ups for prepayment energy; (ii) credit applied to direct-debit accounts; or (iii) direct cash refunds. All transfer modes were administered by electricity suppliers. In addition, the UK government provided direct “cost-of-living” cash transfers to households in receipt of means-tested benefits from summer 2022 until spring 2024.¹¹

2.4 Data

Our main dataset contains transaction-level information on spending and after-tax-and-transfer income, derived from individuals’ bank account and credit card statements. It is provided by analytics firm ExactOne and collected via ClearScore, a fintech company that helps users monitor and manage their finances. A key advantage over single-bank datasets is that users are encouraged to link all accounts and cards, including those jointly held with a spouse/partner. We refer to users as households. The dataset provides a relatively comprehensive picture of households’ incomes and spending patterns and covers just over a quarter of a million UK households over 2019–2023. We provide additional details on the data construction and representativeness in Appendix A.

We focus on households with at least one account that records energy spending and aggregate all linked accounts to the household-year-month level. We construct monthly measures of energy spending (gas + electricity),¹² non-durable spending, and income. All monetary variables are expressed in 2022 prices unless otherwise specified.

¹⁰From October 2022 to March 2023, electricity unit prices under the EPG averaged 42% below the cap and gas unit prices 35% below. Weighting by 2019 average spending shares on electricity and gas yields a 39% average subsidy. Over the full EPG period (October 2022–June 2023), the average implied subsidy was 35%.

¹¹Cost-of-living payments were paid directly into recipients’ bank accounts in six instalments (values ranging from £299 to £326). See Appendix A.2 for details.

¹²We pool gas and electricity spending because most households pay a single dual-fuel bill: in the 2019 LCFS about 80% of electricity customers and 76% of gas customers have dual-fuel bills. As a result, we cannot separate gas and electricity spending. In the UK, 99.9% of households use electricity and 85% use gas (Department for Business Energy and Industrial Strategy, 2022c).

We use billing type to identify households whose energy spending closely tracks usage, based on the recorded payment method (e.g., direct debit or card payment) and the variability of energy spending. We distinguish: (i) smoothed direct debit (69% of ExactOne households), (ii) variable direct debit (6%); (iii) prepayment (25%).¹³ The full ExactOne sample (all billing types) comprises just over 280,000 households from January 2019 to December 2023,¹⁴ while our analysis sample consists of households that pay for all their energy either through variable direct debit or prepayment, yielding about 1.2 million household-month observations for 75,000 households from June 2021 to December 2023.

Representativeness. Compared to the nationally representative LCFS, ExactOne households are slightly younger and more likely to live in the North of England; prepayment contracts are also over-represented. After reweighting on age, region and payment type, the distributions of monthly income, energy spending, and non-durable spending are closely aligned across the two datasets. (see Appendix A). Throughout our analysis we reweight by payment type, and show our results are robust to additionally reweighting by age and region.

Our analysis sample includes households on variable billing (variable direct debit and prepayment). Variable-direct-debit households (i) exhibit much more seasonality in energy spending than for smoothed-direct-debit households—consistent with payments tracking usage—and (ii) have distributions of monthly income and energy spending, as well as age and region, that closely resemble those of smoothed-direct-debit households (see Appendix A). This indicates that variable-direct-debit spending tracks usage and that these households are informative for the broader smoothed-billing population. Prepayment households have lower incomes and lower energy spending than direct-debit households. Finally, Figure 2.2 validates our data and sample construction: crisis-period trends in energy spending in our ExactOne analysis sample align with the UK National Accounts.

3 Exposure to Energy Price Shocks


3.1 Income

We begin by documenting the relationship between energy spending and household income. This helps to understand (i) how the incidence of energy price shocks varies across the income distribution, and (ii) the extent to which income-based transfers can compensate households for such shocks.

¹³Roughly 7% of households pay using “standard credit”, i.e., on receipt of bill for actual usage (monthly or quarterly). We group these with smoothed-direct-debit households.

¹⁴We require that a household is observed for at least six months between June 2021 and December 2023.

Figure 2.2: Trends in energy spending in ExactOne data and National Accounts

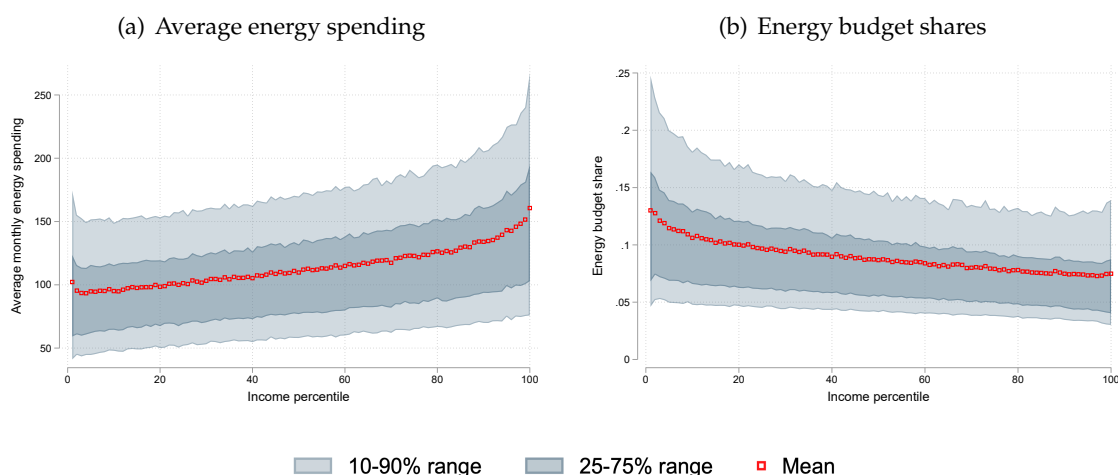


Notes: The figure plots deseasonalised monthly energy spending (indexed to 1 in 2021) from the UK National Accounts and ExactOne data. We deseasonalise the ExactOne data by subtracting calendar-month effects estimated over the pre-crisis period (2019–2020); the National Accounts series uses the Office of National Statistics’ published deseasonalisation. The ExactOne series is constructed using the analysis sample with household fixed effects and is reweighted by payment type. Spending is expressed in 2022 prices.

Figure 3.1 summarises energy spending over the two-year period prior to the crisis (2019–2020). Panel (a) plots the mean, interquartile range, and interdecile range of households’ average monthly energy spending by income percentile. Higher-income households tend to spend more in levels: those in the top income decile spent, on average, £143 per month, compared with £95 for the bottom decile. Panel (b) shows that, although higher-income households tend to spend more in levels, lower-income households tend to devote a larger *fraction* of their budget to energy: about 12%, on average, in the bottom income decile versus 7% for the top decile. Consequently, lower-income households are exposed to larger *proportional* losses when energy prices rise. These patterns remain similar when conditioning on total non-durable expenditure rather than income (see Appendix Figure B.1).

Variation in income nonetheless accounts for only 7% of the cross-sectional variation in energy spending, with substantial dispersion conditional on income. For example, 10% of households in the bottom income decile spent more than £154 per month on energy—exceeding the average among households in the top decile. This is mirrored in budget shares: within the bottom income decile, 10% of households spent more than 20% of their budget on energy. Taken together, this suggests that income-based transfers alone are likely to have limited effectiveness in compensating those most exposed to energy price rises.

Figure 3.1: Energy spending across the income distribution



Notes: Panel (a) summarises the distribution (mean and 10th, 25th, 75th, 90th percentiles) of households’ mean monthly energy spending in the pre-shock period (2019–2020), by income (also measured over 2019–2020). Panel (b) shows the analogous distribution of energy budget shares (energy spending over non-durable spending). Spending is expressed in 2022 prices. Results are very similar when restricting to 2019 only (Appendix Figure B.1), indicating the patterns are not driven by pandemic-specific factors.

3.2 Household Characteristics

What explains the substantial variation in energy spending conditional on income? If the relevant drivers are observable, policymakers could use them to design better-targeted compensatory transfers. The ExactOne data contain limited demographics, so we supplement them with survey data from the LCFS; see Appendix B.2 for details.

In the ExactOne data, a cross-sectional regression of households’ average monthly energy spending on income percentile indicators yields $R^2 = 0.069$, i.e., income explains about 7% of the variation in energy spending. Repeating the exercise with the LCFS gives a very similar estimate. Adding demographic variables that are plausibly observable to government (age and employment status of the household head, number of adults, number of children, and region) increases the R^2 to 0.161. Including housing characteristics, such as number of rooms and local property tax (council tax) payments, raises the R^2 further, but collectively these observables still only explain about 20% of the cross-sectional variation.

This limited explanatory power is not specific to energy. Performing analogous exercises for groceries and vehicle fuel, income plus demographics explain more of the variation (about 35% for groceries and 24% for vehicle fuel), yet substantial residual dispersion remains. These findings suggest that purely income-based (or income-and-demographic-based) transfers can only partially compensate households for shocks to the prices of staple goods.

3.3 Persistence of Energy Usage

Another way to target support toward households most exposed to energy price rises is to base transfers on historical energy usage. This is administratively feasible because payments can be delivered through energy suppliers, who observe past consumption (unlike for other goods, such as food or vehicle fuel). The effectiveness of this approach depends on the persistence of household energy use: if current high use is only weakly predicted by past use, history-based transfers will leave many households exposed. We estimate a one-year autocorrelation coefficient of 0.74 for log monthly energy spending, with an associated R^2 of 0.55 (see Appendix B.3). This indicates substantial persistence—past usage explains roughly half of the variation in subsequent spending—but also meaningful change for some households over time. Consequently, transfers tied to historical usage may improve targeting relative to income-based payments, yet may inadequately support some households whose energy needs rise and over-support those whose needs fall.

4 Household Responses to an Energy Crisis

In this section, we first describe how households' energy spending evolved over the crisis, then estimate price elasticities of demand for energy and the marginal propensity to consume energy out of transfers.

4.1 The Evolution of Energy Spending


Figure 4.1 plots changes in deseasonalised log energy spending from June 2021 to June 2023. Vertical dashed lines mark the dates when the energy price cap was updated and, in April 2023, when the EBSS energy-support transfer programme ended. Spending rises at the 8% and 45% cap increases, but by less than the price changes, indicating inelastic demand.¹⁵ Spending also shifts around the introduction and withdrawal of transfers, indicating an economically meaningful marginal propensity to consume energy out of these funds.

4.2 Price Responsiveness

To estimate how household energy use responds to price changes, we focus on June 2021–September 2022. We exclude the period after the October 2022 cap change because it coincides with the introduction of transfers; we analyse transfer-induced responses below.

¹⁵The response to the 8% cap rise is gradual. Below we show that estimated price elasticities are similar when we use variation arising only from the April 2022 cap rise.

Figure 4.1: Log energy spending over the crisis



Notes: The figure plots deseasonalised log monthly energy spending. Deseasonalisation subtracts calendar-month effects estimated from the pre-crisis period (2019–2020). Spending and cap changes are expressed in 2022 prices; the nominal cap increases are 12% (October 2021), 54% (April 2022), and 27% (October 2022).

We estimate variants of:

$$\log x_{i\tau}^e = \sum_{d \in \{DD, PP\}} \left(\gamma_d \log p_{r(i)\tau}^e + g_d(\mathbf{weather}_{a(i)\tau}; \theta_d) + \text{adjust}_{d\tau} \right) + \zeta_i + \epsilon_{i\tau}. \quad (4.1)$$

where $x_{i\tau}^e$ is deseasonalised energy spending net of the fixed standing charge for household i in year-month τ . $p_{r(i)\tau}^e$ is the marginal (unit) price of energy faced by households in region r at time τ .¹⁶ $g_d(\mathbf{weather}_{a(i)\tau}; \theta_d)$ is a function of local weather in area a ,¹⁷ $\text{adjust}_{d\tau}$ are indicators for the months immediately before and after each cap change (to capture anticipation or lagged adjustment); and ζ_i is a household fixed effect. We allow coefficients γ_d to differ by payment method $d \in \{\text{direct debit (DD)}, \text{prepayment (PP)}\}$.

Our key identification assumption is that—conditional on household fixed effects, seasonal dummies, and flexible local weather controls—no other time-varying determinants of energy demand are systematically correlated with the timing or magnitude of changes in the regulatory price cap. Cap revisions are pre-announced and driven by global whole-

¹⁶We use the Ofgem cap (or the EPG when in place) to measure the electricity and gas unit prices, and we combine them into an energy price index with base (2019–2020) expenditure shares. As electricity and gas quantity shares are stable over time (Appendix C.1), using an index that accommodates substitution between gas and electricity makes very little difference; we also show robustness to household-varying weights.

¹⁷ Weather data are from the UK Met Office, (Appendix A.4). We include 5th-order polynomials in minimum and maximum temperatures, the squared difference between them, rainfall, and humidity: $g_d(\mathbf{weather}_{a\tau}; \theta) = \sum_{p=1}^5 \theta_{\min, d}^p \text{tmin}_{a\tau}^p + \sum_{p=1}^5 \theta_{\max, d}^p \text{tmax}_{a\tau}^p + \theta_{\text{diff}, d} (\text{tmax}_{a\tau} - \text{tmin}_{a\tau})^2 + \theta_{\text{rain}, d} \text{rain}_{a\tau} + \theta_{\text{humid}, d} \text{humid}_{a\tau}$.

sale energy costs rather than contemporaneous household demand, which supports this assumption. Below we show robustness of our estimates in several ways.

Average price elasticities

We estimate a log-log specification so that weather controls enter *proportionally*, rather than in levels. For small price changes, the coefficient on $\log p$ minus one approximates the percent change in quantity associated with a one percent change in price. For large price changes this is not this case. We therefore map the estimates into the implied percentage change in quantity for the April 2022 price increase and divide by the corresponding percentage change in real price to report a finite-change elasticity. We compute separate elasticities for each payment method and then average using shares from the LCFS (85% direct debit and 15% on prepay). Table 4.1 summarises the results.

We estimate that the quantity of energy demanded fell by 15% on average when the real price rose by 45% in April 2022, implying an own-price elasticity of energy of -0.33 . Column (2) instruments the energy price with the discrete cap changes, isolating policy-driven movements, and yields a similar estimate. Column (3) restricts to households that switched supplier between June and December 2021—for whom we therefore know the cap binds—and also yields similar results.¹⁸

A potential concern is that, because our baseline analysis sample consists of prepay consumers and households that opt into variable direct debit, responses could differ from the broader population. To assess this, we use the fact that a subset of smoothed direct-debit customers are with suppliers that regularly review direct debit payments.¹⁹ Consumers were strongly advised to submit regular meter readings around cap changes, so these households' payments are likely to track usage reasonably well. Column (4) reports an elasticity statistically indistinguishable from the baseline, supporting external validity.

Our results are also robust to: using log quantity as the dependent variable; alternative deseasonalisation and weather controls; relying only on the April 2022 cap change; and reweighting to match the population by age and region. In addition, when we augment the specification with Northern Ireland—where the Ofgem cap did not apply—to difference out UK-wide shocks (including cost-of-living pressures) common to both markets, the estimates are unchanged within sampling error (Appendix C.2).

¹⁸Due to escalating costs, several suppliers ceased trading in late 2021, and their consumers were automatically transferred to the standard variable price plans (priced at the cap).

¹⁹Most suppliers review at least quarterly and at price cap changes; we exclude customers of British Gas and EDF, which review less frequently. We also require a supplier switch between June and December 2021 to avoid legacy fixed-price plans.

Table 4.1: Energy price elasticities

	(1)	(2)	(3)	(4)
% Δe over cap change	-14.8 (0.5)	-14.6 (0.5)	-14.9 (1.3)	-16.2 (0.2)
Own-price elasticity	-0.326 (0.011)	-0.320 (0.011)	-0.327 (0.029)	-0.357 (0.005)
N	757,286	757,286	685,579	1,060,153
Instrument with cap changes	No	Yes	No	No
Restricts variable DD to new supplier	No	No	Yes	No
Includes smoothed DD with frequent DD reviews	No	No	No	Yes

Notes: The dependent variable is deseasonalised log energy spending net of the standing charge. For each payment type d , we convert the log–log estimate $\hat{\gamma}_d$ into a percent change in quantity for the April 2022 price change using $\Delta e_d / e_d = \exp((\hat{\gamma}_d - 1)\Delta \log p) - 1$, where $\Delta \log p \equiv \log p_1 - \log p_0$. We also report the corresponding finite-change elasticity: $\epsilon_d = \frac{\Delta e_d / e_d}{(p_1 - p_0) / p_0} = \frac{\exp((\hat{\gamma}_d - 1)\Delta \log p) - 1}{\exp(\Delta \log p) - 1}$. p_1 is the real price after the April 2022 increase and p_0 the real price over Oct 2021–Mar 2022. We average $\Delta e_d / e_d$ and ϵ_d over payment types using LCFS population weights. All specifications include weather controls (5th-order polynomials in local monthly minimum and maximum temperature, the squared difference between maximum and minimum temperature, rainfall, and humidity) and indicators for the months immediately before and after a cap change. Standard errors are shown in parentheses and clustered at the household level and, for elasticities, computed via the delta method.

Heterogeneity in price elasticities

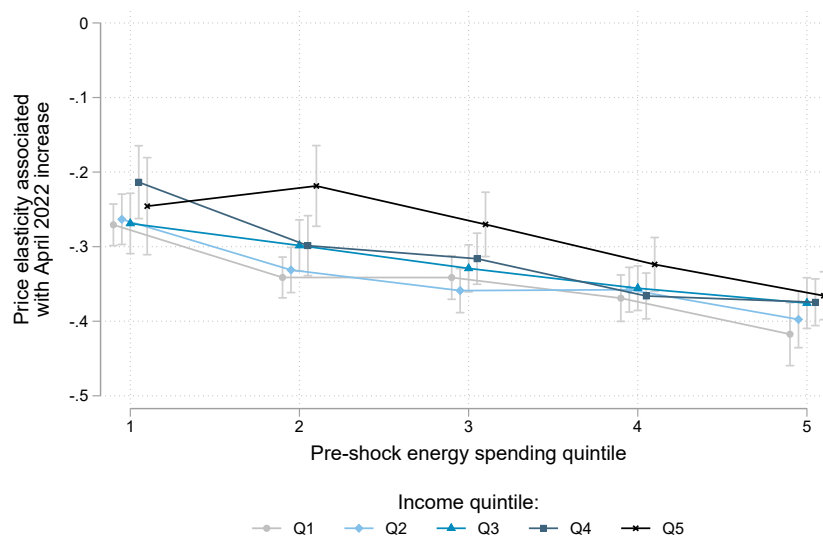
Variation across households in their response to price changes is a key driver of the distribution of welfare effects from energy price shocks. Figure 4.2 shows how consumption responses to the April 2022 price increase vary across groups defined by pre-shock energy spending quintile and income quintile (25 groups in total). Concretely, we estimate equation (4.1) allowing the price coefficient to vary by indicators for income quintile \times pre-shock energy spending quintile and plot the resulting elasticities for each group.

Households with higher pre-crisis energy spending have more elastic demand: the average elasticity varies from -0.25 in the bottom pre-shock spending quintile to -0.39 in the top quintile. Households in the highest income quintile have the least elastic demand; the top-bottom income quintile differences are statistically significant at conventional levels for every pre-shock energy spending quintile except the first.

4.3 Marginal Propensity to Consume Energy Out of Transfers

We estimate households’ marginal propensity to consume energy (MPCE) out of transfers and investigate whether this differs for the energy-support transfers relative to direct cash transfers. One reason for a higher MPCE from the former would be if some of the households that received payment as vouchers would otherwise have spent less on energy than the voucher amount. This appears unlikely in our setting as the large majority of households spent more on energy than the monthly transfer value. A second possibility is a “flypaper

Figure 4.2: *Heterogeneity in energy price elasticities, by pre-shock energy spending and income*



Notes: Each point reports the price elasticity for the April 2022 price increase, by income quintile and pre-shock energy spending quintile. Price responses are allowed to differ by payment type (DD, PP); we average across payment types using LCFS shares that vary by income \times pre-shock spending cell. Pre-shock spending is measured over 2019–2020. Results are robust to constructing the price index with gas and electricity weights that vary across the 25 groups (Appendix C.2)

effect,” whereby money labelled for energy sticks to energy spending (Hines and Thaler, 1995). In household contexts, such non-fungibility has been documented in several settings (see Jacoby, 2002; Choi et al., 2009; Fafchamps et al., 2014).

We focus on prepay households, who received support as vouchers or automatic top-ups on their meters.²⁰ 83% of prepay households had monthly energy spending in excess of the voucher value immediately before the programme. Moreover, prices rose at the same time as the transfer programme began, further reducing the likelihood that the voucher exceeded desired spending. For those whose desired energy spending in a given month was below the voucher value, credit could be added to the meter or redeemed later (until June 2023). Thus, it is unlikely that the transfers acted as a binding constraint forcing households to consume more energy than they otherwise would.

Figure 4.3 plots gross energy spending (inclusive of the energy-support transfer value) for prepayment households. Spending rises sharply during the transfer period and declines once transfers cease. The dashed line shows the counterfactual path implied by our estimated price responses in the absence of transfers (given observed price and weather). Actual spending during the transfer period is substantially above this counterfactual, consistent

²⁰Vouchers/credits could typically be redeemed against both electricity and gas, with flexibility over how to allocate the vouchers. Exceptions were British Gas prepay customers and the small share (< 2%) with different suppliers for electricity and gas, who could not transfer credit to gas. Excluding these households does not affect our results.

with a substantial MPCE out of energy-support transfers. Although households could, in principle, carry forward credit (by accumulating it on the meters or redeeming vouchers later), we observe only a small dip below the counterfactual after the programme ends, indicating little intertemporal shifting.

Figure 4.3: Energy spending over the transfer period, prepayment households



Notes: Vertical dashed lines mark the start and end of the energy-support transfers window (Oct 2022–Mar 2023). The figure plots two series for deseasonalised log energy spending for prepayment household, where deseasonalisation subtracts calendar-month effects estimated from 2019–2020. The solid line is actual deseasonalised spending. The dashed line is a counterfactual predicted from a household fixed-effects panel regression, which includes weather controls, estimates price effects using variation outside the transfer window only, and removes time effects for the transfer period.

We formalise this by estimating:

$$\log x_{i\tau}^e = \gamma \log p_{r(i)\tau}^e + g(\mathbf{weather}_{a(i)\tau}; \theta) + \mu^E \text{transfer}E_\tau + \mu^{EP} \text{posttransfer}E_\tau + \mu^G \text{transfer}G_{i\tau} + \zeta_i + \epsilon_{i\tau} \quad (4.2)$$

where $\text{transfer}E_\tau$ is an indicator equal to 1 during the energy-transfer window (October 2022–March 2023); $\text{posttransfer}E_\tau$ equals 1 for April 2023 to capture any drawdown of accumulated credit; and $\text{transfer}G_{i\tau}$ equals 1 if household i received a cost-of-living payment in month τ or $\tau - 1$. All other variables are defined as in equation (4.1). We estimate equation (4.2) over June 2021–December 2023.

Using the estimated coefficients, we compute the MPCE out of transfers; Table 4.2 reports the results. The MPCE out of energy-support transfers is 0.38 (column (1)), substantially above the average energy budget share for prepay households (14%) over this period. Accounting for any reduction in spending in the month after the transfer window (adding $\text{posttransfer}E_\tau$; column (2)) reduces the estimate by only a small amount (to 0.34), consistent with Figure 4.3. In contrast, the estimated MPCE of the broader cost-of-living transfer

administrated directly by the government is 0.03. This gap is consistent with a pronounced flypaper effect for the energy-support transfers among prepay households.

Table 4.2: *Marginal propensity to consume energy out of transfers*

	(1)	(2)
Energy-support transfers (distributed via suppliers)	0.38 (0.006)	0.34 (0.007)
Cost-of-living transfers	0.03 (0.002)	0.03 (0.002)
N	1,098,240	1,098,240
Adjust for post-transfer spillover	No	Yes

Notes: The dependent variable is deseasonalised log monthly energy spending. In column (1), the MPCE out of energy-support transfers is $MPCE = \Delta_{transferE_\tau} \mathbb{E}[x_{it}^e | Z = \bar{z}] / \mathbb{E}[transfer\ value] = (\exp(\hat{\mu}^E) - 1) \bar{x}_{base}^E / \bar{R}^E$, where \bar{x}_{base}^E is the baseline level of spending during the transfer window, given by: $\bar{x}_{base}^E = \exp(\hat{\eta}_{it}) \times \exp(\hat{\epsilon}_{it})$, with $\hat{\eta}_{it}$ the fitted log outcome evaluated at mean covariates $Z = \bar{z}$ and $transferE_\tau = 0$. \bar{R}^E is the average real monthly transfer scaled by 0.9 to reflect 90% voucher redemption (Department for Energy Security and Net Zero, 2023b). Column (2) nets out the post-transfer window drawdown by computing $\Delta_{posttransferE_\tau} \mathbb{E}[x_{it}^e | Z = \bar{z}] = (\exp(\hat{\mu}^{EP}) - 1) \bar{x}_{base}^{E, April}$ in the same way and subtracting it from the transfer-window effect, scaled for the relative window lengths (one vs six months). For cost-of-living payments, $transferG_{it} = 1$ in the payment and following month. The per-payment $MPCE^G = 2(\exp(\hat{\mu}^G) - 1) \bar{x}_{base}^G / \bar{R}^G$ evaluated at mean covariates over the payment period. Standard errors are in parenthesis and clustered at the household level. All regressions include household fixed effects and weather controls.

Households paying by direct debit received the energy support either as an account credit or a cash refund. We find no evidence of a flypaper effect for transfers paid as cash refunds, despite being labelled “energy bill support” and distributed via energy suppliers. For transfers delivered as account credits, we detect a positive but quantitatively small flypaper effect (Appendix C.3). Taken together, the results indicate that the pronounced flypaper effect we document for prepay households—who received vouchers—is not driven by labelling alone but the interaction of labelling with the payment instrument. This pattern is consistent with recent experimental evidence that marginal propensities to consume vary with the mode of transfer (Boehm et al., 2025).

5 Model of Energy Demand

In this section we develop an empirical model of household energy demand to quantify the welfare effects of a large energy price shock and to evaluate policy responses—including those used during the European Energy Crisis and alternatives. As we are interested in the full distribution of welfare effects, we use a flexible demand specification that captures rich preference heterogeneity across households.

5.1 Household Choice Model

Consider a household's consumption decision over 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 energy price (inclusive of any subsidy) 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. Net available budget is $x \equiv \tilde{x} - f + t$. An optimising household solves the problem:

$$V(p^e, p^n, x; \theta) = \max_{e, n} U(e, n; \theta) \text{ s.t. } p^e e + p^n n \leq x. \quad (5.1)$$

Let $e = e^0(p^e, p^n, x; \theta)$ and $n = n^0(p^e, p^n, x; \theta)$ denote the resulting Marshallian demand functions. The superscript $d = 0$ indicates 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 privately suboptimal choices. Let $(e^1(p^e, p^n, x; \theta), n^1(p^e, p^n, x; \theta))$ denote the (observed) energy and other non-durable choices when the flypaper effect operates; the superscript $d = 1$ indexes suboptimal choices.

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

Assumption 1.

- (a) *The flypaper effect affects utility only through the induced change in the chosen consumption bundle. Therefore the utility attained at suboptimal choices is given by:*

$$U(e^1(\cdot), n^1(\cdot); \theta)$$

- (b) *In the absence of a flypaper effect—i.e., for non-prepay households and for prepay households when no energy-support transfers are in place—households choose the privately optimal bundle that solves equation (5.1).*

Our approach applies Bernheim and Rangel's (2009) choice-based welfare framework. Part (b) identifies environments in which observed choices reveal privately optimal bundles, while part (a) restricts departures so that welfare can be recovered when the flypaper effect operates, without fully specifying the mechanism behind the departure from optimisation. The key restriction is that any impact of the flypaper effect on utility operates only through

the consumption bundle, which means our efficiency-cost calculations are net of intrinsic (e.g., cognitive) costs of making choices other than $(e^1(\cdot), n^1(\cdot))$. This is similar to the assumption in Chetty et al. (2009) in their study of tax salience. Unlike that work, which relies on a first-order approximation, we embed this approach in a fully specified choice model and analyse non-marginal price and policy changes; in our setting this is important for accurately recovering welfare effects.

Our empirical model is based on a flexible form for the expenditure function, so it is convenient to measure welfare using indirect utility. The following proposition shows how we accommodate suboptimal choices:

Proposition 1. *Let $(e^0, n^0) = (e^0(p^e, p^n, x; \theta), n^0(p^e, p^n, x; \theta))$ be the privately optimal bundle solving problem (5.1). Let (e^1, n^1) be a feasible suboptimal bundle with $p^e e^1 + p^n n^1 = x$ and $e^1 > e^0$. Then there exists a unique ϕ such that, with*

$$p^{e'} = (1 - \phi) p^e, \quad x' = x - \phi p^e e^1,$$

the observed energy choice is Marshallian demand at $(p^{e'}, p^n, x')$, i.e.,


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

Define attained utility by

$$\mathbb{V}(p^e, p^n, x, \mathfrak{d}; \theta) = \begin{cases} V(p^e, p^n, x; \theta), & \mathfrak{d} = 0 \text{ (privately optimal)}, \\ V(p^{e'}, p^n, x'; \theta), & \mathfrak{d} = 1 \text{ (privately suboptimal)}. \end{cases}$$

Under continuous, strictly monotone, and strictly convex preferences, such a ϕ exists and is unique; in the $e^1 > e^0$ case it satisfies $\phi \in [0, 1]$ (the $e^1 < e^0$ case is symmetric, involving a compensated increase in the energy price). ϕ is the Slutsky-compensated reduction in the energy price that, together with the budget offset $-\phi p^e e^1$, pivots the budget line through (e^1, n^1) so that it becomes optimal. Figure 5.1 visualises Proposition 1. Starting from the baseline budget line, the suboptimal flypaper bundle (e^1, n^1) lies below the optimal indifference curve. A Slutsky-compensated pivot rotates the budget line through (e^1, n^1) so that it becomes tangent to an indifference curve at that point. With two goods, this single-price pivot plus budget adjustment is equivalent—by homogeneity—to an appropriate rescaling of (p^e, p^n) at fixed x ; with many goods, one can evaluate welfare at the full supporting (virtual) price vector that rationalises the observed bundle (Neary and Roberts, 1980). Although we apply Proposition 1 to the flypaper effect, it applies more broadly to any setting in which households make privately suboptimal choices.

Figure 5.1: Hypothetical budget set that rationalises suboptimal choice



Notes: OAB is the baseline budget at (p^e, p^n, x) . The privately optimal bundle (e^0, n^0) lies on indifference curve \bar{U} . The observed bundle (e^1, n^1) is feasible ($p^e e^1 + p^n n^1 = x$) but yields lower utility $\bar{U}' < \bar{U}$. The virtual budget line CD corresponds to $(p^{e'}, p^n, x')$ with $p^{e'} = (1 - \phi) p^e$ and $x' = x - \phi p^e e^1$; it passes through (e^1, n^1) and is tangent to an indifference curve at the point. We evaluate attained utility at $V(p^{e'}, p^n, x'; \theta)$.

5.2 Empirical Specification

We estimate household energy demand using data for households i and year-months τ on: energy budget shares $w_{i\tau} \equiv \frac{p_{r(i)\tau}^e e_{i\tau}}{x_{i\tau}}$; total net-budgets $x_{i\tau}$ (total non-durable expenditure net of standing charges and inclusive of transfers); prices $(p_{r(i)\tau}^e, p_{\tau}^n)$; and conditioning variables $\mathbf{z}_{i\tau}$.²¹ We also include an indicator $\mathbb{d}_{i\tau}$ equal to one for prepay household-months during the energy-support transfer window. This term captures any flypaper effect (i.e., responses beyond pure income effects) associated with energy-support transfers.²²

We estimate a flexible parametric demand equation using the Exact Affine Stone Index (EASI) system (Lewbel and Pendakur, 2009). We capture preference heterogeneity with a rich set of conditioning variables—including measures of past energy usage. This yields tractable, theory-consistent heterogeneous demands. The model entails specifying a form for the expenditure function that gives rise to *implicit* Marshallian budget share demands:

²¹We construct a (non-energy) non-durables price index p_{τ}^n from the CPI microdata using the official UK CPI methodology (see Appendix A). All financial variables are in real terms.

²²We model estimate $\mathbb{e}^{\mathbb{d}}(p^e, p^n, x; \theta)$ for $\mathbb{d} = \{0, 1\}$. Under Assumption 1, we could omit “suboptimal” observations (prepay household-months during the transfer window) and recover $\mathbb{e}^0(p^e, p^n, x; \theta)$, which suffices for welfare evaluation under observed policy. We instead model suboptimal demands so we can simulate counterfactuals that vary the magnitude of the transfer-induced flypaper effect and quantify heterogeneity in the flypaper effect within the prepayment sample.

$$\omega_{i\tau} = (A + \sum_{l \in \mathcal{Z}_1} A_l z_{i\tau l}) + (B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l}) \times (\log p_{r(i)\tau}^e - \log p_\tau^n) + (C_1 + \sum_{l \in \mathcal{Z}_2} C_{1l} z_{i\tau l}) y_{i\tau} + C_2 y_{i\tau}^2 + D (\log p_{r(i)\tau}^e - \log p_\tau^n) \times y_{i\tau} + (\delta + \sum_{l \in \mathcal{Z}_3} \delta_l z_{i\tau l}) \mathfrak{d}_{i\tau} \quad (5.2)$$

$$y_{i\tau} = \frac{\log x_{i\tau} - (\omega_{i\tau} \log p_{r(i)\tau}^e + (1 - \omega_{i\tau}) \log p_\tau^n) + \frac{1}{2} (B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l}) \times (\log p_{r(i)\tau}^e - \log p_\tau^n)^2}{1 - \frac{1}{2} D \times (\log p_{r(i)\tau}^e - \log p_\tau^n)^2}, \quad (5.3)$$

where

$$\Psi \equiv (A, \{A_l\}_{l \in \mathcal{Z}_1}, B, \{B_l\}_{l \in \mathcal{Z}_2}, C_1, \{C_{1l}\}_{l \in \mathcal{Z}_2}, \{C_2\}, D, \delta, \{\delta_l\}_{l \in \mathcal{Z}_3})$$

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

Equations (5.2) and (5.3) implicitly define the energy demand function, which we write as $\omega_{i\tau} \equiv \omega(p_{r(i)\tau}^e, p_\tau^n, x_{i\tau}, \mathfrak{d}_{i\tau}, \mathbf{z}_{i\tau}; \Psi)$.²³ When $\mathfrak{d} = 0$, by construction, $y_{i\tau} = \log V_{i\tau}$, where $V_{i\tau}$ is realised utility level, and $\omega(\cdot)$ is the budget share form of the optimal demand function. When $\mathfrak{d} = 1$, $\omega(\cdot)$ deviates from the optimal choice. The budget share for other non-durables follows from adding-up: $\omega_{i\tau}^n = 1 - \omega_{i\tau}$.

By including log relative prices in equation (5.2), we ensure preferences satisfy adding-up, zero-degree homogeneity in prices and budget, and Slutsky symmetry. We do not impose the two *inequality* restriction from consumer theory in estimation—concavity of the expenditure function in prices and monotonicity in decision utility—but we check them post-estimation.

The parameters play distinct roles: A shifts the budget-share intercept; B controls the (compensated) price response; C determines the Engel-curve shape; D allows price responses to vary with the Engel curve; and δ captures the flypaper effect of the energy-support transfer for prepay households. We allow each effect to vary with a set of conditioning variables \mathcal{Z}_k with $\mathcal{Z}_3 \subset \mathcal{Z}_2 \subset \mathcal{Z}_1$. The set \mathcal{Z}_3 , which shifts the flypaper effect (δ), contains indicators for decile of the pre-shock energy spending; \mathcal{Z}_2 , which shifts price responsiveness (B) and the first-order Engel curve coefficient (in C), adds an indicator for whether the household prepays for energy; \mathcal{Z}_1 , which shifts the intercept (A), further adds indicators for pre-shock energy spending share decile, as well as indicators for household region (14 in total) and calendar-month, and detailed weather controls (as in equation (4.1)). This structure allows flexible heterogeneity in levels, price responses, Engel-curve shape, and flypaper effects.²⁴

²³This function does not have an analytical form; we solve for it numerically. A sufficient condition for the equations (5.2) and (5.3) to uniquely define $\omega_{i\tau}$ and $y_{i\tau}$ is that the expenditure function is monotone in utility. See Appendix D.1. We check that this condition is satisfied post-estimation.

²⁴We set $\mathfrak{d}_{i\tau} = 1$ for prepay household-months over the energy-support transfer window (October 2022–April 2023) and the first post month (April 2023) to allow for drawdown. For direct-debit households, we include indicators for the transfer and post-transfer periods.

Estimation

We estimate the model by GMM (Hansen, 1982). Let $\mathbf{h}_{i\tau}$ denote the vector of instruments and $\epsilon_{i\tau} \equiv w_{i\tau} - \omega(p_{r(i)\tau}^e, p_{\tau}^n, x_{i\tau}, \mathbf{d}_{i\tau}, \mathbf{z}_{i\tau}; \Psi)$ the model's prediction error. We specify the population moment condition $\mathbb{E}[\epsilon_{i\tau} \mathbf{h}_{i\tau}] = 0$ and obtain parameter estimates $\hat{\Psi}$ by selecting the value of Ψ that minimises the (quadratic form) of the sample analogue of moment conditions (see Appendix D.2 for details). We include in $\mathbf{h}_{i\tau}$ the conditioning variables $\mathbf{z}_{i\tau}$ and relative price $\log p_{r(i)\tau}^e - \log p_{\tau}^n$ (and the interactions that appear in the specification). Given that residential energy prices are adjusted at pre-specified times by a regulator and based on international wholesale costs, we treat prices as exogenous in estimation. To allow for correlation in shocks to energy demand and total budgets, we exclude $\log x_{i\tau}$ from $\mathbf{h}_{i\tau}$, instead including functions of monthly household income.

5.3 Model Estimates and Fit

We estimate the model using data from June 2021–June 2023 and assess out-of-sample fit using July–December 2023. The estimation sample includes prepay and variable-direct-debit households; in all policy simulations, we reweight to match the LCFS shares across these payment types. At the parameter estimates (see Appendix D.3), over 99% of household-month observations satisfy concavity of the expenditure function in prices, and all observations satisfy monotonicity in decision utility. The model also reproduces the pattern of elasticity heterogeneity across pre-shock energy spending and pre-shock income quintiles identified from responses to price-cap adjustments (Figure 4.2).

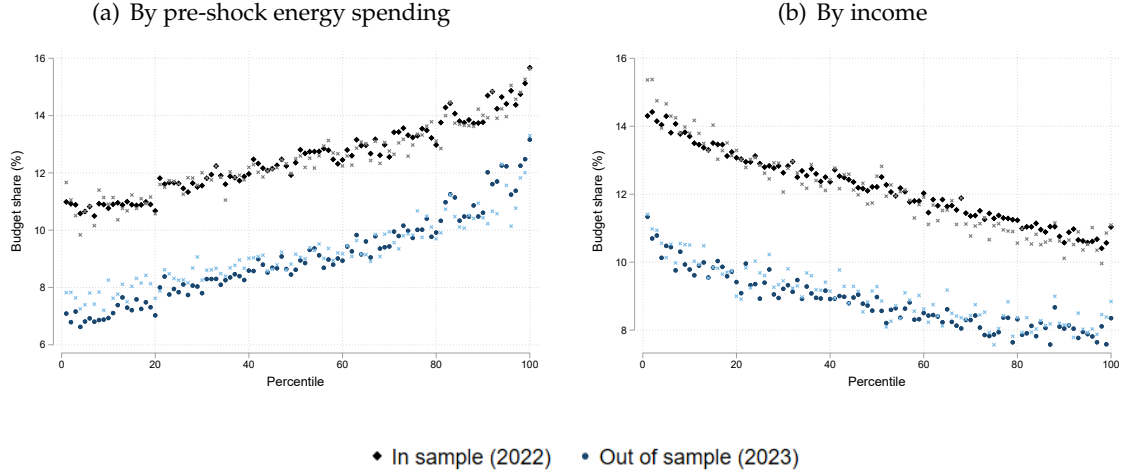
Our model decomposes responses to price changes into substitution and income effects. When an energy-price increase reduces consumption mainly through substitution, the welfare loss is smaller than when the reduction is driven by income effects. On average, for the April 2022 cap rise, the income effect accounts for approximately 10% of the Marshallian response.²⁵

In Figure 5.2 we assess the model's out-of-sample performance by comparing observed energy budget shares in October–December 2023 with model-based predictions. We also report the in-sample fit for October–December 2022, which shows the large decline in budget shares in the out-of-sample period (reflecting the fall in energy prices over 2023). Panel (a) groups households by percentiles of the pre-shock energy-use distribution; panel (b) groups by income percentiles. The model includes only coarse controls for pre-shock energy use

²⁵In Appendix D.3 we illustrate heterogeneity in income effects by plotting Engel curves separately for deciles of the pre-shock energy spending distribution. At low total expenditure, the gap in energy budget shares between the top and bottom deciles exceeds 15 percentage points; it shrinks to under 2 percentage points at the top of the expenditure distribution. These patterns mirror Lewbel and Pendakur (2017), who, using cross-sectional data, model preference heterogeneity via random Barten scales.

(deciles rather than percentiles) and excludes income directly from the demand system. Despite this, it replicates the cross-sectional gradients along both dimensions, including out of sample.²⁶

Figure 5.2: *Winter energy demand in and out of sample*



Notes: Panel (a) plots the average observed (crosses) and predicted (circles) budget shares across percentiles of the pre-shock energy spending distribution during Oct.–Dec. 2023 (out of sample) and Oct.–Dec. 2022 (in sample). Panel (b) shows analogous information across income percentiles. In Appendix D.3, we summarize the model fit jointly across pre-shock energy spending and income and validate the flypaper estimates using a hold-out sample.

6 Welfare Effects of an Energy Crisis and Policy Responses

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

6.1 Money-Metric Household Welfare

Consider household i in period τ facing prices $(p_{r(i)\tau}^e, p_\tau^n)$, with budget $x_{i\tau}$, flypaper indicator $\mathbb{1}_{i\tau}$, and conditioning variables and parameters $\theta_i = (\mathbf{z}_i, \Psi)$. During the crisis, the household's attained utility $\mathbb{V}(p_{r(i)\tau}^e, p_\tau^n, x_{i\tau}, \mathbb{1}_{i\tau}; \theta_i)$ (defined in Proposition 1) depends on policy parameters $\mathbb{P} = (s, t, L)$. s is the energy subsidy rate, so the consumer price relates to the pre-subsidy price $P_{r(i)\tau}^e$ according to $p_{r(i)\tau}^e = (1 - s)P_{r(i)\tau}^e$. The transfer amount t enters

²⁶In the EASI framework, one can treat estimation residuals either as structural preference shocks or as measurement error. Under the former interpretation, in-sample predictions match observed demands by construction. For the predictions in Figure 5.2, we adopt the latter interpretation (treating residuals as measurement error); yet the model recovers the systematic variation in demands with energy need and income, including out of sample. We maintain this interpretation in the simulations below.

the budget as $x = \tilde{x}_{i\tau} - f_{r(i)\tau}^e + t$, where $\tilde{x}_{i\tau}$ is pre-transfer gross budget and $f_{r(i)\tau}^e$ the energy standing (fixed) charge. $L \in \{0, 1\}$ indicates whether the transfer induces a flypaper effect among prepay households; we refer to such a transfer as “labelled” (using this term to capture both the “energy support” designation *and* supplier-based delivery).

For brevity, we write attained utility directly as a function of the policy parameters, $\mathbb{V}_{i\tau}(\mathbb{P})$. When a household’s choice is privately optimal ($\mathfrak{d} = 0$), we recover $\mathbb{V}_{i\tau}(\mathbb{P})$ by solving equations (5.2) and (5.3). When a labelled transfer induces a flypaper effect ($\mathfrak{d} = 1$), we first compute the budget pivot that rationalises the observed bundle as optimal (Proposition 1), and then solve equations (5.2) and (5.3) at that virtual budget set with $\mathfrak{d} = 0$ (see Appendix E.1).

We use a money-metric cardinalisation at pre-crisis prices $(p_{r(i)0}^e, p_0^n)$, denoted $\mathbb{V}_{i\tau}^{MM}(\mathbb{P})$.²⁷ The money-metric (equivalent variation) loss for household i in month τ from the energy-price shock under policy \mathbb{P} is $\mathcal{L}_{i\tau}(\mathbb{P}) \equiv (\tilde{x}_{i\tau} - f_{r(i)0}^e) - \mathbb{V}_{i\tau}^{MM}(\mathbb{P})$, where $(\tilde{x}_{i\tau} - f_{r(i)0}^e)$ is their pre-transfer budget net of the pre-shock fixed fee. Total loss over the shock window $\tau \in \{\underline{\tau}, \dots, \bar{\tau}\}$ is $\mathcal{L}_i(\mathbb{P}) = \sum_{\tau=\underline{\tau}}^{\bar{\tau}} \mathcal{L}_{i\tau}(\mathbb{P})$, interpretable as willingness-to-pay (at pre-shock prices) to avoid the shock under policy \mathbb{P} . We denote *proportional losses* by $l_i^y(\mathbb{P}) \equiv \frac{\mathcal{L}_i(\mathbb{P})}{Y_i}$ where Y_i is pre-shock household income (average monthly income over the preceding tax year, multiplied by the number of months in the shock window).

6.2 The European Energy Crisis and Observed Policy Response

Household losses

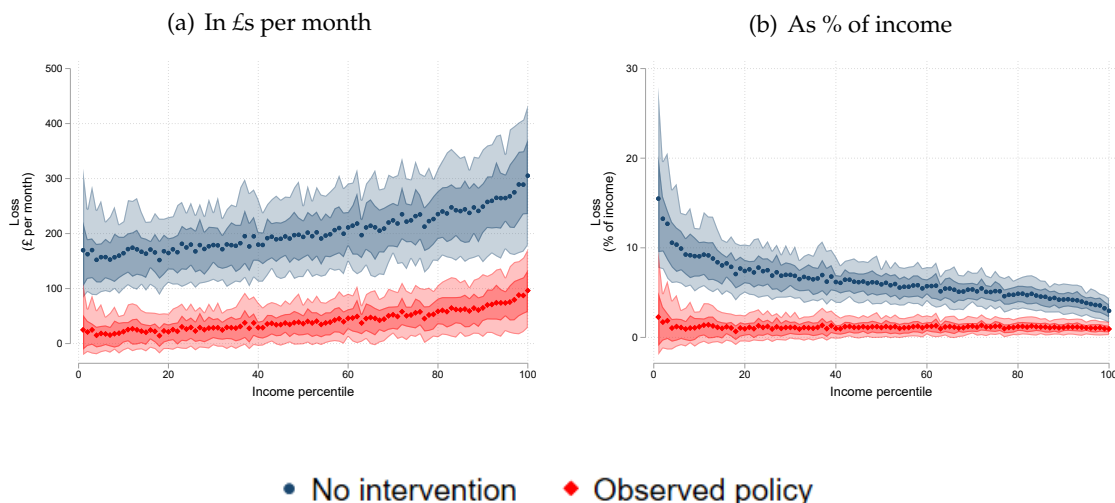
In Figure 6.1 we summarise the distribution of money-metric (panel a) and proportional (panel b) losses that households incurred during the height of the European Energy Crisis (October 2022 to March 2023). Markers show average losses; dark and light shading indicate the interquartile and interdecile ranges. Table 6.1 reports average and aggregate losses.

No government intervention. In the absence of government intervention, households would have been worse off by £207 per month on average (6.1% of income) worse off, implying an aggregate equivalent-variation loss over the six-month period of £35.2bn. Losses would also have been highly unequal: at the 95th percentile, proportional losses equal 11% of income. Although proportional losses decline with income on average, there is substantial variation *within* income groups. For example, within the bottom income decile, the 10th and 90th percentiles of proportional losses would have been 6 and 17%, respectively. The shock would have led to considerable hardship, with the number of households in “energy poverty”

²⁷Let $V_i^{\text{Pre}}(x) \equiv V(p_0^e, p_0^n, x; \theta_i)$ and let $\chi_i^{\text{Pre}}(u)$ be its inverse (expenditure) function. Then $\mathbb{V}_{i\tau}^{MM}(\mathbb{P}) \equiv \chi_i^{\text{Pre}}(\mathbb{V}_{i\tau}(\mathbb{P}))$.

(spending more than 10% of their after-housing-costs income on energy, see Appendix E.2) rising to over 9.6 million (compared with 3.9 million in 2019).

Figure 6.1: *Distribution of household losses by income*



Notes: For each household-month in Oct. 2022-Mar. 2023, we compute the loss from the price shock using equivalent variation, valued at pre-shock prices, and then average these losses over months for each household. The figures summarise how the distribution of household losses varies across percentiles of the Apr. 2021-Mar. 2022 income distribution. 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 implemented 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.

Table 6.1: *Average losses*

	No intervention	Observed policy	Under no behavioural response	
			No intervention	Observed policy
Monetary (£pm)	206.77 [204.66, 208.98]	43.38 [41.89, 44.83]	329.52 [322.07, 335.04]	81.19 [77.96, 83.58]
Proportional to Y (%)	6.10% [6.05, 6.16]	1.15% [1.11, 1.19]	9.75% [9.55, 9.90]	2.28% [2.20, 2.35]
Aggregate (£bn)	35.23 [34.87, 35.61]	7.39 [7.14, 7.64]	56.15 [54.88, 57.09]	13.84 [13.28, 14.24]
Number of households in energy poverty	9.6m	7.3m	21.7m	11.8m

Notes: The table reports average monthly losses and aggregate losses, constructed as described in the notes to Figure 6.1. “Under no behavioural response” refers to simulating welfare costs under the assumption that household demands remain at their pre-shock levels. Energy poverty is defined as spending more than 10% of after-housing-costs income on energy; see Appendix E.2 for further details. 95% confidence bands are reported in square brackets.

Observed policy response. Figure 6.1 shows that the UK policy response substantially reduced household losses, to £43 per month (1.1% of income). The intervention was particularly

effective at supporting the most exposed households: losses at the 95th percentile of the proportional-loss distribution are 3% of income, compared with 11% without intervention. The policy also flattened the relationship between mean proportional losses and income. It prevented 2.3 million households from falling into energy poverty—40% of the increase that would otherwise have occurred.

Importance of behavioural responses. The final two columns of Table 6.1 report aggregate losses from the energy crisis, both without policy intervention and under the observed policy response, when we *ignore* household behavioural responses. Doing so substantially overestimates aggregate equivalent-variation losses: by 59% in the absence of intervention and by 87% under the implemented policy response. For a shock of this magnitude, household substitution responses meaningfully mitigate welfare losses.

Policy efficiency costs

The revenue cost of the policy response was £30.9 billion, and it generated efficiency costs totalling £3.8 billion (12% of revenue). We decompose these costs in Table 6.2.²⁸

Efficiency costs of the subsidy. 70.3% of total efficiency costs arise directly from the dead-weight loss from subsidising energy—i.e., supporting households through a subsidy raises their utility by less than providing the same funds as a lump-sum transfer. This effect is linked to households' willingness to substitute away from energy when its price rises: larger substitution responses imply higher efficiency costs of a subsidy.

Efficiency costs of the transfer delivery. 11.5% of total efficiency costs are attributable to the flypaper effect among prepay households induced by the energy-support transfers. This operates through two channels. First, there is a direct private utility cost from suboptimal choices, accounting for 1.9% of total efficiency costs and borne by 15% of households on prepayment contracts. Among affected households, the flypaper effect entails average excess energy spending of £24 per month and an average monthly welfare loss of £3.²⁹ Second, a larger loss—more than fivefold—arises from a fiscal spillover, as the transfers stimulate additional consumption of a subsidised good.

²⁸We measure efficiency costs as the reduction in aggregate money-metric utility under the observed policy relative to a scenario in which each household receives a revenue-equivalent amount as a lump-sum transfer (see Appendix E.4). Because we base money-metric utility at *pre-shock* prices, we also value revenue and efficiency costs at pre-shock prices. Two alternative, natural denominations are no intervention post-shock prices and observed post-shock prices. The former yields an equivalent-variation measure; the latter yields a compensating-variation measure, which has the drawback that the money-metric cardinalisation becomes policy-dependent (see Auerbach, 1985). Efficiency costs as a share of revenue are very similar across these denominations.

²⁹The average compensated price reduction needed to rationalise the suboptimal choices is 21%. The implied excess MPCE is $24/66 = 0.36$, consistent with the descriptive evidence in Table 4.2.

Costs of extra carbon emissions. The remaining 18.2% of efficiency costs reflect higher carbon emissions due to the subsidy and flypaper-induced increase in energy use. We measure the social cost of these additional emissions by multiplying each household’s policy-induced increase in energy consumption by α , the monetised atmospheric externality per unit of energy consumption.³⁰

Table 6.2: *Efficiency costs*

	Total efficiency cost	Source of efficiency cost			Carbon externality
		Price signal	Labelling:		
			choice distortion	fiscal spillover	
Aggregate (£bn)	3.76	2.64	0.07	0.36	0.69
	[3.67, 3.83]	[2.56, 2.72]	[0.06, 0.08]	[0.34, 0.38]	[0.67, 0.70]
Contribution:		(70.28%)	(1.87%)	(9.64%)	(18.21%)
		[69.44, 71.12]	[1.64, 2.12]	[9.00, 10.27]	[17.98, 18.37]

Notes: The table reports the efficiency cost associated with the implemented UK policy (a 39% energy subsidy and a £66-per-month transfer) over October 2022–March 2023, and decomposes it into four mutually exclusive and exhaustive components. Row (1) reports aggregate numbers and row (2) reports the percentage contribution from each source. See Appendix E.4 for decomposition details. 95% confidence bands are reported in square brackets.

Value of public funds

The policy response reduced aggregate household losses by £27.8 billion, raised the social cost of carbon emissions by £0.7 billion, and had a public revenue cost of £30.9 billion. The implied benefit to households, net of carbon externalities, per £1 of revenue spent, is $(27.8 - 0.7)/30.9 = 0.88$. This “value of public funds” is comparable to magnitudes reported in Hendren and Sprung-Keyser (2020) for programmes such as housing vouchers and welfare-to-work.

An important distinction from the *marginal* value of public funds (see Hendren and Sprung-Keyser, 2020; Hahn et al., 2024) is that our measure incorporates behavioural responses in households’ willingness to pay for a large policy package. This matters quantitatively in our setting. If, instead, we ignore behavioural responses—i.e., retain only the first-order (mechanical) budget effect—the implied value is 0.78. That calculation omits both the welfare gains from substitution toward energy at the subsidised price and the welfare losses from the flypaper effect among prepay households. Quantitatively, the former

³⁰We measure α using the average carbon emissions per £1 of residential energy spending in Q1 2023 and use a social cost of carbon of £59 per tonne (Department for Energy Security and Net Zero, 2023d); see Appendix E.3 for details. We then convert this cost to its value at pre-shock prices to align with our money-metric utility cardinalisation (Appendix E.4).

dominates the latter, so the first-order measure understates household benefits and, as a result, overstates policy-induced inefficiencies.

6.3 Policy Design

In this section, we embed our model of household behaviour in a social-welfare framework to quantify the policy trade-off between targeting assistance to households most affected by a price shock and minimising inefficiencies. We focus on an energy price shock of the same magnitude as 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 use different amounts of public funds. To make comparisons, we restrict attention to policies that entail the same level of public expenditure *inclusive of the monetised cost of carbon emissions*. This ensures that our comparisons internalise emission damages in the resource cost, while keeping distributional considerations focused on contemporaneous household losses driven by energy prices and transfers.

Denote by $\mathbb{P}^O = (s^O, t^O, L^O)$ the policy response implemented by the UK during the European Energy Crisis; by $e_i(\mathbb{P}^O) = \sum_{\tau=\underline{\tau}}^{\bar{\tau}} e_{i\tau}(\mathbb{P}^O)$ the energy consumption of household i over the energy crisis period; by $e_i(\emptyset)$ their energy consumption in the absence of policy intervention; and by $x_i^e(\mathbb{P}^O) = \sum_{\tau=\underline{\tau}}^{\bar{\tau}} P_{r(i)\tau}^e e_{i\tau}(\mathbb{P}^O)$ their subsidy-exclusive energy spending over the crisis. The total public resources devoted to the policy response are given by:

$$\bar{R} \equiv s^O \sum_{i=1}^N x_i^e(\mathbb{P}^O) + N \times 6t^O + \alpha \sum_{i=1}^N \left(e_i(\mathbb{P}^O) - e_i(\emptyset) \right), \quad (6.1)$$

where N is the number of UK households and, as we focus on the six-month period between October 2022 and March 2023 when a subsidy and transfer were in place, the monthly transfer t^O is pre-multiplied by six. The parameter α converts the policy-induced additional energy consumption into the social cost of the resulting carbon emissions.³¹

We consider the policy menu $(s, t, L = 1)$, which nests the policy implemented by the UK during the European Energy Crisis. This entails a universal, labelled transfer, so we refer to

³¹More 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 imply a negative value for α . As the suppliers in the UK market use two-part tariffs, which in principle facilitate fixed-cost recovery without this additional distortion to marginal prices, we assume the subsidy-exclusive energy price coincides with suppliers' private marginal cost.

this as *labelled (universal)*. We also consider alternative policies that expend no more (carbon externality-inclusive) public resources than \bar{R} , focusing on four counterfactual policy menus:

1. *Unlabelled (universal)*, $(s, t, L = 0)$ – a subsidy for energy consumption and a universal transfer that is not labelled.
2. *Proportional to inverse income*, $(s, (t/Y_i), L = 0)$ – a subsidy and transfer that is proportional to inverse household income (measured based on the preceding tax year).
3. *Proportional to past usage*, $(s, (t \times E_i), L = 0)$ – a subsidy and transfer that is proportional to a household’s monthly energy usage averaged over October 2021–March 2022 (the same six calendar months one year earlier).
4. *Proportional to past usage over income*, $(s, (t \times E_i/Y_i), L = 0)$ – 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.³²

Social welfare

To compare the impact of alternative policies on the distribution of household losses, we use a social loss function that aggregates household-level losses into a single summary measure. We adopt a criterion under which, if a social planner could directly allocate losses across households, it would choose to equalise *proportional losses*, $l_i^y(\mathbb{P})$, across all households. In practice, however, the planner is restricted to a limited set of instruments and faces a trade-off between using these instruments to reduce dispersion in proportional losses and the efficiency costs induced by the policy instruments. We model the planner as choosing policy to minimise a convex transformation of households’ proportional losses:

$$\mathcal{W}(\mathbb{P}) = \frac{1}{N} \sum_{i=1}^N \frac{1}{\psi} (\exp(\psi \times l_i^y(\mathbb{P})) - 1); \quad \psi > 0. \quad (6.2)$$

\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, as is standard in the optimal tax literature (e.g., Saez, 2002; Allcott et al., 2019). Aversion to loss

³²For 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 within income groups is captured by the convex transformation of losses, which allows the planner to place 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 governed by the parameter ψ . As $\psi \rightarrow 0$, the planner becomes indifferent to inequality in proportional losses, and minimising equation (6.2) is equivalent to maximising the sum of money-metric utilities with Pareto weights equal to inverse income. We set ψ so that observed policy (s^O, t^O) is optimal within the set of policies $(s, t, L = 1)$ that consist of a subsidy and a labelled universal transfer and that expend the same amount of public resources as observed policy (\bar{R}) . We first hold ψ at this value, before discussing the robustness of our results to varying ψ .³³

Let $\zeta^{\mathbb{P}}$ denote the equi-proportional loss: the common proportional loss which, if borne by every household, yields the same value of equation (6.2) as policy \mathbb{P} . It is a monotone transformation of $\mathcal{W}(\mathbb{P})$, given by $\zeta = \frac{1}{\psi} \log(\psi \mathcal{W} + 1)$, so it simply rescales \mathcal{W} into units interpretable as the corresponding equi-proportional loss. We report $\zeta^{\mathbb{P}}$ as our summary measure of social losses.³⁴

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

$$\zeta^{\mathbb{P}} = \underbrace{\frac{1}{\bar{Y}} \mathcal{L}^{LS}}_{\text{uncompensated losses}} + \underbrace{\frac{1}{\bar{Y}} (\bar{\mathcal{L}}^{\mathbb{P}} - \mathcal{L}^{LS})}_{\text{efficiency costs}} + \underbrace{\left(\zeta^{\mathbb{P}} - \frac{\bar{\mathcal{L}}^{\mathbb{P}}}{\bar{Y}} \right)}_{\text{targeting costs}}. \quad (6.3)$$

Equation (6.3) separates total social losses into (i) uncompensated losses that are unavoidable given the resource constraint, (ii) efficiency costs arising from policy-induced distortions, and (iii) targeting costs due to households experiencing different proportional losses.

\mathcal{L}^{LS} denotes average equivalent-variation losses under household-specific, non-labelled lump-sum transfers with revenue cost \bar{R} (set to equalise proportional losses across households). This term represents the portion of losses that cannot be compensated given the resource constraint (equation (6.1)), even under an idealised policy that induces zero efficiency costs and achieves the social loss-minimising distribution of household-level losses.³⁵

The second term measures the efficiency costs of policy \mathbb{P} —the difference between average household-level money-metric losses under policy \mathbb{P} and uncompensated losses.

³³In Appendix E.4, we plot the associated welfare weights. For households with losses of £43 per month (the average under the observed policy), the planner places 2.7 times more weight per £1 of loss on a household at the 25th income percentile compared to one at the 75th percentile. For a household with median income, the planner places 2.1 times more weight per £1 of loss at the 75th monetary-loss percentile compared to one at the 25th percentile.

³⁴By definition, ζ solves $\frac{1}{N} \sum_{i=1}^N \frac{1}{\psi} [\exp(\psi \zeta) - 1] = \mathcal{W}$, yielding $\zeta = \frac{1}{\psi} \log(1 + \psi \mathcal{W})$. Because this mapping is monotone, any policy \mathbb{P}^* that minimises \mathcal{W} also minimises ζ . Reporting the equi-proportional loss is in the spirit of Atkinson’s (1970) equally distributed equivalent income.

³⁵We set the *LS* benchmark to have revenue cost equal to \bar{R} , rather than a public resource cost (i.e., including the value of induced carbon emissions) equal to \bar{R} . This ensures that the full cost of carbon emissions induced by policy \mathbb{P} —including those arising from income effects—is captured in the efficiency-cost term.

These efficiency costs reflect the deadweight loss associated with the allocation of public resources for household support, rather than the efficiency costs of raising these funds, which remain constant across all policies we consider.³⁶ Dividing the first and second terms by average income, \bar{Y} , normalises them so that they are measured as fractions of income, ensuring they share the same units as the welfare-equivalent loss measure ζ^{IP} .

The final term captures targeting costs—social losses that arise when a policy fails to achieve equi-proportional losses. These are equal to the difference between the welfare-equivalent constant loss level, ζ^{IP} , and the average money-metric loss level scaled by average income.

The efficiency–targeting trade-off

Figure 6.2(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 *labelled (universal)* policy menu.³⁷ Moving leftwards along the line corresponds to a higher subsidy, which improves targeting by reducing the proportional losses of those most affected by price increases but comes with increased efficiency costs. The cross denotes the observed UK policy, which, by construction, is the social-loss-minimising policy (given our calibration of ψ). The shaded grey area delineates the combinations of targeting and efficiency costs that result in lower social losses than those under the observed policy. The social-loss-equivalent constant proportional loss under this policy is 1.99% of income. Panel (b) decomposes this into the contribution from uncompensated losses (0.57 percentage points), efficiency costs (0.59 percentage points) and targeting costs (0.83 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 (52.5%), the *unlabelled (universal)* policy menu dominates the *labelled (universal)* one. 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 *unlabelled (universal)* 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 17.7% lower efficiency costs.

³⁶This separation of the incidence and efficiency costs of how funds are spent from how they are raised echoes 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).

³⁷The δ parameters in the demand equation (5.2) 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 halve them). This means that, for households subject to the flypaper effect, we keep their excess marginal propensity to consume energy fixed as the transfer size varies.

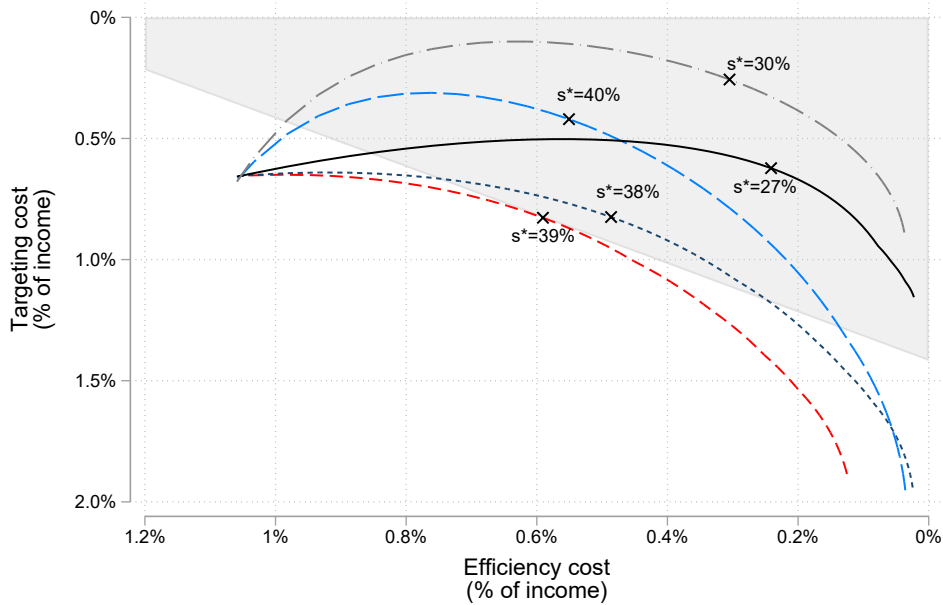
The lighter blue line shows the effect of specifying transfers to be inversely proportional to income. This policy menu significantly improves on those based on a universal transfer. The optimal policy under this menu results in total social losses of 1.54%, around four-fifths of those under the *unlabelled (universal)* menu. 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 compared to the observed UK policy. Under such a policy, many households still face large proportional losses, resulting in high targeting costs.

Another strategy is to base transfers on past energy use. The black and grey lines represent the policy menus when transfers are solely a function of past energy use and when they are a function of past energy in proportion to income, respectively. Both approaches perform substantially better than a policy that 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 1.13% of income. Of this, 0.57 percentage points 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 contribute much less strongly to increased carbon emissions, thereby 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, a policy that relies solely on transfers based on past usage with zero subsidy—represented by the rightmost points of the black and grey lines—outperforms the observed UK policy.

Figure 6.2: Counterfactual policy responses


(a) Efficiency–targeting trade-off



Subsidy combined with transfer that is:

- Labelled
- Unlabelled
- Prop. to $(1/Y)$
- Prop. to E
- Prop. to (E/Y)
- x Loss minimizing policy

(b) Social loss-minimising policy



Notes: Panel (a) reports the targeting and efficiency costs associated with the policy menus described in the text as $s \geq 0$ and $t \geq 0$ are varied holding public resources fixed. Moving from left to right along each line corresponds to lowering the subsidy rate and raising the transfer value. The cross marks the social loss-minimising policy, with the corresponding optimal subsidy rate indicated. Panel (b) decomposes total social losses under the loss-minimising policy for each menu into uncompensated losses (those remaining if public funds were allocated lump-sum to equalise proportional losses), efficiency costs and targeting costs. See Appendix E.4 for decomposition details.

Discussion

This section presents complementary analyses that clarify and extend our main results. We examine how the efficiency–targeting trade-off depends on household price responsiveness; to what extent income-based transfers can offset inequality in proportional losses; an alternative “dual” formulation of the planner’s problem that fixes social losses rather than resource costs; and how the quantitative results vary with social preferences. In each case, our main qualitative conclusions remain unchanged. We provide detail in Appendix E.5.

Price responsiveness. A key driver of the efficiency–targeting trade-off is how sensitive households are to energy price changes: stronger substitution responses both reduce welfare losses from the price shock and increase the efficiency costs of subsidising energy. To assess sensitivity of policy conclusions to household price responsiveness, we rescale the demand parameters governing price sensitivity, $(B, \{B_l\}_{l \in Z_2})$, to (i) half the aggregate demand response to the energy price rise and (ii) increase it by 50%, adjusting the constant term A so that model predictions absent the shock remain unchanged.

As expected, optimal subsidy rates decline with greater price responsiveness. For instance, under the universal transfer, the optimal subsidy is about one-third lower than in the baseline when demand is 50% more elastic, and roughly one-third higher when demand is 50% less elastic. Across all policy menus, there remains a clear efficiency–targeting trade-off generating a positive optimal subsidy rate, while better-targeted transfers consistently improve policy outcomes. Moreover, even with demand 50% less elastic than in our baseline, ignoring behavioural responses still substantially overestimates aggregate equivalent-variation losses from the crisis—by 52% under the less-elastic calibration, compared with 87% at our estimated parameters. These results confirm that our central qualitative conclusions are robust to substantial variation in demand elasticity.

Offsetting income-based transfers. Our framework isolates the welfare impact of inequality in the joint incidence of the shock and policy response in the targeting-costs term, which reflects heterogeneity in proportional losses both across and within income groups. In principle, inequality across income groups can be offset by income tax adjustments. To assess how much such adjustments could reduce targeting costs, we re-compute social losses under (i) the subsidy and unlabelled universal transfer, $(s, T, L = 0)$, and (ii) the subsidy and past-energy-use-based transfer, $(s, T \times E_i, L = 0)$, when combined with revenue-neutral, income-contingent transfers that equalise proportional losses across income percentiles.³⁸

³⁸These offsetting adjustments are similar in spirit to those proposed by Kaplow (2011), but adapted to (i) account for the joint incidence of the shock and policy, and (ii) accommodate heterogeneity within income groups. See Appendix E.5.

At the optimal policy (marked in Figure 6.2), income-based adjustments make essentially no difference to targeting costs under the universal transfer, since that policy already produces nearly flat average losses across income groups. When transfers are based on past energy use, income-based adjustments reduce targeting costs by about half, as they offset the higher average losses otherwise borne by lower-income households. Yet even under this flexible scheme, targeting costs remain substantial—exceeding efficiency costs—reflecting heterogeneity in exposure to shock conditional on income and past energy use.

Dual of social planner's problem. Figure 6.2 illustrates how social losses vary across policies holding total public resources constant at the level of the observed policy. Alternatively, one can ask how total public resource costs vary across different policy menus if social losses are held constant at the level under observed policy. In Appendix E.5, we show that this dual formulation yields the same ranking over policy options, and that the UK could have achieved the same level of social losses at 17.6% lower public resource cost under the menu *proportional to past usage over income* ($s, (t \times E_i/Y_i), L = 0$).

Social loss convexity. Our precise quantification of social losses depends on the social preference parameter ψ . Appendix E.5 shows how results vary with ψ , and demonstrates that, as long as the planner exhibits some 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 qualitative findings hold, including the ranking of policy menus, the presence of a positive optimal subsidy, and the relatively poor performance of income-based transfers alone.

Vertical equity concerns. Our baseline social loss function (equation (6.2)) incorporates vertical equity by scaling money-metric losses by inverse income. To explore stronger vertical equity concerns, Appendix E.5 introduces welfare weights that assign greater importance to a given proportional loss in inverse proportion to income.³⁹ Optimal policy continues to entail a substantial positive subsidy (of 25% if transfers are proportional to past energy usage and inverse income, compared with 30% under the baseline calibration), and tying transfers to income or past energy use still improves policy performance. The main quantitative change relative to our baseline is that stronger vertical equity concerns increase the welfare gains from linking transfers to income.

³⁹Under this specification, a given proportional loss for a household with £5,000 monthly income is assigned five times the weight of one with £1,000 income, and the first pound of loss receives twenty-five times the weight—compared with five times under our baseline. This efficiency-targeting decomposition continues to apply exactly. However, because this weighting places greater normative importance on the proportional losses of lower-income households, the targeting-costs term can become negative, reflecting that certain patterns of unequal losses (where poorer households lose less) are socially beneficial rather than costly.

Government's information set. The scope for better-targeted transfers depends on what information is observable and usable for policy design. Historic residential energy consumption data are available from suppliers, but delivery via suppliers can induce flypaper effects. Basing transfers on the interaction of past usage and income requires observing the joint distribution at the household level; in some settings, treating usage-based transfers as taxable income can proxy this data-linkage requirement. Our policy menus map to these different information regimes. Incorporating past usage into transfer design substantially improves targeting and reduces reliance on a subsidy. Where past-consumption data are unavailable, our results imply a larger role of a price subsidy in the optimal policy mix.

7 Conclusion

Sudden increases in the cost of living can cause significant hardship and strain both the wider economy and political stability. Effective policy therefore requires understanding not only average effects, but also how they vary across households, especially when relief must be deployed rapidly under limited information.

Our results show that policymakers can reduce reliance on price subsidies by using more targeted transfers that exploit information on households' circumstances. History-dependent policies, however, can create dynamic incentives—if used repeatedly, they may encourage higher consumption in normal times. To mitigate this, rapidly deployed transfers can be conditioned on harder-to-manipulate observables (e.g., dwelling characteristics) that are predictive of energy use. Nonetheless, the substantial variability in household welfare losses—both across and within income groups—means a temporary subsidy will often remain a valuable component of a relief package.

Reducing reliance on fossil fuels purchased on international commodity markets would have the dual benefits of lowering carbon emissions and reducing exposure to wholesale price volatility, along with the associated costs of relief packages.⁴⁰ Energy taxes will likely be needed to incentivise demand reductions and green investments. When such taxes 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 central to its political feasibility.

⁴⁰A separate geopolitical externality—namely, that higher European gas demand can translate into revenue for Russia—is not modelled here. Because retail-price subsidies weaken households' incentives to substitute away from energy, they raise aggregate demand and associated payments, pushing against the use of price subsidies in favour of targeted transfers. In our framework, this channel can be incorporated as an additional external cost—for example, by augmenting the social cost of carbon to reflect geopolitical considerations.

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.
- Atkinson, A. B. (1970). On the Measurement of Inequality. *Journal of Economic Theory* 2, 244–263.
- Auclert, A., H. Monnerie, 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*.
- Battle, 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.
- Bento, A. M., L. H. Goulder, M. R. Jacobsen, and R. H. Von Haefen (2009). Distributional and Efficiency Impacts of Increased US Gasoline Taxes. *American Economic Review* 99(3), 667–699.
- 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 Redistributive 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). Measuring cost of living inequality during an inflation surge. *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.
- 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, December). Quarterly domestic energy customer numbers. <https://www.gov.uk/government/statistical-data-sets/quarterly-domestic-energy-price-stastics>.

- Department for Energy Security and Net Zero (2023d). 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*.
- 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 Method 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.

- Hendren, N. and B. Sprung-Keyser (2020). A Unified Welfare Analysis of Government Policies. *Quarterly Journal of Economics* 135(3), 1209–1318.
- 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.
- 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. *Brookings Papers on Economic Activity* 1, 77–153.
- Kaplow, L. (2011). *The Theory of Taxation and Public Economics*. Princeton University Press.
- 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.
- 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.
- Met Office, D. Hollis, M. McCarthy, M. Kendon, T. Legg, and I. Simpson (2025). HadUK-Grid gridded and regional average climate observations for the UK. <https://catalogue.ceda.ac.uk/uuid/4dc8450d889a491ebb20e724debe2dfb/>.
- 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.
- Neary, J. and K. Roberts (1980). The theory of household behaviour under rationing. *European Economic Review* 13(1), 25–42.
- 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 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.
- 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>.
- 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.
- Ray-Chaudhuri, S., T. Waters, T. Wernham, and X. Xu (2023). Living standards, poverty and inequality in the UK: 2023. (R265).
- 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 Institutional Setting and Data

A.1 Price Cap and Energy Price Guarantee

Figure A.1 shows the energy price cap and Energy Price Guarantee for (i) prepay consumers and (ii) consumers paying by direct debit, averaging across regions. Levels differ slightly across payment types, but the series move together closely.

Although the cap is set separately by region, regional differences are small. In April 2022, for example, the average nominal cap increase was 54%, ranging from 50% in the South West to 57% in London. In our analysis, each household is assigned the cap for its region, identified from its postcode sector in the ExactOne data.

Figure A.1: *Energy price cap, energy price guarantee and cheapest available price plans for prepay consumers*



Notes: The figure reports the annual bill at Ofgem’s “typical” consumption values (12,000kWh of gas and 2,900kWh of electricity), in nominal terms. The “average of the cheapest X pricing plans” for $X=\{\text{direct debit, prepay}\}$ is a simple mean of the lowest-cost plans by each of the 10 cheapest suppliers (one plan per supplier), and includes fixed-price plans; only plans generally available to consumers are included. Small regional differences in the cap exist across Great Britain; the figure shows the average value.

A.2 Cost of Living Payments

Cost of living payments were paid directly into recipients' bank accounts in six instalments (with values ranging from £299 to £326) from summer 2022 to the spring 2024. Four payments fall within our analysis window:

- £326 paid between 14 and 31 July 2022
- £324 paid between 8 and 23 November 2022
- £301 paid between 25 April and 17 May 2023
- £300 paid between 31 October and 19 November 2023

Each payment was made within a relative short window (c. 20 days). Eligibility was determined by receipt of specified benefits or tax credits during a qualifying period preceding each payment window. The specific payment amounts and tight disbursement windows make the transfers straightforward to identify in the bank account data (following Ray-Chaudhuri et al. (2023)).

A.3 Bank Account Data

A.3.1 Sample Construction

Our main dataset is the *ExactOne Transactional Dataset* collected by ClearScore. ExactOne contains records extracted from individuals' bank and credit-card statements. For each transaction, we observe the date, amount, merchant/payee, and a description. Where available, we also observe the payment instrument (direct debit, card, or standing order). ExactOne assigns a transaction to one of 150 categories (e.g., "Energy (Gas, Elec, Other)", "Food, Groceries, Household", "Entertainment, TV, Media"). We focus on ClearScore users who are responsible for paying their households' energy bills, identified as those with at least one account that records energy payments.

In a first step, we classify energy transactions into (i) likely prepayment, (ii) direct debit, and (iii) other payment types (primarily standard-credit bills paid monthly or quarterly on receipt). We also identify receipt of the Energy Bill Support Scheme (EBSS) cash refunds as payments of £66 or £67 from an energy supplier in any months from October 2022 through March 2023.

We define "likely prepayment" as transactions that (i) are multiples of £5 and less than £100; (ii) are not coded as direct debit in ExactOne nor described as direct debit in the transaction narrative; and (iii) are not single, regular monthly payments made on a fixed calendar date. We validate this classification using two suppliers with known prepayment

focused business models. For Boost (a prepay-only supplier), 93% of transactions are classified as likely prepayment. For Utilita (which specialises in prepayment but does offer some standard credit tariffs), the corresponding share is 85%. For the other suppliers that serve both direct-debit and prepay customers, the share classified as likely prepayment ranges from 5% and 20%.

Direct debit transactions are identified using the bank-code field and/or the transaction description. Most customers who pay for energy by direct debit do so monthly; around 80% have a single monthly payment covering both gas and electricity. The remainder have separate direct debits for gas and electricity, sometimes with different suppliers. We define a continuous-payment spell as an interval during which at least one direct-debit payment occurs every 60 days (i.e., no gaps exceeding 60 days). Within such spells, we classify direct debits as variable if the payment amount changes in at least half of the months in the spell (equivalently, a change occurs on average at least every 2 months). This pattern is consistent with households on smart meters (automatic usage reporting) or those submitting regular meter readings, allowing suppliers to align monthly bills with actual consumption. We refer to the remaining direct debits as smoothed.

We retain spells of continuous energy payments lasting for at least six months with a single supplier and a single payment method. Continuity requires at least one payment every two months for prepay or direct-debit users; for other payment types (standard credit), at least one payment every four months (to allow for quarterly billing). We classify a spell as prepayment if, in a majority of its months, more than 50% of energy transactions are flagged as likely prepayment. Otherwise, spells are classified as direct debit or standard credit according to the household's observed payment methods.⁴¹

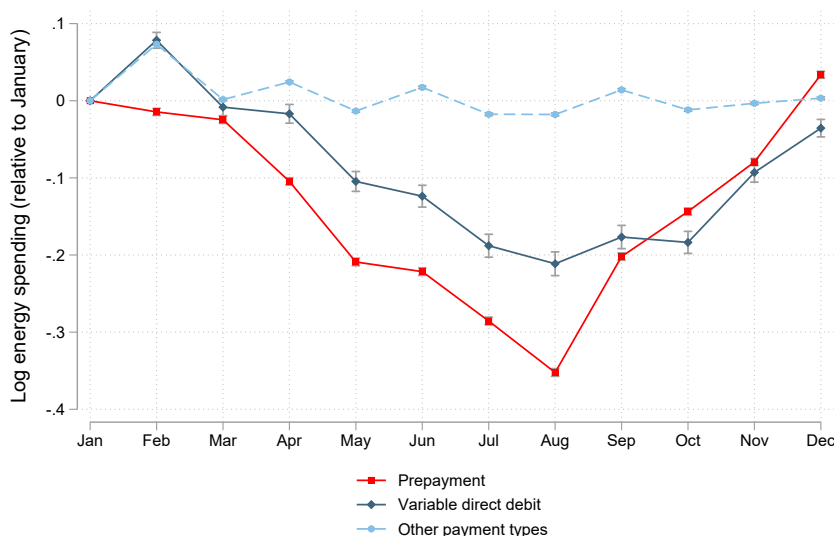
We collapse the data to the household-year-month level, excluding around 1% of households that either (i) have more than two energy suppliers or (ii) are served by a supplier that sells only heating oil. When a household uses two payment types with different suppliers within the same month (e.g., prepay with one and direct debit for another, or two separate direct debits) we retain information on both payment methods.

We classify a household-year-month as prepayment if all energy payments in that month are identified a prepay. We classify a household-year-month as variable direct debit if all energy payments in that month are direct debits and those debits are classified as variable. For each household-year-month we aggregate total energy spending across suppliers/accounts. Our analysis sample excludes household-year-months with mixed payment methods within the month.

⁴¹We classify a spell as direct debit if each month in the spell includes at least one direct-debit payment. Otherwise the spell is classified as standard credit. If some months include direct-debit payments and others do not, the spell is excluded from our sample.

The energy spending of households on variable billing exhibits much stronger seasonality than under smoothed billing. Figure A.2 plots average log deviations from the annual mean by calendar month in the pre-crisis period (2019–2020). For prepayment households and those on variable direct debits, seasonality is pronounced, with spending lower in summer and higher in winter. For all other payment types—most of which are smoothed direct debits—spending is approximately flat over the year. This pattern supports using spending on the variable-billing sample as a measure of energy usage.

Figure A.2: *Seasonality of energy spending by payment type*



Notes: The figure plots estimated deviations in log energy spending by calendar month relative to January, pooled over 2019–2020. Estimates come from regressions with household fixed effects; 95% confidence intervals are shown. The red line is for the prepayment sample, the dark blue line for the variable direct debit sample, and the light dashed line for all other payment types.

For each household–year–month, we construct measures of non-durable spending and income from all transactions recorded on the household’s linked ExactOne accounts. Non-durable spending includes expenditures on energy, groceries, vehicle fuel (gasoline), discretionary leisure (e.g., going out, entertainment), personal services (e.g., haircuts), phone and TV subscriptions, other household bills, and transport. We measure income as the sum of inflows to the account(s), excluding transfers from other accounts.

In a final step, we trim outliers at the 1st and 99th percentiles (by household–year–month) of energy budget share, non-durable spending, and income. We also require households to be present in the pre-crisis period (2019–2020), enabling construction of pre-crisis energy spending, and to appear for at least six months between June 2021 and December 2023 (our main estimation window).

A.3.2 Representativeness

Table A.1 reports the composition of payment types in our main sample. Prepayment households account for 25% of the sample—higher than the 15% share in the nationally representative Living Costs and Food Survey (LCFS). We reweight all results to match the spending share of prepay households.

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

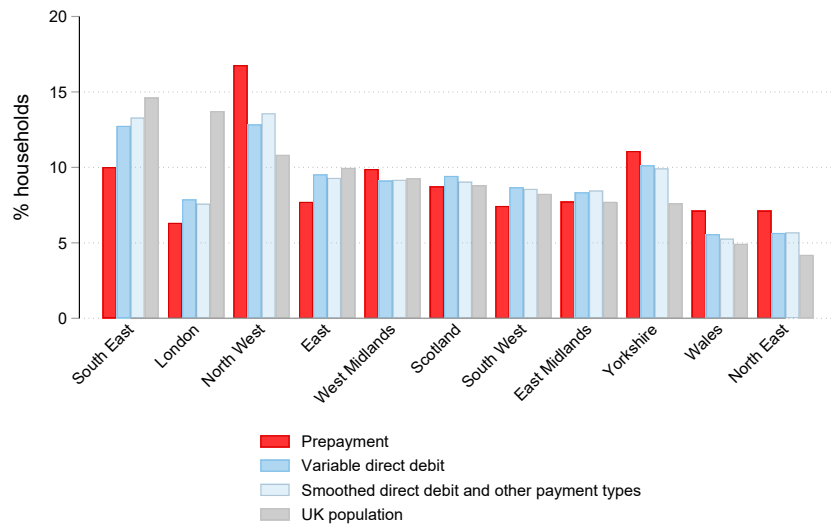
	Analysis sample?	Households		Mean energy spend (£)
		No.	%	
Prepayment	✓	70,181	24.8	107.8
Variable direct debit	✓	17,811	6.3	120.6
Smoothed direct debit & other payment modes	✗	195,283	68.9	119.7
Total		283,275	100.0	

Notes: The table reports the number and share of households in the full ExactOne sample (2019–23) by payment type, constructed as described in Appendix A.3.1. The final column reports mean monthly energy expenditure in 2019 by payment type. Households can switch payment type over time; our analysis sample comprises those paying by prepayment or variable direct debit.

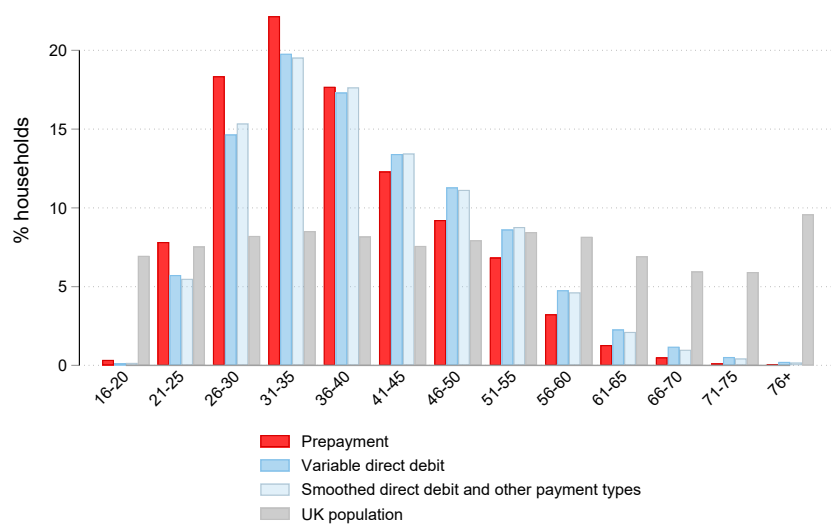
Figure A.3 compares the age and regional distributions of the ExactOne sample (by payment type) with those of the UK population (Office for National Statistics, 2021). Relative to the population, the ExactOne data over-represents younger individuals and residents of the North West, and under-represents residents of London and the South East. We construct weights by region and 5-year age bands and show our results are robust to reweighting the sample match the UK population.

Figure A.3: Age and geographic sample composition

(a) Region



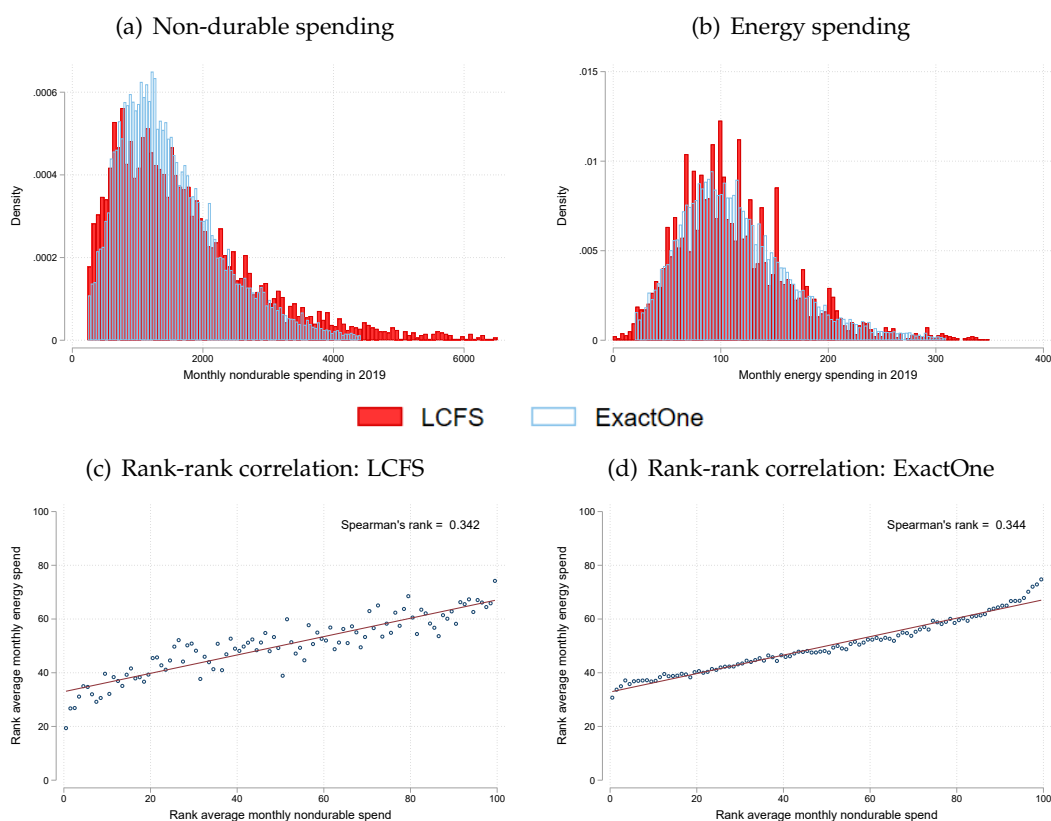
(b) Age



Notes: The top panel shows the share of households by payment type in the ExactOne sample and in the UK population by region. The bottom panel plots the corresponding shares by 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 is that the ExactOne data are substantially less noisy and exhibit little “bunching” at rounded monthly amounts, unlike the LCFS survey data.

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

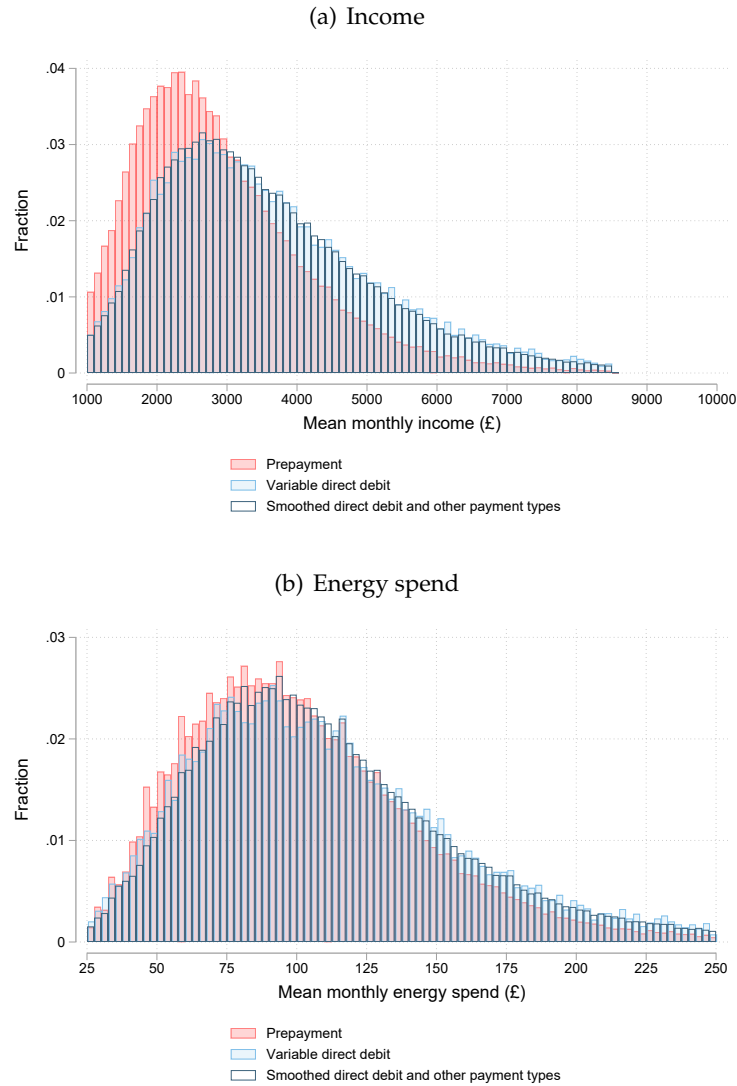


Notes: The top two panels compare the distributions of monthly non-durable spending (left) and energy spending (right) in ExactOne and the LCFS. Distributions are for 2019; for ExactOne we use each household's mean across the months the household is observed. The bottom panels show rank-rank correlations between households' average monthly non-durable and energy spending in the LCFS (left) and ExactOne (right). ExactOne observations are reweighted by age and region to match the UK population. All amounts are in 2022 prices.

Figure A.2 shows that variable-billing households exhibit pronounced seasonality in energy spending, consistent with spending closely tracking usage. For our analysis it is also important that these households are otherwise representative of the broader population. Within the variable-billing sample, we distinguish prepayment and variable-direct-debit households. Prepayment households have lower incomes and lower energy spending on average (Figure A.5) and are also younger and more likely to live in northern regions (Figure A.3), compared with direct-debit-households. Variable-direct-debit households look very similar, in observable, to those on smoothed direct debits and other payment types. Figure A.5 shows near-identical income and energy spending distributions, and Figure A.3 shows broadly comparable age and regional profiles. Taken together, these patterns suggest that variable-direct-debit households are broadly representative of the more common, smoothed-billing population. Providing we reweight the prepayment and direct-debit subsamples to match their total spending shares (as we do throughout), our sample is representative of

the wider UK population. Figure 2.2 shows consistency between energy spending in the ExactOne analysis sample and aggregate spending measured using the National Accounts over 2019–2023.

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



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

A.3.3 Estimation Periods

Table A.2 summarises the estimation windows used across the analysis, together with associated numbers of observations and households. Estimation windows begin from June 2021 onwards, when the energy price cap became binding for the vast majority of households. For estimation of price elasticities (Section 4.2) we use June 2021–September 2022, to avoid

confounding from transfers introduced in October 2022. To estimate the marginal propensity to consume energy (Section 4.3), we use June 2021–December 2023 (the entire period after the price cap became binding). For demand estimation, we use June 2021–June 2023 for model estimation, with out-of-sample validation using October–December 2023.

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

	No. households	No. obs
<i>Analysis sample over full crisis period</i>		
June 2021 – December 2023	74,428	1,261,414
<i>Price elasticities (Section 4.2)</i>		
June 2021 – September 2022	68,640	757,286
<i>Marginal propensity to consume energy (Section 4.3)</i>		
June 2021 – December 2023 (prepay only)	60,739	1,098,240
<i>Demand model (Section 5.3)</i>		
June 2021 – June 2023 (in-sample)	72,582	1,086,158
October 2023 – December 2023 (out-of-sample)	23,998	59,179

Notes: The table shows number of households and observations in each estimation window used in our analysis.

A.4 Weather Data

We use monthly minimum and maximum temperatures, humidity and total rainfall from the UK Met Office to control for seasonal and local weather changes affecting energy demand (Met Office et al., 2025). Weather station observations are interpolated to a 5km × 5km grid. We average these variable to Lower-layer Super Output Areas (LSOA; mean population ≈ 1,500) and merge them into ExactOne using each household’s residential LSOA.

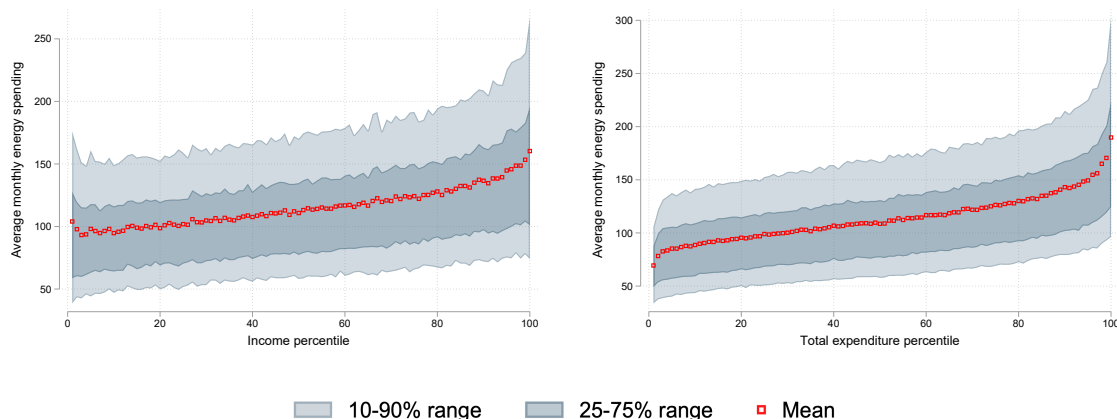
B Supplementary Details on Exposure to Price Shocks

B.1 Pre-Crisis Energy Spending

Figure B.1(a) replicates Figure 3.1(a) from the main paper using only 2019 data, thereby excluding the COVID-19 period. The two plots are very similar. Panel (b) summarises the joint distribution of energy spending and total non-durable spending; the relationship closely mirrors that between energy spending and income.

Figure B.1: *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, by income. The figure is constructed using data from 2019. Panel (b) summarises the distribution (mean and 10th, 25th, 75th, 90th percentiles) of households’ mean monthly energy spending, by mean monthly total non-durable spending (both measured over 2019–2020). Spending is expressed in 2022 prices.

B.2 Explanatory Power of Household Characteristics

Table B.1 shows R^2 from regressions of households’ average monthly spending by category—energy, groceries and vehicle fuel—on income, demographics, and housing characteristics. Higher R^2 indicates greater explanatory power of these observables for that spending category.

Table B.1: *Explanatory power of household income and demographics, by spending categories*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Energy				Groceries		Vehicle fuel	
R-squared	0.069	0.065	0.158	0.196	0.209	0.352	0.178	0.248
N	162,641	5,306	5,306	5,306	5,306	5,306	5,306	5,306
Dataset	ExactOne	LCFS	LCFS	LCFS	LCFS	LCFS	LCFS	LCFS
Income controls	✓	✓	✓	✓	✓	✓	✓	✓
Demographic controls	✗	✗	✓	✓	✗	✓	✗	✓
Housing controls	✗	✗	✗	✓	✗	✗	✗	✗

Notes: The table reports R^2 from OLS regressions of households’ average monthly spending by category—energy, groceries and vehicle fuel—on (i) income (percentile indicators); (ii) demographics (indicators for the number of adult males, adult females and children; age band and employment status of the head of the household; and region); and (iii) housing characteristics (number of rooms; property tax payments). Regressions use 2019 data; ExactOne in column (1) and the LCFS in columns (2)–(8).

B.3 Persistence of Energy Spending

We study persistence by estimating autocorrelation regressions of energy spending on its one year lag. Let x_{iyp}^e denote average energy spending by households i in year y and period $p = \{\text{summer, winter}\}$ (summer = April–September; winter = October–March). We estimate:

$$\log x_{iyp}^e = \rho_0 + \rho_1 x_{iy-1p}^e + \tau_y + \epsilon_{iyp} \quad (\text{B.1})$$

where τ_y are year effects and the lag matches the same period p one year earlier.

Table B.2 reports the $\hat{\rho}_1$ and partial R^2 from equation (B.1) estimated using two samples. Column (1) uses the pre-crisis period (2019–2020) and omits year effects; column (2) uses 2019–2023 and includes year effects. In both cases we estimate $\hat{\rho}_1 \approx 0.7$, with the lagged term explaining roughly half of the variation in current energy spending.

Table B.2: *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
N	82,864	422,276
Time period	2019-20	2019-23
Year effects	No	Yes

Notes: The table shows estimates of $\hat{\rho}_1$ and the partial R^2 from equation (B.1). The samples consist of households present for all six months in each time period.

C Supplementary Details on Prices, Elasticities and MPCs

C.1 Measuring the Energy Price and Quantities

C.1.1 Price Index


We construct a fixed-weight Laspeyres energy price index with base year 2019. Let w^g denote the mean 2019 energy expenditure share on gas in 2019. For household i in region $r(i)$ at time t , the index is:

$$p_{r(i)t} = w^g p_{r(i)t}^g + (1 - w^g) p_{r(i)t}^{el} \quad (\text{C.1})$$

where $p_{r(i)t}^g$ and $p_{r(i)t}^{el}$ are the regional unit prices of gas and electricity.

A potential concern with this index is substitution since gas prices rose more than electricity prices during the crisis. However, Figure C.1 shows that electricity's share of total energy is stable before and during the crisis, supporting the fixed-weight specification.

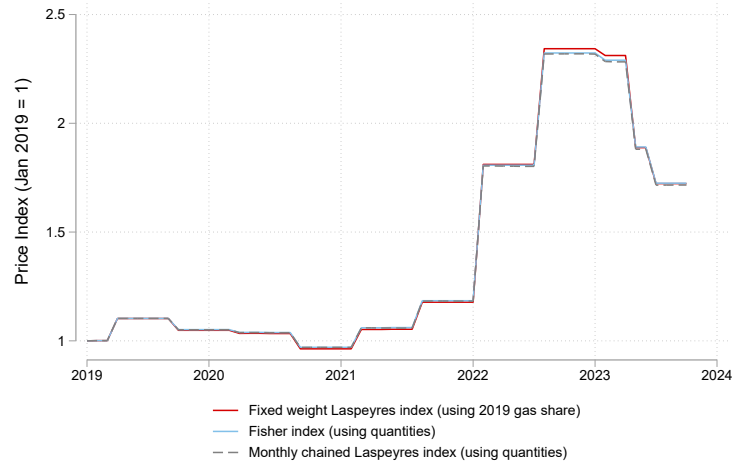
Figure C.1: *Quantity share of total energy from electricity*



Notes: The figure reports electricity quantity shares of total energy (gas + electricity) using data from Department for Energy Security and Net Zero (2024b), which is quarterly and seasonally adjusted. Quantity is measured in tonnes of oil equivalent.

Consequently, our fixed-weight Laspeyres index (using the 2019 annual-average gas expenditure share) closely tracks alternatives that allow for substitution. Figure C.2 shows that our baseline index is very similar to (i) a monthly chained Laspeyres index constructed with seasonally adjusted gas and electricity quantities as time-varying weights, and (ii) a Fisher index—the geometric average of a Laspeyres index (using January 2019 base quantity weights) and a Paasche index (using current-period quantity weights). The Fisher index permits substitution between electricity and gas, and as a superlative index, provides a second-order approximation to an arbitrary homothetic sub-cost-of-living index (Diewert, 1976).


Figure C.2: *Alternative price indexes*



Notes: Price indexes are calculated using data on residential gas and electricity consumption (quantities) from Department for Energy Security and Net Zero (2024b) or gas and electricity expenditure shares from the LCFS (2019). The “fixed weight Laspeyres index (using 2019 gas share)” is the index used for our main results.

Household-varying price indexes. To accommodate heterogeneity across households in their use of electricity and gas, we also construct a Laspeyres index (equation C.1) with expenditure shares that vary by income quintile and energy spending quintile (calculated using the 2019 LCFS). Figure C.3 shows that gas expenditure shares are broadly similar across quintiles (around 0.45), except for households in the bottom energy spending quintile, who allocate a smaller fraction of their energy spending to gas.

Figure C.3: *Share of energy spending on gas, by income quintile and energy spending quintile*



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

C.1.2 Spending and Quantities

Variable energy spending. Let $\tilde{x}_{i\tau}^e$ denote the total energy expenditure of household i in year-month τ . We obtain variable (usage-related) spending, $x_{i\tau}^e$ as

$$x_{i\tau}^e = \tilde{x}_{i\tau} - F_{i\tau}^g - F_{i\tau}^{el},$$

where $F_{i\tau}^g$ and $F_{i\tau}^{el}$ are the standing charges for gas and electricity, respectively.

Implied spending from energy-support transfers. The UK government provided £400 in bill rebates under the Energy Bill Support Scheme (EBSS), paid in six monthly instalments of £66/£67 from October 2022 to March 2023. Direct-debit customers either received these as credit added to their account or as cash refunds each billing period. In the bank-account data, we identify cash refunds as incoming credits of £66/£67 in the EBSS months; approximately 20% of direct debit customers received refunds. Standard credit and prepayment customers with smart meters received automatic credits to their meters. Prepayment customers with traditional meters were sent vouchers redeemable through 30 June 2023; aggregate statistics suggest that over 90% of these vouchers were redeemed (Department for Energy Security and Net Zero, 2023b). Let $R_{i\tau} \in \{66, 67\}$ denote the EBSS amount in month τ . We define the transfer-inclusive measure of energy spending,

$$\tilde{\tilde{x}}_{i\tau}^e = \begin{cases} x_{i\tau}^e & \text{if customer received a cash refund} \\ x_{i\tau}^e + 0.9 \times R_{i\tau} & \text{if customer is on prepayment} \\ x_{i\tau}^e + R_{i\tau} & \text{otherwise} \end{cases}$$

where the 0.9 factor for prepayment reflects the observed voucher redemption rate.

Quantity measurement. Quantity, $q_{i\tau}^e$ is then calculated as follows:

$$q_{i\tau}^e = \frac{\tilde{\tilde{x}}_{i\tau}^e}{p_{r(i)\tau}}$$

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

C.2 Price Elasticities


C.2.1 Northern Ireland Price Variation

A potential concern with our baseline specification is that seasonal and weather variables may not fully absorb time-varying factors that both correlate with changes in the regulatory price cap and affect energy demand. As a robustness check, we augment the sample with

households from Northern Ireland (NI), which was not subject to the Great Britain (GB) price cap.⁴² We exclude roughly half of NI households that purchase heating oil (a common residential heating fuel in NI but rare in GB), focusing on households using electricity and gas so as to be comparable to the GB sample.

Figure C.4(a) illustrates the differential price paths: NI prices rose much more gradually, with no discrete step increases in October 2021 and April 2022. Figure C.4(b) shows that in early 2021 the (not seasonally adjusted) spending series move in parallel across GB and NI; from April 2022, spending jumps in GB as the cap increases, while NI shows no corresponding break (indeed, spending falls with spring seasonality). Using this differential geographic price variation and including a full set of year–month dummies, Table C.1 shows an elasticity estimate very close to our main specification.

Figure C.4: *Energy prices and spending in Northern Ireland relative to the rest of the UK*



Notes: Panel (a) compares the average energy price in Northern Ireland (NI) versus the rest of the UK (GB) between January 2021 to September 2022. The NI series is a composite electricity-gas index from Consumer Council for Northern Ireland (2025). Panel (b) plots the change in mean monthly spending (not seasonally adjusted) for households in NI (excluding those observed purchasing heating oil) versus GB. All payment types are included (i.e., we do not restrict to the variable-billing sample).

C.2.2 Elasticity Robustness Checks


Table C.1 reports robustness of our elasticity estimates to alternative specifications. Column (1) uses $\log e$ (quantity) as the dependent variable. Columns (2)–(4) show estimates when we: (i) omit weather controls; (ii) use raw (non-deseasonalised) log spending; and (iii) replace the fifth-order polynomials in local minimum and maximum temperatures with second-order polynomials. Column (5) restricts identifying variation to the April 2022 price cap increase; because the estimation window narrows, we use second-order temperature polynomials (all other features match the baseline). Column (6) uses household-specific price indices

⁴²Northern Ireland has a separate regulator and participates in an all-island wholesale electricity market with the Republic of Ireland. Price increases are scrutinised as announced by suppliers rather than imposed via a market-wide cap, so adjustments are more gradual rather than discrete jumps.

with gas/electricity weights varying by pre-shock income and energy spending quintiles. Column (7) reweights observations by age and region to match the UK population. Column (8) include households paying by standard credit. Column (9) augments equation (4.1) with year-month fixed effects and identifies the price coefficients from differential geographic price movements between Great Britain and Northern Ireland. In all cases we obtain similar result to our baseline estimate.

Figure C.5 replicates Figure 4.2 from the main paper, except that the energy price index uses gas and electricity weights that vary by pre-shock income and energy spending quintiles (computed using the LCFS). The resulting elasticities across the 25 (5×5) groups exhibit the same qualitative patterns as in the baseline.

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



Notes: Each point reports the price elasticity for the April 2022 price increase, by income quintile and pre-shock energy spending quintile. Price responses are allowed to differ by payment type (DD, PP); we average across payment types using LCFS population shares that vary by income \times pre-shock spending cell. Pre-shock spending is measured over 2019–2020. The energy price index uses gas and electricity weights that vary by pre-shock income and energy spending quintiles (see Appendix C.1 for details).

C.3 MPCEs for Direct-Debit Households

In Section 4.3, we show that prepay households have a substantially higher marginal propensity to consume energy (MPCE) out of energy-support transfers (delivered from suppliers as meter top-ups or vouchers) than out of cash transfers paid directly from the government. Here, we conduct a similar analysis for direct-debit-households, and find some evidence their MPCEs vary by transfer type, but the differences are much smaller than for prepayment households.

Table C.1: Energy price elasticities, robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
% Δe over cap change	-14.4 (0.5)	-14.7 (0.4)	-15.6 (0.6)	-15.1 (0.5)	-14.2 (0.8)	-15.0 (0.5)	-16.5 (0.7)	-15.1 (0.5)	-14.4 (2.5)
Own-price elasticity	-0.317 (0.011)	-0.323 (0.008)	-0.342 (0.013)	-0.331 (0.010)	-0.312 (0.017)	-0.330 (0.011)	-0.363 (0.015)	-0.332 (0.010)	-0.317 (0.054)
N	757,286	757,286	757,286	757,286	757,286	757,286	757,286	899,540	824,155
Dependent variable	$\log e$	$\log x^e$	$\log x^e$	$\log x^e$	$\log x^e$	$\log x^e$	$\log x^e$	$\log x^e$	$\log x^e$
Deseasonalised	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	No
Polynomial in temp	5th	None	5th	2nd	2nd	5th	5th	5th	5th
Uses only April cap change	No	No	No	No	Yes	No	No	No	No
HH-specific index weights	No	No	No	No	No	Yes	No	No	No
Reweighting by age & region	No	No	No	No	No	No	Yes	No	No
Includes standard credit hh	No	No	No	No	No	No	No	Yes	No
Includes NI & yr-mn effects	No	No	No	No	No	No	No	No	Yes

Notes: Each column reports the percentage change in energy quantity e across the April 2022 cap change and the implied elasticity from equation (4.1). For each payment type d , let $\hat{\gamma}$ be the coefficient on $\log p$. Define $\hat{\gamma}' = \hat{\gamma} - 1$ when the dependent variable is $\log x$, and $\hat{\gamma}' = \hat{\gamma}$ when the dependent variable is $\log e$ (column (1)). We compute the percent change in quantity for the April 2022 price change as: $\Delta e_d / e_d = \exp(\hat{\gamma}'_d \Delta \log p) - 1$ where $\Delta p \equiv \log p_1 - \log p_0$. We also report the corresponding finite-change elasticity: $\epsilon_d = \frac{\Delta e_d / e_d}{(p_1 - p_0) / p_0} = \frac{\exp(\hat{\gamma}'_d \Delta \log p) - 1}{\exp(\Delta p) - 1}$. p_1 is the real price after the April 2022 increase and p_0 is the real price over Oct 2021–March 2022. We average $\Delta e_d / e_d$ s and ϵ_d s over payment type using LCFS population weights. All specifications, except column (2), include weather controls—fifth-order polynomials in local monthly minimum and maximum temperature (2nd-order in columns (4)–(5)), the squared difference between maximum and minimum temperature, rainfall and humidity—and indicators for the months immediately before and after a cap change. Standard errors (shown in parentheses) are clustered at the household level; standard errors for elasticities are computed via the delta method.

C.3.1 Transfers Received as Cash

Households paying their energy bills by direct debit received the £400 energy-support transfers either as credit on their bills or as cash refunds from their supplier. We begin with the cash-refund group. Panel (a) of Figure C.6 plots their log spending (solid line). For this figure, and the regressions that follow, we pool smoothed- and variable-direct-debit sample: once the variable-direct-debit sample is split into cash- versus credit-recipients, cell sizes shrink and estimates become noisier. The figure shows a steady increase in log energy spending over the crisis period among households receiving the transfer in cash.

The dashed line in Panel (a) of Figure C.6 plots predicted spending for October 2022–March 2023 based on price responses estimated outside this period. We construct this counterfactual exactly as for prepayment households (shown in Figure 4.3): regress deseasonalised log spending on log prices, weather controls, and a transfer-period indicator, then set the indicator to zero when plotting the transfer months. To allow direct-debit payments to reflect past or anticipated prices, we also include lags and leads of log price alongside the contemporaneous term. Actual and predicted spending series align closely, indicating little evidence of flypaper effects for this group.

To quantify this, we estimate a variant of equation (4.2) for households that received the energy-support transfer as cash, adding lags and leads of log price as controls. We allow for the possibility that there was higher consumption during the transfer window, but this only led to higher bills in the following quarter, by including a post-transfer indicator for April–June 2023. We then calculate MPCEs analogously to the prepayment case, and we also compare them with MPCEs out of cost-of-living payments. Column (1) of Table C.2 shows that the implied MPCE out of energy-support transfers is small and close to zero, providing no evidence of a flypaper effect for this group.


C.3.2 Transfers Received as Credit

We assess the impact of transfers credited to energy accounts by comparing households that received account credits with those that received cash refunds (Figure C.6(b)). For credit recipients, we construct transfer-inclusive spending by adding the monthly transfer amount to observed energy expenditure, yielding a measure comparable to cash recipient's spending. Prior to the transfers, the two groups' spending evolved similarly. With the October 2022 price rise and start of the transfers, transfer-inclusive spending for credit recipients rose noticeably relative to cash recipients; when the transfers ended in March 2023, transfer-inclusive spending for the credit group fell sharply. By late 2023, the gap between the two groups had returned to its pre-transfer level. This pattern is consistent with credit recipients


initially accumulating account surpluses and subsequently drawing them down after the transfer period.

Figure C.6: *Energy spending over the transfer period, direct-debit households*

(a) Direct-debit cash recipients



(b) Direct-debit cash versus credit recipients



Notes: The dashed vertical lines mark the start (Oct 2022) and end (Mar 2023) of the period during which energy-support transfers were delivered via suppliers. In panel (a): the solid line shows actual deseasonalised log spending; the dashed line shows counterfactual log spending predicted from a household fixed-effects panel regression, with weather controls, where the price coefficient is estimated using only variation outside the transfer window and the transfer indicator is set to zero for Oct 2022-Mar 2023. Panel (b) shows deseasonalised log spending inclusive of transfers for direct-debit households, comparing recipients of cash refunds with those that received account credit.

To estimate whether intertemporal smoothing fully explains the spending increase during the transfer window, we estimate a differences-in-differences specification, that compares

credit recipients (treatment group) with cash recipients (control group):

$$x_{i\tau}^e = \sum_{\tau} \text{time}_{\tau} \times \text{credit}_i + \sum_{\tau} \text{time}_{\tau} + \zeta_i + v_{i\tau}, \quad (\text{C.2})$$

where time_{τ} are year-month effects, and credit_i is an indicator for whether the transfers were paid as credit. The interaction terms $\sum_{\tau} \text{time}_{\tau} \times \text{credit}_i$ trace the spending difference between credit and cash recipients. Summing the interaction coefficients and dividing by the value of transfers value yields an estimate of the MPCE for those receiving transfers as credit (see column (2) of Table C.2). There is a small difference in cumulative spending between these two groups. For credit recipients, we estimate an MPCE from energy-support transfers that is 0.09 higher than for cash recipients. This is consistent with a modest flypaper effect, far smaller than the effect we estimate for prepayment households.

Table C.2: *MPCE estimates for transfers (households paying by direct debit)*

	(1)	(2)
Energy-support transfers (distributed via suppliers)	-0.04 (0.007)	0.09 (0.005)
Cost-of-living transfers	-0.02 (0.003)	
N	940,121	3,367,227
Sample	Cash only	Cash and credit

Notes: The dependent variable is deseasonalised log energy spending (column (1)) and deseasonalised energy spending (column (2)). Column (1) shows implied marginal MPCE out of energy-support transfers and cost-of-living payments for direct-debit households that received energy-support transfers as cash refunds. Column (2) shows the difference in the MPCE out of energy-support transfers, for direct-debit households, between credit and cash recipients. Standard errors are in parentheses, clustered at the household-level. Regressions include weather controls and household fixed effects. Column (1) controls for lags and leads of log real energy price.

D Supplementary Details on the Energy Demand Model

D.1 Empirical Demand Model

Let i index households, τ index year-months, and $r(i)$ denote the household's supply region. The price vector is $(p_{r(i)\tau}^e, p_{\tau}^n)$ where $p_{r(i)\tau}^e$ is the regional marginal residential energy price and p_{τ}^n is the price for other non-durables. Let $x_{i\tau}$ denote the household's total net budget (net of the standing charge—fixed fee—and inclusive of any energy-support transfers). We collect conditioning variables $\mathbf{z}_{i\tau}$ and parameters Ψ in $\theta_{i\tau} = (\mathbf{z}_{i\tau}, \Psi)$. $\mathbb{1}_{i\tau}$ is an indicator for whether a household is influenced by a flypaper effect; this equals 1 for prepay households during the energy-support transfer window.

Empirical counterpart of problem (5.1)

Consider the case when $\mathfrak{d} = 0$. This corresponds to privately optimal choices (and to observed choice behaviour for non-prepay households and for prepay households outside the energy-support transfer period). Let $u_{i\tau} = V(p_{r(i)\tau}^e, p_{\tau}^n, x_{i\tau}, \theta_{i\tau})$ denote indirect utility and let the expenditure function be $x_{i\tau} = \chi(p_{r(i)\tau}^e, p_{\tau}^n, u_{i\tau}, \theta_{i\tau}) \equiv V^{-1}(p_{r(i)\tau}^e, p_{\tau}^n, u_{i\tau}, \theta_{i\tau})$. We specify this as:

$$\begin{aligned} \log \chi = & \log u_{i\tau} + \log p_{\tau}^n + (A + \sum_{l \in \mathcal{Z}_1} A_l z_{i\tau l}) (\log p_{r(i)\tau}^e - \log p_{\tau}^n) + \frac{1}{2} (B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l}) \times (\log p_{r(i)\tau}^e - \log p_{\tau}^n)^2 \\ & + \left((C_1 + \sum_{l \in \mathcal{Z}_2} C_{1l} z_{i\tau l}) \log u_{i\tau} + C_2 (\log u_{i\tau})^2 \right) \times (\log p_{r(i)\tau}^e - \log p_{\tau}^n) \\ & + \frac{1}{2} D (\log p_{r(i)\tau}^e - \log p_{\tau}^n)^2 \times \log u_{i\tau} \end{aligned} \quad (\text{D.1})$$

Writing the log expenditure function directly in terms of the log price difference between the two goods ensures (in a two-good system) that the implied demands satisfy adding-up, homogeneity in prices and budgets and Slutsky symmetry.

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

$$\begin{aligned} \omega_{i\tau} = & (A + \sum_{l \in \mathcal{Z}_1} A_l z_{i\tau l}) + (B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l}) \times (\log p_{r(i)\tau}^e - \log p_{\tau}^n) + (C_1 + \sum_{l \in \mathcal{Z}_2} C_{1l} z_{i\tau l}) \log u_{i\tau} + \\ & C_2 (\log u_{i\tau})^2 + D (\log p_{r(i)\tau}^e - \log p_{\tau}^n) \times \log u_{i\tau} \end{aligned}$$

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

$$\omega_{i\tau} = (A + \sum_{l \in \mathcal{Z}_1} A_l z_{i\tau l}) + (B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l}) \times (\log p_{r(i)\tau}^e - \log p_{\tau}^n) + (C_1 + \sum_{l \in \mathcal{Z}_2} C_{1l} z_{i\tau l}) y_{i\tau} + \quad (\text{D.2})$$

$$\begin{aligned} & C_2 y_{i\tau}^2 + D (\log p_{r(i)\tau}^e - \log p_{\tau}^n) \times y_{i\tau} \\ y_{i\tau} = & \frac{\log x_{i\tau} - (\omega_{i\tau} \log p_{r(i)\tau}^e + (1 - \omega_{i\tau}) \log p_{\tau}^n) + \frac{1}{2} (B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l}) \times (\log p_{r(i)\tau}^e - \log p_{\tau}^n)^2}{1 - \frac{1}{2} D \times (\log p_{r(i)\tau}^e - \log p_{\tau}^n)^2}, \end{aligned} \quad (\text{D.3})$$

where $y_{i\tau} = \log u_{i\tau}$. The equations are *implicit* Marshallian budget-share demand and indirect utility; Together they define the (privately optimal) Marshallian budget share $\omega_{i\tau} \equiv \omega^0(p_{r(i)\tau}^e, p_{\tau}^n, x_{i\tau}, \mathbf{z}_{i\tau}; \Psi)$.

D.1.1 Empirical demands

Our demand specification (equations (5.2) and (5.3)) accommodates choices when $\mathfrak{d} = 1$, corresponding to a flypaper effect by adding $(\delta + \sum_{l \in \mathcal{Z}_3} \delta_l z_{i\tau l}) \mathfrak{d}_{i\tau}$ to equation (D.2). For completeness we repeat equations (5.2) and (5.3) here:

$$\begin{aligned} \omega_{i\tau} &= (A + \sum_{l \in \mathcal{Z}_1} A_l z_{i\tau l}) + (B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l}) \times (\log p_{r(i)\tau}^e - \log p_\tau^n) + (C_1 + \sum_{l \in \mathcal{Z}_2} C_{1l} z_{i\tau l}) y_{i\tau} + \\ &\quad C_2 y_{i\tau}^2 + D (\log p_{r(i)\tau}^e - \log p_\tau^n) \times y_{i\tau} + (\delta + \sum_{l \in \mathcal{Z}_3} \delta_l z_{i\tau l}) \mathfrak{d}_{i\tau} \\ y_{i\tau} &= \frac{\log x_{i\tau} - (\omega_{i\tau} \log p_{r(i)\tau}^e + (1 - \omega_{i\tau}) \log p_\tau^n) + \frac{1}{2} (B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l}) \times (\log p_{r(i)\tau}^e - \log p_\tau^n)^2}{1 - \frac{1}{2} D \times (\log p_{r(i)\tau}^e - \log p_\tau^n)^2}, \end{aligned}$$

Together these equations define the budget share energy demand $\omega_{i\tau} \equiv \omega(p_{r(i)\tau}^e, p_\tau^n, x_{i\tau}, \mathfrak{d}_{i\tau}, \mathbf{z}_{i\tau}; \Psi)$. We compute $\omega(\cdot)$ numerically by solving the two-equation fixed-point problem (5.2) and (5.3). A sufficient condition for a unique solution $(\omega_{i\tau}, y_{i\tau})$ at $(p_{r(i)\tau}^e, p_\tau^n, x_{i\tau}, \mathfrak{d}_{i\tau}, \mathbf{z}_{i\tau})$ is monotonicity of the expenditure function in utility (i.e., equation (D.5)); we verify this post estimation (the denominator term, $1 - \frac{1}{2} D (\log p_{r(i)\tau}^e - \log p_\tau^n)$ is bounded away from zero in our data).

D.1.2 Inequality constraints

Two additional regularity conditions implied by consumer theory, which we do not impose during estimation, are that the function $\chi(p_{r(i)\tau}^e, p_\tau^n, u_{i\tau}, \theta_{i\tau})$ is (i) concave in prices and (ii) increasing in utility. In our two-good set up, these translate into the following inequality restrictions:

$$\omega_{i\tau}^2 - \omega_{i\tau} + B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l} + D y_{i\tau} < 0 \quad (\text{D.4})$$

$$(\log p_{r(i)\tau}^e - \log p_\tau^n) \left((C_1 + \sum_{l \in \mathcal{Z}_2} C_{1l} z_{i\tau l}) + 2C_2 y_{i\tau} + \frac{1}{2} D (\log p_{r(i)\tau}^e - \log p_\tau^n) \right) > -1, \quad (\text{D.5})$$

which we check are satisfied post estimation.

D.2 Estimation

GMM. Let $w_{i\tau}$ be the observed energy budget share and $\mathbf{p}_{r(i)\tau} = (p_{r(i)\tau}^e, p_\tau^n)$ the price vector. We impose the population moments

$$\mathbb{E}[g_{i\tau}(\Psi)] = \mathbf{0}, \quad g_{i\tau}(\Psi) \equiv (w_{i\tau} - \omega(\mathbf{p}_{r(i)\tau}, x_{i\tau}, \mathfrak{d}_{i\tau}, \mathbf{z}_{i\tau}; \Psi)) \mathbf{h}_{i\tau}, \quad (\text{D.6})$$

where $\mathbf{h}_{i\tau}$ is the instrument vector. The estimator solves

$$\hat{\Psi} = \arg \min_{\Psi} \bar{\mathbf{g}}(\Psi)^\top W \bar{\mathbf{g}}(\Psi), \quad \bar{\mathbf{g}}(\Psi) = \frac{1}{N} \sum_{i,\tau} g_{i\tau}(\Psi),$$

with W positive semidefinite. We choose $|\mathbf{h}_{i\tau}| = |\Psi|$ (just-identified), so the estimator is equivalently the method of moments solution to $\mathbf{g}(\hat{\Psi}) = \mathbf{0}$, and the choice of W has no effect on $\hat{\Psi}$.

We use the following instrument vector:

$$\mathbf{h}_{i\tau} = \begin{pmatrix} 1 \\ \{\mathbf{z}_{i\tau}\}_{Z_1} \\ \Delta p_{i\tau} \\ \Delta p_{i\tau} \times \{\mathbf{z}_{i\tau}\}_{Z_2} \\ \widetilde{\{\log \text{inc}_{i\tau}\}_r} \\ \widetilde{\log \text{inc}_{i\tau} \times \Delta p_{i\tau}} \\ \mathbb{d}_{i\tau} \times \{\mathbf{z}_{i\tau}\}_{Z_3} \end{pmatrix}, \quad \Delta p_{i\tau} \equiv \log p_{r(i)\tau}^e - \log p_{i\tau}^n.$$

This instrument set excludes the household's total budget $x_{i\tau}$ and instead includes functions of (deflated) monthly income:

$$\widetilde{\log \text{inc}_{i\tau}} \equiv \log \text{inc}_{i\tau} - \left(\bar{w} \log p_{r(i)\tau}^e + (1 - \bar{w}) \log p_{i\tau}^n \right),$$

where \bar{w} is a fixed energy budget share used to form a Stone price index.

Starting values. We initialise GMM with 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 (5.2)) by IV given this guess, updating $y_{i\tau}$ (equation (5.3)), and continuing until convergence. For each iteration, when estimating equation (5.2), we instrument $\log y_{i\tau}$ using a generated instrument defined by equation (5.3), but with $\text{inc}_{i\tau}$ in place of $x_{i\tau}$.

Standard errors. We compute standard errors clustered by household, allowing for arbitrary within- i dependence. Let T_i denote the number of periods household i is observed and define the household-level average moments as $g_i(\hat{\Psi}) = \frac{1}{T_i} \sum_{\tau=1}^{T_i} g_{i\tau}(\hat{\Psi})$. Let $V = \frac{1}{N} \sum_i g_i(\hat{\Psi}) g_i(\hat{\Psi})'$ denote the variance of the average moment conditions and $D = \overline{\nabla g_{i\tau}}$ be the sample average of the gradient of the moments evaluated at estimated parameters. The asymptotic variance matrix of the estimates is then:

$$\text{Var}(\hat{\Psi}) = (D'WD)^{-1} (D'WVWD) (D'WD)^{-1},$$

which, with $|\mathbf{h}_{i\tau}| = |\Psi|$, simplifies to $\text{Var}(\hat{\Psi}) = D^{-1}V(D^{-1})'$. For statistics based on $\hat{\Psi}$ (e.g., efficiency costs), we construct Monte Carlo confidence intervals by taking 100 draws from the asymptotic variance matrix of the estimates, recomputing the statistic for each draw, and using the empirical quantiles.

D.3 Estimates and Model Fit

D.3.1 Parameter estimates

In Table D.1 we report the parameter estimates. Column (1) reports estimates of the baseline parameters; we model the energy demand Engel curve as a quadratic—adding higher-order terms does not materially alter its shape. The remaining columns report interaction effects for the budget-share intercept, price and first-order Engel curve term with a prepayment indicator and indicators for pre-shock energy spending deciles. We do not include interaction effects for the second-order Engel curve term or the price-Engel curve term; including these has negligible impact on our results. Consistent with the evidence in Section 4.3, we include a flypaper term for the energy-support transfers among prepayment households and allow it to vary across pre-shock energy spending deciles.

Table D.1: *Parameter estimates*


		×prepay	×pre-shock energy spending decile								
			2	3	4	5	6	7	8	9	10
Constant											
(A)	0.9898	-0.0899	0.0679	0.1051	0.1330	0.1548	0.1931	0.2119	0.2446	0.2860	0.2964
	(0.0504)	(0.0113)	(0.0122)	(0.0121)	(0.0124)	(0.0129)	(0.0133)	(0.0137)	(0.0148)	(0.0166)	(0.0209)
Price											
(B)	0.2240	-0.0274	-0.0033	-0.0027	-0.0046	-0.0031	-0.0017	-0.0038	-0.0017	-0.0016	-0.0060
	(0.0080)	(0.0012)	(0.0016)	(0.0016)	(0.0016)	(0.0016)	(0.0017)	(0.0016)	(0.0017)	(0.0019)	(0.0024)
Implicit utility											
(C ₁)	-0.2380	0.0113	-0.0090	-0.0135	-0.0168	-0.0192	-0.0241	-0.0260	-0.0298	-0.0344	-0.0341
	(0.0140)	(0.0016)	(0.0018)	(0.0018)	(0.0018)	(0.0018)	(0.0019)	(0.0019)	(0.0021)	(0.0023)	(0.0029)
(C ₂)	0.0149	-	-	-	-	-	-	-	-	-	-
	(0.0010)										
Price× Implicit utility											
(D)	-0.0237	-	-	-	-	-	-	-	-	-	-
	(0.0011)										
Flypaper effect											
(δ)	-	0.0172	-0.0014	-0.0014	-0.0006	-0.0017	-0.0015	-0.0012	-0.0019	-0.0031	-0.0056
		(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 (5.2) and (5.3)). 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, 5th-order polynomials in local monthly minimum and maximum temperature, the squared difference between maximum and minimum temperature and local monthly rainfall.

D.3.2 Elasticities and Engel curves

Figure D.1 plots model-based elasticities that summarise heterogeneity across pre-shock energy spending and pre-shock income quintiles, showing a similar pattern to elasticities estimated using responses to price cap adjustments (see Figure 4.2).

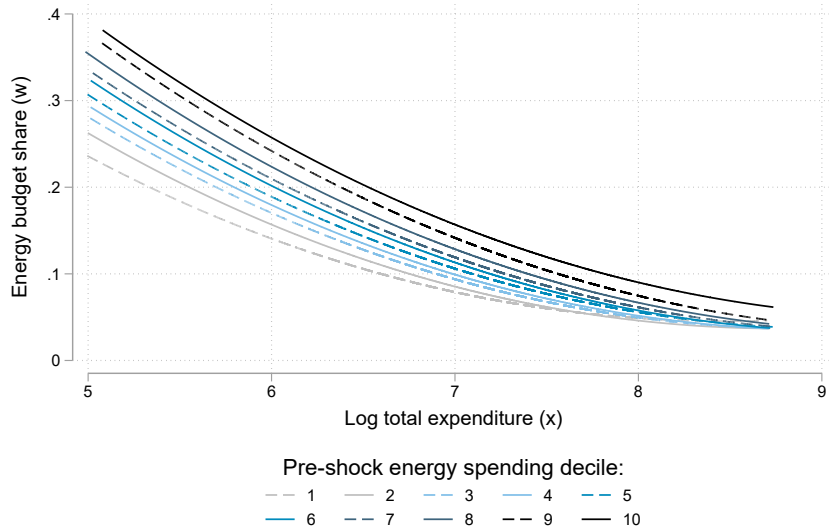
Figure D.1: *Heterogeneity in elasticities, model-based estimates*



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

In Figure D.2 we summarise heterogeneity in energy Engel curves. Each line show the relationship between the energy budget share, ω , and log total net expenditure ($\log x$). We draw a separate line for each pre-shock energy spending decile. We hold prices and weather controls at their mean level, and indicators for prepay and pre-shock energy spending share decile at their pre-shock within-decile means (computed over April–September 2022). The figure shows pronounced heterogeneity in Engel curves. The magnitude of this heterogeneity is similar to that reported in Lewbel and Pendakur (2017), who, using Canadian cross-section data, model preference heterogeneity via random Barten scales.

Figure D.2: *Heterogeneity in Engel curves*

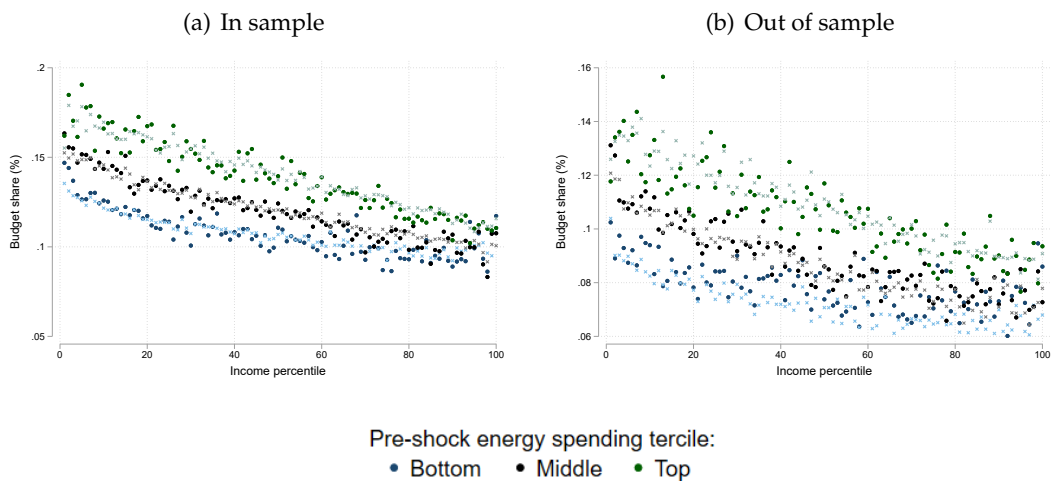


Notes: Model-implied Engel curves by decile of pre-shock energy spending. Prices/weather fixed at sample means; indicator controls set to within-decile means (April–September 2022).

D.3.3 Out-of-sample validation

Figure D.3 shows the joint variation in budget shares across income percentiles and pre-shock energy-use terciles, with panel (a) showing the in-sample (October–December 2022) fit and panel (b) the out-of-sample (October–December 2023) fit.

Figure D.3: *Winter energy demand jointly by income and pre-shock energy spending*

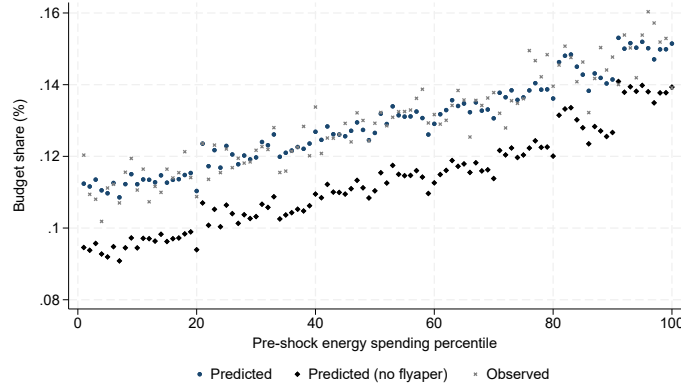


Notes: Panel (a) plots the average observed (crosses) and predicted (circles) budget shares across percentiles of the income distribution, separately by tercile of the pre-energy spending distribution, for October–December 2022 (in-sample). Panel (b) shows analogous information for October–December 2023 (out-of-sample).

D.3.4 Flypaper Effect Validation

To validate the flypaper effect, we use a hold-out sample during the energy-support transfer payment window. Specifically, we re-estimate the model after randomly holding out 25% of prepay observations over Oct 2022–Mar 2023. We then predict energy shares for the hold-out sample (i) with the flypaper parameters switched on, and (ii) with them switched off. Figure D.4 shows observed versus predicted energy budget shares by baseline (pre-sample) energy spending. Predictions that include the flypaper effect track the data closely, whereas turning it off leads to systematic under-prediction.

Figure D.4: *Hold-out sample validation of flypaper effect*



Notes: Figure plots average observed (crosses) and predicted budget shares (blue=flypaper effect on; black=flypaper effect off) for the hold-out prepayment sample across percentiles of the pre-shock energy spending distribution during Oct. 2022–Dec. 2023.

E Supplementary Details on the Welfare Framework

E.1 Computing Household-Level Welfare

Optimal choice. When a household's choice is privately optimal ($d = 0$), we obtain attained utility at the budget set $(p_{r(i)\tau}^e, p_{\tau}^n, x_{i\tau})$ by solving equations (5.2) and (5.3) for $(\omega_{i\tau}, y_{i\tau})$. Attained utility is given by $\exp(y_{i\tau})$.

Suboptimal choice. When a household's choice is subject to a flypaper effect ($d = 1$), we solve for the budget pivot $\phi_{i\tau}$ that rationalises the observed (suboptimal) choice as optimal (Proposition 1). Let $e_{i\tau}^1 \equiv \omega(p_{r(i)\tau}^e, p_{\tau}^n, x_{i\tau}, 1, \theta_{i\tau}) \frac{x_{i\tau}}{p_{r(i)\tau}^e}$ be the suboptimal energy quantity and $e_{i\tau}^0(p_{r(i)\tau}^e, p_{\tau}^n, x_{i\tau}, \theta_{i\tau}) \equiv \omega(p_{r(i)\tau}^e, p_{\tau}^n, x_{i\tau}, 0, \theta_{i\tau}) \frac{x_{i\tau}}{p_{r(i)\tau}^e}$ be the optimal Marshallian demand. We solve for the budget rotation $\phi_{i\tau} \in [0, 1)$ such that

$$e_{i\tau}^1 = e_{i\tau}^0(p_{r(i)\tau}^e(1 - \phi_{i\tau}), p_{\tau}^n, x_{i\tau} - \phi_{i\tau} p_{r(i)\tau}^e e_{i\tau}^1, \theta_{i\tau})$$

using the iterative algorithm:

$$\phi_{i\tau}^{(l)} = \phi_{i\tau}^{(l-1)} + \log e_{i\tau}^1 - \log e^0(p_{r(i)\tau}^e(1 - \phi_{i\tau}^{(i-1)}), p_{i\tau}^n, x_{i\tau} - \phi_{i\tau}^{(l-1)} p_{r(i)\tau}^e e_{i\tau}^1, \theta_{i\tau}),$$

stopping when $\|\phi_{i\tau}^{(i)} - \phi_{i\tau}^{(i-1)}\| < 10e^{-3}$. Once we have obtained $\hat{\phi}_{i\tau}$ we solve equations (5.2) and (5.3), with $d = 0$, at the virtual budget set $(p_{r(i)\tau}^e(1 - \hat{\phi}_{i\tau}), p_{i\tau}^n, x_{i\tau} - \hat{\phi}_{i\tau} e_{i\tau}^1)$ to obtain $y_{i\tau}$; attained utility as $\exp(y_{i\tau})$.

Money-metric utility. The preceding subsections outline the mapping from observed choices to attained utility,

$$\mathbb{V}(p_{r(i)\tau}^e, p_{i\tau}^n, x_{i\tau}, d, \theta_{i\tau}) = \begin{cases} V(p_{r(i)\tau}^e, p_{i\tau}^n, x_{i\tau}, \theta_{i\tau}) & \text{if } d_{i\tau} = 0 \\ V(p_{r(i)\tau}^e(1 - \phi_{i\tau}), p_{i\tau}^n, x_{i\tau} - \phi_{i\tau} p_{r(i)\tau}^e e_{i\tau}^1; \theta_{i\tau}) & \text{if } d_{i\tau} = 1. \end{cases}$$

We cardinalise utility in money-metric terms using the expenditure function (equation (D.1)): money-metric utility is the minimum expenditure at reference (pre-shock) prices $(p_{r(i)0}^e, p_0^n)$ required to reach the attained utility level,

$$\mathbb{V}^{MM}(p_{r(i)\tau}^e, p_{i\tau}^n, x_{i\tau}, d_{i\tau}, \theta_{i\tau}) = \chi(p_{r(i)0}^e, p_0^n, \mathbb{V}(p_{r(i)\tau}^e, p_{i\tau}^n, x_{i\tau}, d_{i\tau}, \theta_{i\tau}), \theta_{i\tau})$$

Note, if the household faces pre-shock marginal prices $(p_{r(i)0}^e, p_0^n)$, fixed fee $f_{r(i)0}^e$ and no transfer (so $d = 0$) then:

$$\begin{aligned} \mathbb{V}^{MM}(p_{r(i)0}^e, p_0^n, \tilde{x}_{i\tau} - f_{r(i)0}^e, 0, \theta_{i\tau}) &= \chi(p_{r(i)0}^e, p_0^n, \mathbb{V}(p_{r(i)0}^e, p_0^n, \tilde{x}_{i\tau} - f_{r(i)0}^e, 0, \theta_{i\tau}), \theta_{i\tau}) \\ &= \chi(p_{r(i)0}^e, p_0^n, V(p_{r(i)0}^e, p_0^n, \tilde{x}_{i\tau} - f_{r(i)0}^e, \theta_{i\tau}), \theta_{i\tau}) \\ &= \tilde{x}_{i\tau} - f_{r(i)0}^e \end{aligned}$$

Money-metric utility is the minimum expenditure at pre-shock prices $(p_{r(i)0}^e, p_0^n)$ required to attain the (possibly suboptimal) post-shock utility level. Denote a policy menu $\mathbb{P} = (s, t, L)$, where s is an energy-price subsidy, t a transfer to the household, and where $L \in \{0, 1\}$ indicates whether the transfer is labelled (so that $d_{i\tau} = L \cdot \mathbb{1}\{\text{prepay}_i\}$). Under \mathbb{P} the household faces the budget set $(p_{r(i)\tau}^e, p_{i\tau}^n, x_{i\tau}) = ((1 - s)P_{r(i)\tau}^e, p_{i\tau}^n, \tilde{x}_{i\tau} - f_{r(i)\tau}^e + t)$ For notational parsimony, we write $\mathbb{V}_{i\tau}^{MM}(\mathbb{P}) \equiv \mathbb{V}^{MM}(p_{r(i)\tau}^e, p_{i\tau}^n, x_{i\tau}, d_{i\tau}, \theta_{i\tau})$. The equivalent variation (money-metric) loss for household i in period τ from the price shock under \mathbb{P} is:

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

and the cumulative loss over the shock window $\tau \in \{\underline{\tau}, \dots, \bar{\tau}\}$ is $\mathcal{L}_i(\mathbb{P}) = \sum_{\tau=\underline{\tau}}^{\bar{\tau}} \mathcal{L}_{i\tau}(\mathbb{P})$. For some results we scale by income: letting Y_i denote average household monthly income over April 2021–March 2022 multiplied by the number of months in the shock window, we report $l_i^y(\mathbb{P}) = \frac{\mathcal{L}_i(\mathbb{P})}{Y_i}$.

E.2 Energy Poverty

We use an affordability-based 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.⁴³ 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 reliably observe housing costs for the whole sample in the ExactOne data, so we use the Living Costs and Food Survey (LCFS) 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 ratio on indicator variables for deciles of income, non-durable spending, and energy spending 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:

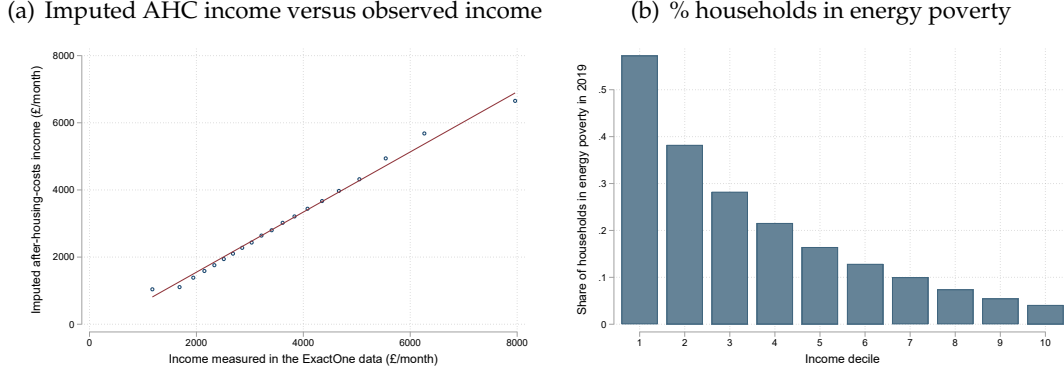
$$y_{it}^{AHC} = \tilde{y}_{d(i)}^{AHC} x_{it}.$$

Figure E.1(a) shows the relationship between imputed AHC income and the income measure in the ExactOne data. We define an energy-poverty indicator equal to one if energy spending, e_{it} as a share of after-housing-cost income, y_{it}^{AHC} , exceeds 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 E.1(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.

⁴³There 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 E.1: *After-housing costs income and energy poverty*



Notes: The left-hand panel shows the relationship between our imputed after-housing-costs (AHC) income measure, constructed as described in the text, and households' 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.

E.3 Carbon Emissions

We calculate carbon emissions per £1 of residential energy spending in the first quarter of 2023 as follows. We take total residential consumption of electricity and gas in that period from Tables 4.1 and 5.2 of Department for Energy Security and Net Zero (2024b) and multiply these quantities by the corresponding unit prices to obtain total electricity and gas spending. We then calculate the share of each £1 of energy spending accounted for by electricity and gas (46% is electricity and 54% is gas). Next, we obtain the number of kWh of electricity and gas per £1 spent. Multiplying these kWh-per-£1 figures by the carbon intensity of grid-average electricity and gas consumption yields CO2 emissions per £1 of electricity and gas spending. The carbon intensities are taken from Tables 9 and 10 of Department for Energy Security and Net Zero (2023a). Summing across the two fuels gives CO2 emissions per £1 of energy spending, yielding an estimate of 1.24kg of CO2 per £1. We then assume a social cost of carbon of £59 per tonne (Department for Energy Security and Net Zero, 2023d) to obtain the monetary cost of the carbon externality.

E.4 Computing Efficiency Costs and Social Welfare

E.4.1 Efficiency Costs

The observed government policy response ($\mathbb{P}^O = (s^O, t^O, L^O)$) expended public resources inclusive of carbon emissions given by \bar{R} defined by equation (6.1), which we repeat here:

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

where $e_i(\mathbb{P}^O) = \sum_{\tau=\underline{\tau}}^{\bar{\tau}} e_{i\tau}(\mathbb{P}^O)$ is household i 's total energy consumption over the crisis, $e_i(\emptyset)$ is their consumption in the absence of any government policy response, $x_i^e(\mathbb{P}^O) = \sum_{\tau=\underline{\tau}}^{\bar{\tau}} P_{r(i)\tau}^e e_{i\tau}(\mathbb{P}^O)$ 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 (6.1) as:

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

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 $\bar{\alpha} = \frac{\alpha}{6N} \sum_{i=1}^N (e_i(\mathbb{P}^O) - e_i(\emptyset))$ is the average monthly social cost 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. $\mathbb{V}_{i\tau}^{MM}(s^O, t^O, L^O)$: money-metric utility under observed policy
2. $\mathbb{V}_{i\tau}^{MM}(0, R_{i\tau} + \bar{\alpha}, 0)$: money-metric utility under an individual-specific transfer equal to $R_{i\tau} + \bar{\alpha}$, the monetary value of funds given to the household plus the average value of the carbon externality.
3. $\mathbb{V}_{i\tau}^{MM}(0, R_{i\tau}, 0)$: money-metric utility under a transfer equal to $R_{i\tau}$, the monetary value of funds given to the household
4. $\mathbb{V}_{i\tau}^{MM}(s^O, t^O, 0)$: money-metric utility under the observed subsidy rate and transfer, but under no transfer labelling
5. $\mathbb{V}_{i\tau}^{MM}(s^O, \tilde{t}, 0)$ 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) chosen to expend the same amount of government revenue as observed policy.

We measure aggregate efficiency costs of the implemented policy as

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

and decompose it according to:

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

E.4.2 Counterfactual Policies

We consider a series of counterfactual policies that all expend the same amount of public resources, inclusive of the social costs of carbon emissions, \bar{R} . 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 E_i by computing the Marshallian budget-share energy demand for each household i in each period τ in the six-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-specific marginal energy price, and averaging over the six 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 $\mathcal{W}(\mathbb{P}) = \frac{1}{N} \sum_{i=1}^N \frac{1}{\psi} (\exp(\psi \times l_i^Y(\mathbb{P})) - 1)$; $\psi > 0$. 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 money-metric (equivalent-variation) loss \mathcal{L} as the relative weight the planner places on each £1 of loss they suffer. This is given by their contribution to the social loss function scaled by their money-metric loss:

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

In Figure E.2 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 level sets for our calibrated value of ψ . Panel (b) shows them for the limiting case when $\psi \rightarrow 0$; 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. Comparing level sets across the two panels illustrates how the convexity of the social loss function alters the standard vertical welfare weights.

Figure E.2: *Level set of average welfare weights*



Notes: Figure shows combinations of income and money-metric 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 money-metric 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:

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

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

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

$$\zeta^{\mathbb{P}} = \underbrace{\frac{1}{\bar{Y}} \mathcal{L}^{LS}}_{\text{uncompensated losses}} + \underbrace{\frac{1}{\bar{Y}} (\bar{\mathcal{L}}^{\mathbb{P}} - \mathcal{L}^{LS})}_{\text{efficiency costs}} + \underbrace{\left(\zeta^{\mathbb{P}} - \frac{\bar{\mathcal{L}}^{\mathbb{P}}}{\bar{Y}} \right)}_{\text{targeting costs}}.$$

where \mathcal{L}^{LS} denotes average equivalent-variation losses under household-specific, unlabelled, lump-sum transfers with revenue cost \bar{R} that equalise proportional losses across households, $\bar{\mathcal{L}}^{\mathbb{P}}$ is average money-metric losses under the policy \mathbb{P} , and \bar{Y} is average income.

E.5 Sensitivity Analysis

E.5.1 Price Responsiveness

To illustrate how policy conclusions vary with households' responsiveness to energy prices, we rescale the parameters governing compensated price responses, $(B, \{B_l\}_{l \in \mathcal{Z}_2})$ by a constant factor, sc , and adjust the budget share constant term A by adding to it:

$$A'_{i\tau} = (1 - sc) \times \left(B + \sum_{l \in \mathcal{Z}_2} B_l z_{i\tau l} \right) \times \left(\log p_{r(i)0}^e - \log p_{\tau}^n \right).$$

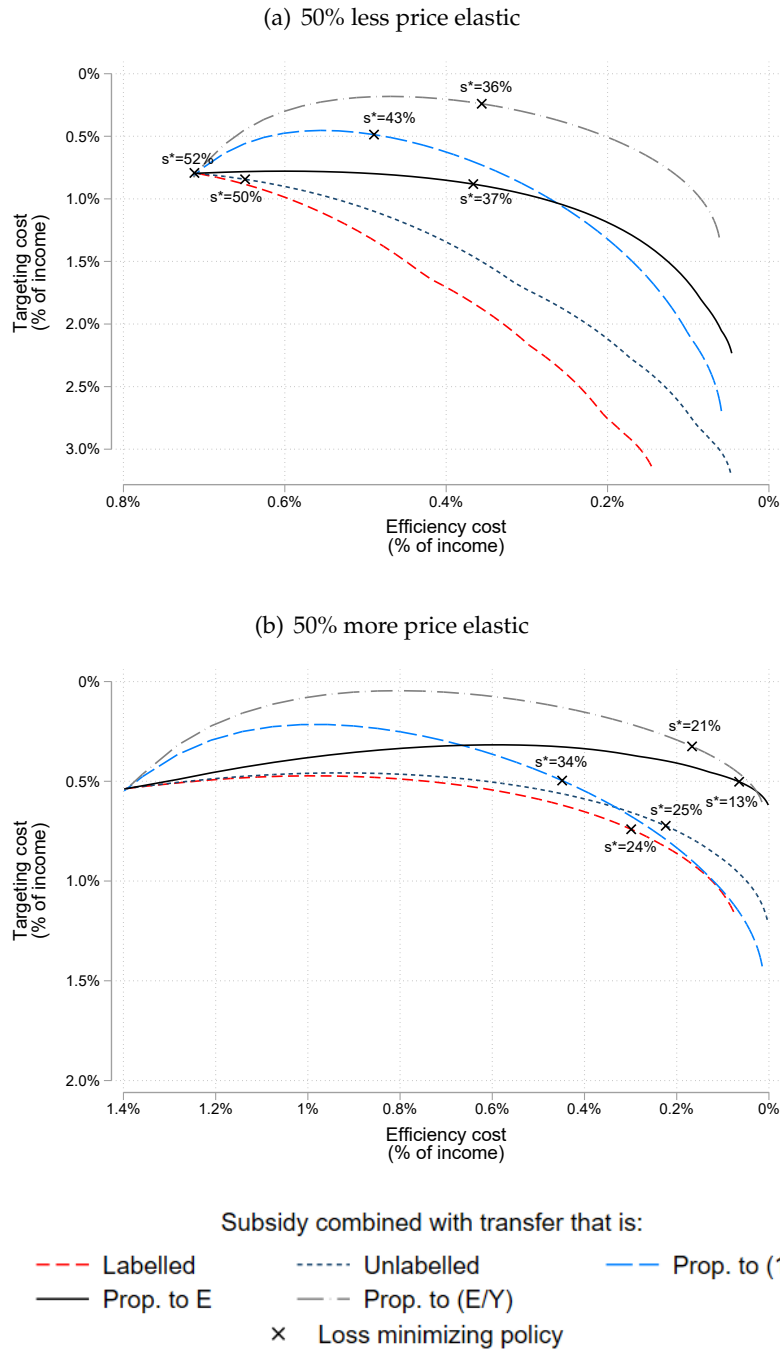
This adjustment changes the degree of price responsiveness while holding fixed the household-level demands at pre-shock energy prices.

We calibrate the scaling parameter such that the aggregate percentage reduction in energy demand resulting from the price rise from $p_{r(i)0}^e$ to $p_{r(i)\tau}^e$ —that is, from the no-shock to the observed (subsidy-inclusive) energy price—is 50% lower (sc_l) and higher (sc_h) than under our estimated parameters.

For each calibrated parameter set, we replicate the analysis underlying Figure 6.2, holding social preferences (specifically, ψ) fixed at their baseline value. We set the value of the public resource constraint, \bar{R} , to the implied cost of observed policy, \mathcal{P}^0 , under the calibrated preferences (which results in higher/lower total resources under less/more elastic demand).

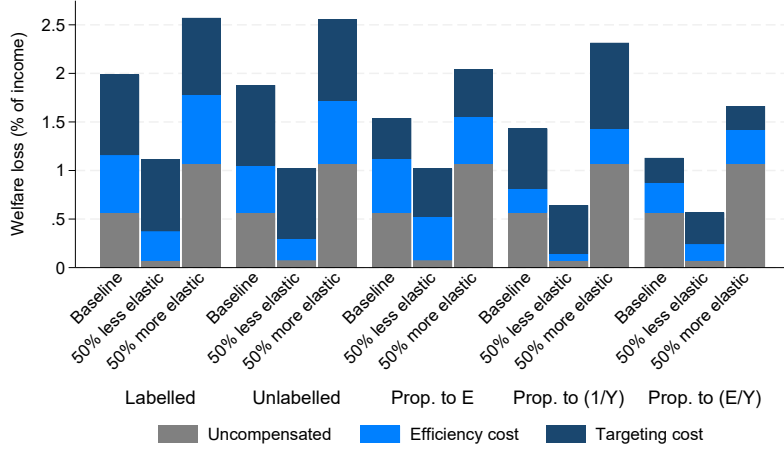
Figures E.3 and E.4 display the resulting efficiency–targeting trade-offs and decompositions. With less elastic demand, total social losses are larger, efficiency costs become relatively less important compared to targeting costs, and optimal subsidy rates rise compared with the baseline. With more elastic demand, the opposite holds: total losses fall, efficiency costs rise in relative importance, and the optimal subsidy rates decrease. These patterns follow the expected direction and confirm that, while quantitative magnitudes depend on demand elasticity, the qualitative efficiency–targeting trade-off—and hence the role of transfer design—remain robust.

Figure E.3: Efficiency–targeting trade-off, under more/less elastic energy demand



Notes: Figure 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 (a) corresponds to demand parameters calibrated to be 50% less elastic than at baseline estimates; panel (b) corresponds to parameters calibrated to be 50% more elastic than baseline.

Figure E.4: Social loss-minimizing policy, under more/less elastic energy demand



Notes: Figure decomposes total social losses under the loss-minimising policy for each menu into uncompensated losses, efficiency costs and targeting costs. It shows baseline results (repeating information from Figure 6.2(b)) and results when energy demand is calibrated to be 50% less and more elastic.

E.5.2 Offsetting Income-Based Transfers

The targeting costs reported in our counterfactual policy analysis (Figure 6.2) reflect heterogeneity in proportional losses both across and within income groups. Differences in average losses across income groups could, in principle, be offset by income tax adjustments.

To quantify the extent to which such adjustments reduce targeting costs, we re-compute social losses under (i) the subsidy and unlabelled universal transfer, $(s, T, L = 0)$, and (ii) the subsidy and past-energy-use-based transfer, $(s, T \times E_i, L = 0)$, when combined with revenue-neutral, income-contingent transfers that equalise average proportional losses across percentiles of the income distribution. With these transfers, all remaining targeting costs arise from within-income-group heterogeneity in losses. As we focus here on non-labelled transfers, we suppress the L term.

Specifically, consider the augmented policy $\mathbb{P}' = (s, T + \mathcal{T}^y)$ where T is the universal or past-energy-use-based transfer, and $\mathcal{T}^y = (\mathcal{T}_1^y, \dots, \mathcal{T}_G^y)$ are income-group-specific, revenue-neutral transfers. We set \mathcal{T}^y to satisfy:

$$\begin{aligned} \bar{l}^{\mathbb{P}'} \times \mathbb{E}_g[y] &= \mathbb{E}_g[\mathcal{L}(s, T + \mathcal{T}_g^y)] \quad \forall g \\ \sum_g \mathcal{T}_g^y &= 0, \end{aligned} \tag{E.1}$$

where expectations are taken over households within income group g ; and $\bar{l}^{\mathbb{P}'}$ denotes the common proportional loss level consistent with policy \mathbb{P}' . We jointly solve these equations for $(\mathcal{T}^y, \bar{l}^{\mathbb{P}'})$.

In words, we choose the income-based transfers that (i) equalise average proportional losses across income groups, and (ii) are revenue-neutral. These offsetting adjustments are similar to those proposed in Kaplow (2011), but adapted here to account for the joint incidence of the energy price shock and policy, as well as heterogeneity within income groups.

To clarify the connection, consider the simplified case with (i) no within-income group heterogeneity and (ii) no exogenous shock, so that comparisons are between pre- and post-policy outcomes (and $\bar{l} = 0$). Let (p_0, x_{g0}) and (p_1, x_{g1}) denote the pre- and post-policy prices and total budgets (net of fixed fees) for income group g . Recall money-metric losses are defined as $\mathcal{L}_g(p_1, x_{g1}, \mathcal{T}_g^y) \equiv x_{g0} - \chi_g(p_0, V(p_1, x_{g1} + \mathcal{T}_g^y))$. Then equation (E.1) can be written as:

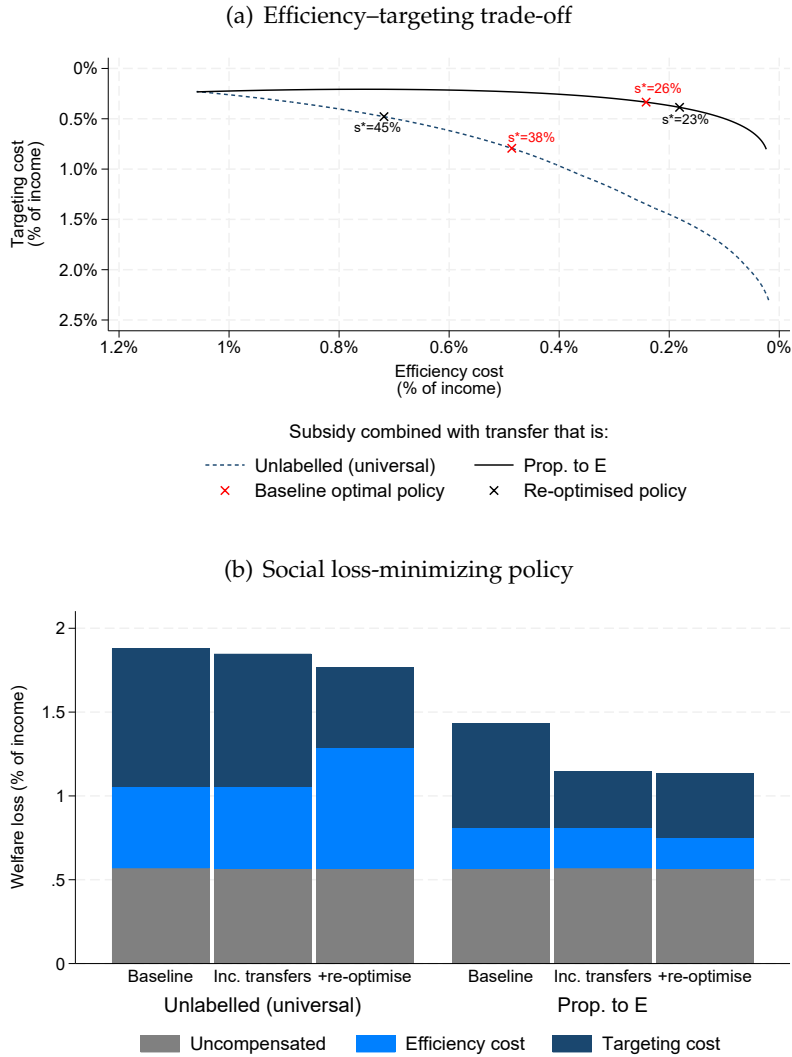
$$\begin{aligned} 0 &= x_{g0} - \chi_g(p_0, V(p_1, x_{g1} + \mathcal{T}_g^y)) \\ \rightarrow V_g(p_0, x_{g0}) &= V_g(p_1, x_{g1} + \mathcal{T}_g^y) \\ \rightarrow \mathcal{T}_g^y &= \chi_g(p_1, V_g(p_0, x_{g0})) - x_{g1}. \end{aligned}$$

Thus, as in Kaplow (2011), the offsetting income-based transfers correspond to the compensating variations required to restore each group's pre-policy welfare. Equation (E.1) extends this logic to a setting with uncompensated losses and within-group heterogeneity.

Figure E.5 summarises the impact of these offsetting transfers on the efficiency–targeting trade-off and overall social losses. We report results for policy menus (s, T) and $(s, T \times E_i)$. In panel (b), in each case, the first bars report results for the menu-specific optimal subsidy rate in the absence of income-based transfers (reproducing Figure 6.2(b)); the second bar holds that subsidy rate fixed while adding income-based transfers; and the third bar re-optimises the subsidy rate allowing for income-based transfers.

Under the baseline optimal policy, income-based adjustments have almost no effect on targeting costs in the universal transfer case, since that policy already produces nearly flat average losses across income groups. When transfers are based on past energy use, income-based adjustments reduce targeting costs by about half, as they offset the higher average losses otherwise borne by lower-income households. Re-optimising the subsidy rate after incorporating offsetting transfers produces only minor quantitative changes in total social losses.

Figure E.5: Counterfactual policy responses, allowing for offsetting income-based transfers



Notes: Panel (a) shows the efficiency–targeting trade-off under policies that combine a subsidy with either a universal transfer (s, T) or a past-energy-use-based transfer ($s, T \times E_i$), both with income-based offsetting adjustments that equalise average proportional losses across income groups. Panel (b) decomposes social losses at the menu-specific optimal policies into uncompensated losses, efficiency costs, and targeting costs. For each policy menu, the first bar reports results for the baseline optimal subsidy rate, the second adds income-based transfers while holding that subsidy rate fixed, and the third re-optimises the subsidy rate allowing for the income-based adjustments.

E.5.3 Dual of the Planner’s Problem

The results reported in the paper (Figure 6.2) 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 (6.2) subject to the constraint (6.1). 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 E.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 that result in the same social loss level as 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 17.6% (or £6bn) lower public cost. The social ranking of the policies in Table E.1 are the same as those implied by Figure 6.2.

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

	(1)	(2)	(3)	(4)	(5)
	% public resource cost savings under:				
	$(s^O, t^O, L = 1)$	$(s, t, L = 0)$	$(s, t/Y_i, L = 0)$	$(s, t \times E_i, L = 0)$	$(s, t \times E_i/Y_i, L = 0)$
$s=s^*$	£34.44bn	+1.5%	+8.3%	+7.6%	+17.6%
$s=0$		-11.6%	-20.1%	+4.8%	+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.

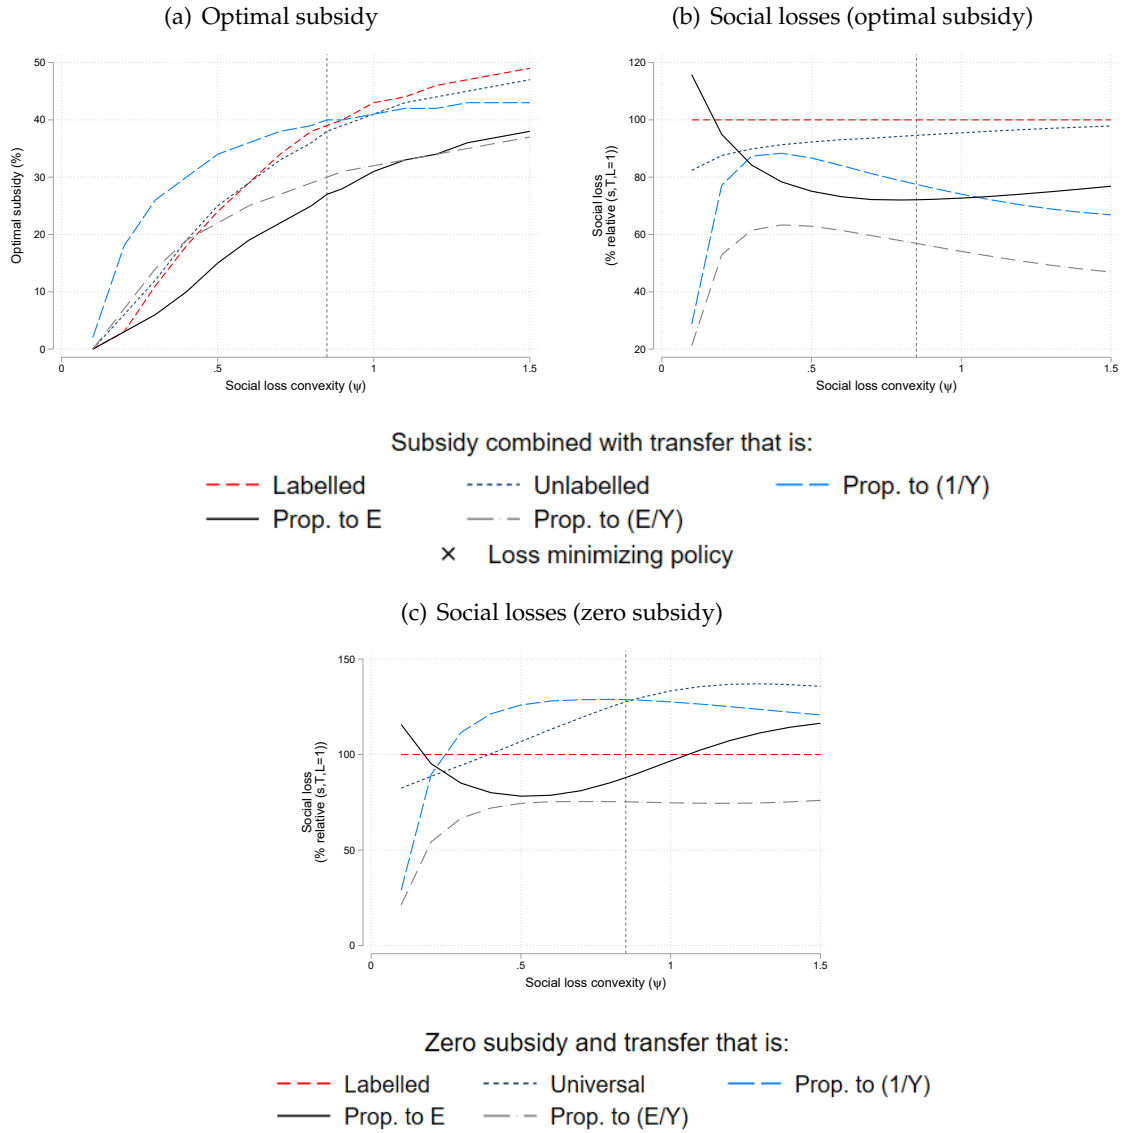
E.5.4 Social Loss Convexity

In Figure E.6(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 E.6(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 E.6: *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 .

E.5.5 Vertical Equity Concerns

Our baseline social loss function (equation (6.2)) incorporates vertical equity by scaling money-metric losses by inverse income. Here, we explore stronger vertical equity concerns by introducing welfare weights that assign greater importance to a given proportional loss

in inverse proportion to income. Specifically, we define the social loss function:

$$\mathcal{W}(\mathbb{P}) = \frac{1}{N} \sum_{i=1}^N \frac{1}{\psi} g_i \times (\exp(\psi \times l_i^y(\mathbb{P})) - 1), \quad (\text{E.2})$$

where


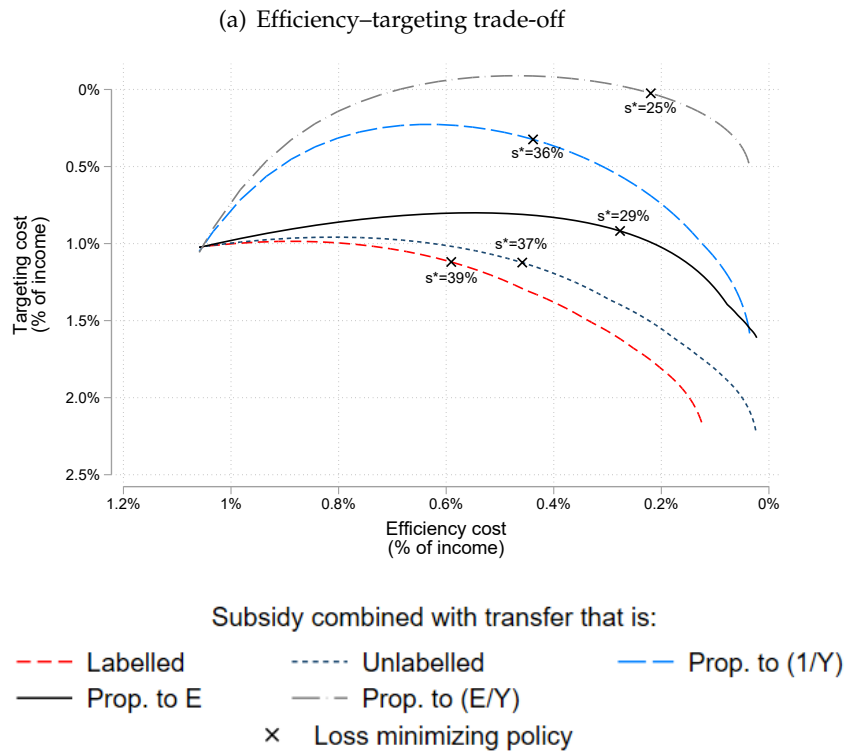
$$g_i \propto \frac{1}{Y_i}, \quad \mathbb{E}[g] = 1$$

and we hold ψ at its baseline value.

Under this specification, a given *proportional* loss for a household with monthly income of £5,000 is assigned five times the weight of one with £1,000 income, and the first pound of loss receives twenty-five times the weight—compared with five times under our baseline. The equi-proportional equivalent loss measure, $\xi^{\mathbb{P}}$, and efficiency–targeting decomposition continues to apply exactly. However, because this weighting places greater normative importance on the proportional losses of lower-income households, the targeting-costs term can become negative, reflecting that certain patterns of unequal losses (where poorer households lose less) are socially beneficial rather than costly.

Figure E.7 replicates Figure 6.2 under the modified social loss function (equation E.2). The results shows that an efficiency–targeting trade-off remains, that optimal policy continues to entail a substantial positive subsidy, and that tying transfers to income or past energy use still improves policy performance. The main quantitative change relative to our baseline is that stronger vertical equity concerns amplify the welfare gains from linking transfers to income.

Figure E.7: Counterfactual policy responses



Panel (a) reports the targeting and efficiency costs associated with the policy menus described in the text as $s \geq 0$ and $t \geq 0$ are varied, holding public resources fixed. The cross marks the social loss-minimising policy, with the corresponding optimal subsidy rate indicated. Panel (b) decomposes total social losses under the loss-minimising policy for each menu into uncompensated losses (those remaining if public funds were allocated lump-sum to equalise proportional losses), efficiency costs, and targeting costs. Results are evaluated using the modified social loss function in equation (E.2), which places greater weight on proportional losses of lower-income households.