Distributional effects of the European energy crisis

Peter Levell, Martin O'Connell and Kate Smith*

March 2024

Abstract

We estimate the distributional effects of the unprecedented rise in European energy prices over 2021-2023. Using bank account data for a sample of UK households we show that there were significant energy consumption falls, consistent with an average price elasticity of around -0.45, with proportionally larger responses for those with high pre-crisis spending. We also document evidence of a labeling effect associated energy bill rebates, part of the government's policy response. Using estimates of a flexible model of energy demand, we show that the introduction of a large energy price subsidy, along with bill rebates, limited welfare losses, though the labeling effect created an inefficiency that, if avoided, would have reduced average monetary losses by a further 33%. We show that the UK government's reliance on a relatively large subsidy can be rationalized by a social welfare function that places high weight on avoiding large monetary losses.

Keywords: Energy prices, price subsidies, cash transfers, labeling, targeting support, household demand
JEL classifications: D12, D31, D61, H23, H31, H53, Q41, Q48
Acknowledgements: 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 reference ES/M010147/1). All errors and omissions remained 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

By February 2022, when Russia invaded Ukraine, European wholesale gas spot prices were over seven times higher than a year earlier, feeding through to unprecedented increases in the price of natural gas and electricity in many countries. A number of governments responded with support packages aimed at protecting households from the consequences of these price rises, typically through a combination of price subsidies and transfers. In designing these packages policymakers face the challenge of targeting support to the neediest in a cost-effective way. Price subsidies can be effective at supporting those with large energy needs, but, depending on their design, they can also entail efficiency costs through incentivizing consumption, and they can channel resources disproportionately to those relatively well-off. Transfers avoid distorting consumption behavior, but may be less well targeted at those most in need of support, particularly if exposure to the price shock varies along dimensions that differ from those along which conventional redistribution is typically conducted (i.e., income, age or disability). In this paper we estimate the distributional impact of large rises in energy prices, and we study the design of support packages aimed at mitigating the welfare consequences of the shock.

We focus on the UK, one of the worst affected countries, where the cost of residential energy rose over four-fold between early 2021 and early 2023. We use individual-level panel data that includes monthly spending on energy over period of price rises. The data, collected by a fintech company, contain information of each households' bank account inflows and outflows across linked accounts, including credit and debit card transactions, standing orders and balances, for a period of at least three years. In addition to energy spending, we also observe salary and benefit payments and spending on other goods and services. In the UK the majority of residential electricity and gas bills entail a fixed access fee and single marginal usage price, both of which are subject to a regulatory price cap, which became binding for almost all suppliers during the energy price crisis. By focusing on a sample of households with variable bills (where spending tracks usage) and exploiting the binding regulatory cap, we are able to measure energy consumption for a sample of over 100,000 households.

We begin by documenting heterogeneity in exposure to energy price rises. Using pre-shock data we show that higher income households, on average, consume more energy but they allocate a smaller fraction of their total spending to energy. Therefore, while in monetary terms higher incomes households are more exposed than less well off households to price rises, exposure rises less than proportionately with total household spending. However, conditional on income, there is substantial variation in energy consumption, with some low income households being relatively heavy consumers of energy. In response to small price changes, these patterns of exposure are likely be a reasonable guide to the resulting welfare effect. However, UK households experienced a series of large incremental price increases, driven by the changes in the regulatory price cap, of 12% in October 2021, 54% in April 2022 and 27% in October 2022 (with the latter representing the net impact of a price rise and the introduction of a price subsidy), meaning consumption responses are likely to be non-negligible.

We use the ratcheting profile of prices, which result from the periodic increases in the binding regulatory price cap, to trace out energy spending and consumption responses. Each price increase is associated with an upwards jump in average energy spending, and a decline in consumption, which imply an average own price elasticity of energy demand of -0.48. We show that consumption responses are largest among the highest energy users – for instance, the average energy own price elasticity for those in the bottom quintile of the pre-shock energy spending distribution is -0.4, rising to -0.6 for those in the top quintile. This pattern holds within quintile of the income distribution, while, conditional on pre-shock energy spending, high income households tend to be less price sensitive.

The energy price cap increased by a 80% in October 2022, taking it to over four times its pre-shock level. At this point the UK introduced a price subsidy for energy which meant that households instead experienced a 27% price rise. Concurrently, it introduced a rebate on energy bills of £67 per month, which was in place until the end of March 2023. The projected combined cost of these measures was £33 billion (over 1% of UK GDP). We show that, on average, energy consumption exhibits only a small drop in October 2022, and when the energy rebates were withdrawn (with no change in prices) consumption fell by over 10%. This pattern implies that, on average, households allocated about 30% of their rebate to energy spending. We compare this with the marginal propensity to consume energy out of a different (conditional) transfer the government introduced over the same time period. Unlike the energy bill rebate, this second transfer was not specifically labeled as addressing rising energy costs, but rather was given to households in response to the wider "cost-of-living crisis". Households that received these cost-of-living payments, allocated around 5% of them, on average, to energy spending. The much higher marginal propensity to consume energy from the rebates is consistent with them influencing consumption through a labeling effect that is in addition to the income effect associated the monetary values of the rebates. This is similar to the labeling effect from "Winter Fuel Payment" – a labeled cash transfer provided annually to older households – document by Beatty et al. (2014).

We next estimate of model energy demand, which enables us to quantify the incidence of energy price increases and subsidies and bill rebates, and to model counterfactual policy responses. We specify an Exact Affine Stone Index (EASI) demand model to capture households' decisions over how to allocate their total spending between energy and alternative goods (Lewbel and Pendakur (2009)). We use the panel dimension of our data to allow energy price responses and Engel curves to vary by pre-shock energy spending. Our demand estimates replicate the descriptive patterns of how (Marshallian) price responses vary by pre-shock energy consumption and income. We decompose these price responses into a substitution effect (reflecting the increased relative price of energy) and an income effect (reflecting the impact of lower purchasing power). We find that the total expenditure elasticity of energy varies from 0.4, on average for households in the bottom half of the income distribution, to 0.8 for households in the top income quintile and bottom quintile of the pre-shock energy spending distribution. For all households energy is a normal good (a necessity), meaning that the substitution and income effects are reinforcing, with the former being quantitatively most important.

The labeling effect associated with the energy bill rebate means that, during the period that rebates were administered, observed choices are not necessarily utility-maximizing. We estimate how labeling impacts energy demands, and, show, under the assumption that households make optimal choices in the absence of labeling, how to measure utility associated with distorted choices. Our method entails computing the rotation of the budget constraint through observed choices that rationalizes them as optimal, and evaluating the true (normative) indirect utility function at the price and total expenditure values implied by the rotated budget constraint. This approach leverages the existence of an observed choice domain under which decisions are optimal (see Bernheim and Rangel (2009)) and avoids the need to specify a behavioral model for suboptimal decisions.

We use the model to compute the distribution of energy consumption changes and equivalent variations associated with the combination of the energy price shock and policy response. In the absence of government intervention, energy consumption by October 2022 would have been around 52% less than at pre-shock prices, and the associated welfare loss would have been £143 per month on average for the bottom quintile of the income distribution, rising to £223 for top quintile, primarily as a result of average energy consumption being positively correlated with incomes. The combination of the subsidy and bill rebates that the government introduced lowered the average fall in energy consumption to 35% and limited average welfare losses to £15 per month for the bottom income quintile, rising to £63 for the top quintile. However, the unconditional distribution of losses was dispersed with 5% of household experience monthly losses of at least £108 per month.

Administering a transfer as a labeled bill rebate led households to over consumer energy. This entails a direct cost to households associated with failing to choose utility maximizing choices. It also entails a fiscal externality due the encouragement of consumption of a subsidized good. We show that quantitatively this second channel is more important. Specifically, had households received a non-labeled transfer of the same monetary amount as the rebate, their average monthly welfare loss would have fallen by £2. However, if instead households had received an unlabeled transfer that expended the same resources as that of observed policy, their monetary loss would have fallen by £10 on average. Therefore the welfare loss from labeling primarily arises from the fiscal costs from over consumption, rather the from the direct utility costs.

In choosing the balance between subsidizing the price of energy and providing households with an unconditional transfer, the policymaker faces a trade-off. While raising the value of the subsidy helps mitigate losses of those particularly exposed to energy price shocks, on average, it benefits better-off households more than poorer households. The optimal balance between these instruments depends on the value the policymaker places on limiting large monetary losses and channeling resources to lower income households. To quantify this trade-off we specify a social welfare function with both vertical welfare weights (allowing for the policymaker to value limiting losses to the poor more than to the rich) and horizontal welfare weights (capturing aversion to large losses). Most income groups prefer, on average, a relatively low subsidy (and large transfer). Therefore a policymaker that is neutral to sizes of losses, conditional on income, would favor a low (or zero) energy subsidy. However, this leaves some households exposed to very large losses. The more the policymaker is concerned with limiting large losses, the larger is the optimal subsidy. In practice the UK government choice a subsidy equivalent to 38.5% of the price of energy; our framework rationalizes this through a high social value on limited large losses.

A recent set of papers study the European energy price crisis, including its likely distributional impacts based on pre-crisis data (Bachmann et al., 2022; Fetzer et al., 2023) and its implications for macroeconomic stabilization policy (Auclert et al., 2023; Dao et al., 2023). We contribute to this emerging literature by using micro spending data over the course of the crisis to estimate the distribution of energy consumption responses and welfare losses. In doing this we provide new estimate of energy price elasticities, that are based on large, salient and persistent price increases (see Labandeira et al. (2017) for a recent meta-study). Our work also relates to a literature that documents the distribution and efficiency cost implications of non-linear pricing in energy markets (Borenstein, 2012; Borenstein and Davis, 2012; Hahn and Metcalfe, 2021).

The rest of this paper is structured as follows. In Section 2 we describe the relevant features of the UK energy market and the government's policy response to the European energy price crisis, and we outline our main dataset. In Section 3 we present estimates of the distribution of energy consumption changes in response to price rises and the introduction energy bill rebates. In Section 4 we outline our empirical model of household energy choice. In Section 5 we present model estimates and describe the incidence of observed and counterfactual policy. A final section concludes.

2 Setting and data

2.1 Global energy price rises

Global demand for energy surged following the end of COVID-19 lockdowns and as political tensions with Russia, then the world's largest exporter of natural gas, gradually worsened. This led to significant increases in wholesale gas prices in Europe in the winter of 2021 and early 2022.

Figure 2.1 shows trends in wholesale gas prices in the UK market. Day-ahead prices for natural gas in the UK market spiked following Russia's invasion of Ukraine in February 2022 at 314 pence per therm, having risen seven-fold between March 2021 and March 2022. Prices eventually peaked in August 2022 at 356 pence per therm as Russia interrupted its pipeline exports to Europe, before it halted exports all together in September. Thereafter, spot prices fell back slightly but remained both high and volatile. These increases had a particularly large effect on energy prices in the UK, owing to the UK's relatively strong dependence on natural gas for both domestic heating and electricity generation. Retail energy prices remained high long after the decline in spot wholesale gas prices in late 2022, owing to the widespread use of forward price contracts by energy suppliers.

2.2 UK domestic energy market

UK domestic energy consumption consists of electricity, and gas for heating (a small minority of households rely on heating oil). Energy bills typically consist of unit charges for electricity and gas (charges per kilowatt-hour of use), and fixed standing charges that are independent of use. For the majority of households the fixed fee standing charge comprises only a small fraction of their overall bill – for instance betweeen April and September 2022, standing charges for gas and electricity accounted for 14% of a typical consumer's bill. Non-linear pricing, such as increasing block pricing where the marginal price increases in usage, are not a feature of the UK market. The majority of UK households (86% in 2012 Department for Energy and Climate Change (2013)) face a unit charge that is fixed throughout the day. The majority of the remaining households are on contracts that entail a different marginal price for day-time and night-time consumption.

Figure 2.1: *Wholesale gas prices in the United Kingdom, 2019-2023*

Notes: Data from Ofgem (2023). Figures show monthly average day-ahead wholesale gas prices in pence per therm. Russia invaded Ukraine on 24th February 2022.

Approximately 60% of households pay for energy via fixed direct debit, i.e., the payment of a bill, typically monthly, which is based on their expected energy use. Energy suppliers review direct debit payment amounts every 3-6 months, based on households' actual consumption measured from meter readings. Therefore households paying by fixed direct debit typically spend part of the year in credit (summer) or debit (winter).

The remaining 40% of households have bills that correspond much more closely to their actual energy usage each billing period. Some of these households are on "variable" direct debits, which adjust each period according to their energy consumption. This requires them to either have a smart meter, which automatically transmits usage to the supplier, or send in regular meter readings. Other households pay via standard credit, which means they only pay for energy once they have received a bill for their actual use. The final group have pre-payment meters, which require topping up (either online or in shops) before the credit can be used for energy consumption.[1] The amount of credit that can be held on such meters is limited (typically around £200-250 for electricity meters) meaning that prepayment customers must typically top up their meters regularly. 85% of those that recording spending on gas or electricity prepayment meters in the Living Costs and Food Survey report expect-

[1] Common reasons for using pre-payment include: a desire to more closely monitor and control their energy costs; because they are renters, and their landlord has arranged for them to use prepay (for example to avoid their tenants leaving the property with unpaid energy bills); or, because they are in arrears with their energy provider and a prepayment meter is being used to collect payments.

ing their payment to cover one month or less. In Appendix A we describe these different payment methods in more detail.

2.3 Government policies

Energy price controls

A distinctive feature of the UK energy market is that a price cap (officially known as the "default tariff cap") has been in place since January 2019. The cap is administered by the energy regulator Ofgem and sets maximum values for the unit prices and standing charges that suppliers are allowed to charge domestic consumers for electricity and gas (although it is often presented as a cap on the cost of a 'typical' energy consumer's bill based on average annual consumption of electricity and gas).

The stated aim of the cap is to limit firm profits and to prevent those households that do not shop around for cheaper tariffs from overpaying. The cap level is based on Ofgem's estimates of supplier costs. Until October 2022, Ofgem updated the cap every 6 months to reflect changes in suppliers' costs (most likely driven by wholesale price changes). After October 2022 Ofgem updated the cap every 3 months as a response to the increasingly volatile wholesale market. Ofgem decides the level of the cap in each period and announces it approximately one month in advance.

As wholesale energy prices rose rapidly, the price cap became increasingly binding on suppliers. Figure 2.2 shows the value of a bill at capped electricity and gas prices for Ofgem's definition of a typical consumer, alongside the average annual bills this consumer would pay under the 10 cheapest tariffs on the market. In January 2020, the average of the 10 cheapest tariffs was around 70% of the cost of a bill at the default tariff cap. By January 2021, this had risen to 90%, and from October 2021 it had risen to 99%. At this point, the vast majority of consumers were paying the maximum prices specified by the cap. The fact that the cap binds in this way is useful for us, as it means know with a high degree of certainty the tariffs consumers paid the energy they consumed from mid-2021 onwards, enabling us to obtain accurate measures of the quantities of energy they are consumed.[2]

[2] A potential complication in using the prices set by the Ofgem cap to infer prices paid, and calculate quantities consumed, is that the cap does not apply to tariffs on fixed price contracts that consumers may have agreed before the energy price cap increases occurred. Historically, these offered consumers lower prices than variable rate contracts. However, the price advantage of fixed contracts disappeared from mid-2021 onwards as energy prices rose, leading to a substantial reduction in their share of the market. In June 2021, 42% of gas customers and 40% of standard electricity consumers were on fixed term contracts. By December 2022, these figures had both fallen to 25% (Department for Business Energy and Industrial Strategy (2022)). Fixed tariffs are in any case much less prevalent among the set of consumers with variable energy consumption that we use for our main analysis. The share of pre-pay consumers on fixed tariffs is much lower than for direct debit consumers, at less than 1% for both electricity and gas tariffs (Department for Business Energy and Industrial Strategy, 2022).

Figure 2.2: *Energy price cap, energy price guarantee and cheapest available tariffs, 2019-2023*

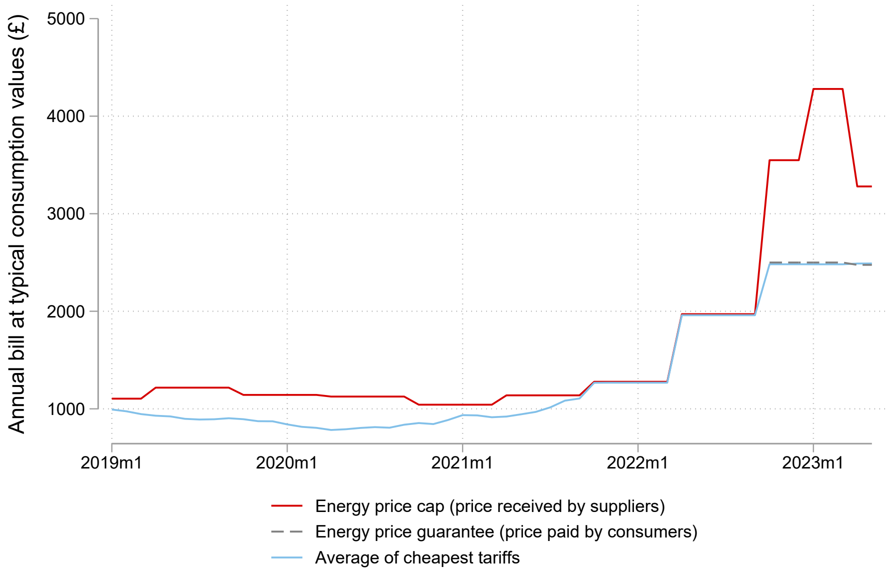

Notes: Data from Ofgem (2023). Figures are costs of an annual bill at 'typical' consumption values of gas and electricity (12,000kWh of gas and 2,900kWh of electricity) for dual fuel direct debit consumers. The average of the cheapest tariffs is a simple average of direct debit tariffs from the 10 suppliers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

In September 2022, as households faced the prospect of an 80% increase in the energy price cap in October 2022, the government introduced an "Energy Price Guarantee" (EPG) to limit further price increases. The EPG capped the unit rates and standing charges that suppliers could charge, and paid a subsidy to suppliers to make up for the difference between the EPG rates and the energy price cap set by Ofgem. It was set at a rate to cap the costs for a typical user at £2,500 from October 2022 to April 2023 (later extended until June).

The dashed line in Figure 2.2 shows the cost of the typical bill under the EPG. The average cost of an energy bill for those paying via direct debit at typical consumption values rose by 54% in April 2022. It then rose a further 27% in October; in the absence of the EPG, it would have risen by 80%. The EPG also protected consumers from a further 21% price increase in the energy price cap in January 2023. Unit prices for electricity under the EPG from October 2022 to June 2023 were on average 40% below the cap price, while unit gas prices were on average 29% lower. Weighting by average shares of spending on electricity and gas from 2019 this implies a 38.5% subsidy on the marginal cost of energy. While the EPG reduced unit prices, it left fixed fee standing charges for both electricity and gas essentially unaffected.[3]

[3]The cap prices and prices set under the EPG differed across users according to their payment type. However the price changes over this period were very similar across payment methods. Figure A.1 in the Appendix shows the cap levels for pre-payment users in the same way as Figure 2.2.

Rebates and transfers

In addition to the energy price subsidy provided by the EPC, the government provided further support through general rebates on energy bills. The government paid rebates on household energy bills of £400 in monthly instalments of either £66 or £67 from October 2022 until March 2023. In most cases, those paying by direct debit had these payments directly credited to their account, with the rest receiving the payments as cash refunds. Those paying by pre-payment either had funds directly added to their meter (if using a newer 'smart' meter), or received vouchers in the post that could be redeemed to add credits to their meters (if using older 'traditional' meters).

On top of measures that were explicitly labeled as addressing the rising costs of energy, the government also provided households with direct cash payments to help with the general increase in the cost of living. These were, in various ways, targeted to those likely to be in most need. Households in receipt of means tested benefits, pensions or disability benefits received "cost of living" payments of varying amounts that were paid directly into their bank accounts in instalments from summer 2022 until spring 2024.[4] Households also received a £150 rebate on their council tax in the summer and autumn of 2022 if they lived in homes with Council Tax bands A-D (loosely targeting lower value properties).

The total costs of all the Energy Price Guarantee, rebates on energy bills, cost of living payments and council tax rebates was £45.1 billion across 2022-23, equivalent to 1.4% of GDP over this period (Office for Budget Responsibility (2023)). Most of this was accounted for by the energy support measures: £20.3 billion was spent on the EPG, and a further £12.7 billion on the energy rebates.

2.4 Bank account data

Our main dataset contains information on spending and incomes drawn from individual-level bank account and credit card statements. The data are collected by two fintech companies. We use historic data covering the period 2015-19, which was collected by Moneydashboard, and data that cover the energy price shock, over 2019-23, which was collected by ClearScore. Both Moneydashboard and ClearScore are designed to assist users in monitoring and managing their finances; Moneydashboard provides users with access to a budgeting app, and ClearScore additionally provide users with access to their credit report and

[4] Payments of £650 were paid to those in receipt of means tested benefits in two instalments in the summer of 2022 and in November, with an additional £300 for those receiving pensions and £150 for those receiving disability benefits that were both paid in the winter. Further cost of living payments of £900 for those on means tested benefits were later announced to be paid in three instalments from spring 2023 to spring 2024, alongside further payments of £300 for those receiving pensions to be paid in winter 2023 and £150 for those receiving disability benefits that were paid in the summer 2023.

scores, and to Open Banking data. We undertake the majority of our analysis with the ClearScore dataset. We use the Moneydashboard data as one of a number of alternative ways of controlling for seasonality in energy spending. Therefore, in this section we describe the ClearScore data. We provide further details about the ClearScore data, and on the Moneydashboard data, in Appendix A.

Anyone that signs up to ClearScore links in their bank account details, from which their entire transaction histories over the previous three years are extracted. This gives us the transaction-level data (as shown on bank statements) for almost half a million UK users over the period 2019 to 2023. One of the advantages of the ClearScore data, relative to similar data provided by single banks, is that users are encouraged to link all of their accounts and credit cards. This gives us a relatively complete picture of user's incomes and spending patterns. As users are encouraged to link all of their accounts, including those held jointly with their spouse, we refer to them as households.

We use a sample of households that have at least one account that records spending on energy, aggregating the data (including across accounts) to the household-year-month level. We construct measures of monthly energy (gas and electricity combined) spending,[5] monthly non-durable spending, and monthly income. Our ClearScore sample includes 11.1 million observations covering 464,000 consumers over the period 2019-23. Younger individual and those that live in northern regions of the UK are over represented in the data – we therefore reweight our sample to match UK population figures. Appendix A contains more details on the sample construction and representativeness.

The energy bill rebates that households received between October 2022 and March 2023 were administered in one or two ways. Some households had their accounts directly credited with the rebate amount (or for those that use pre-paymeant, received vouchers). Other households, that pay by direct debit, were paid a refund directly into the bank account they use to pay their energy bills. For households in this latter group we observe the payment of the refund. For households in the former we group we adjust their observed energy spending to account for the value of the bill rebate.

The ClearScore data provide us with a detailed picture of spending patterns over the period of elevated energy prices. In Figure 2.3 we show that the energy spending patterns in the data align well with other data source. In panel (a) we show that the distribution of energy spending in 2019 in the ClearScore data is similar to that in the UK's national expenditure survey (the LCFS); the principle difference is, unlike in the survey data, the

[5]We focus on combined gas and electricity spending as 70% of UK electricity consumers and 80% of UK gas consumers have duel bills, meaning they pay for both gas and electricity together. This means we cannot distinguish between spending on gas and electricity. Note, that 80% of household use both electricity and gas (Ofgem, 2018).

ClearScore data do not exhibit any rounding bias. In panel (b) we show that the aggregate energy spending trend in the ClearScore data is similar to that in the UK National Accounts.

Figure 2.3: *Comparison of energy spending in ClearScore with other data*

(a) Cross-sectional distribution

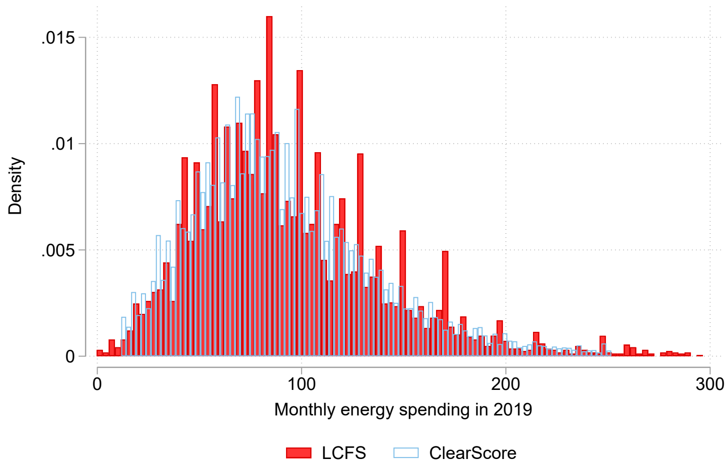

(b) Time trends

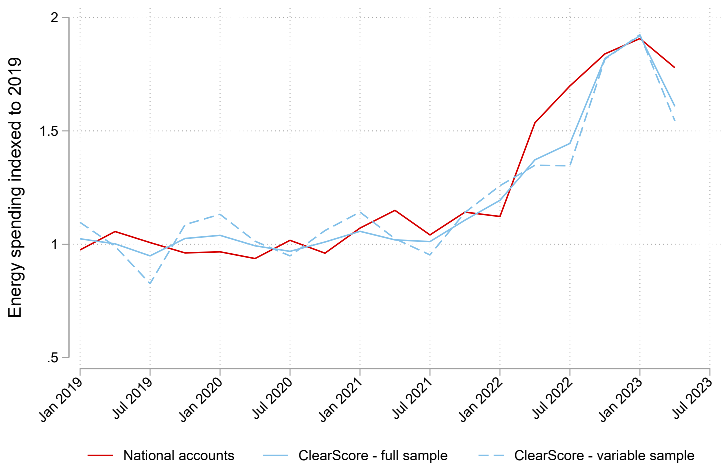

Notes: The top panel plots the distribution of monthly energy spending (in 2019) in the ClearScore data and in the UK's national expenditure survey, the Living Costs and Food Survey. The bottom panel shows the evolution of average monthly energy spending in the ClearScore data compared with that measured in the UK's National Accounts. In both pictures the ClearScore data are reweighted by age and region to match the UK population.

As described above, some households have variable billing and other fixed direct debits. We distinguish between these two groups based on the payment method we observe households using (e.g., direct debit or card payment) and variability in the amount they spend (see Appendix A for details). We undertake most of our analysis on the subsample of households that use variable billing (approximately 40% of the full sample). For this group, monthly energy spending closely aligns with their monthly usage. We focus on the period when the regulatory price cap was binding. This leaves with a sample of 112,340 households and 1.44million year-months covering June 2021 to May 2023. Figure 2.3(b) shows that the trend in energy spending for this variable sub-sample closely track that of the full sample, albeit with more seasonal variation (as expected).

For each household-month in our analysis sample, we measure the quantity of energy consumed by dividing spending with an energy price index. Specifically, we first subtract the fixed standing charges, for both electricity and gas, from total monthly energy spending. We then divide this amount with a fixed-weight price index based on the unit prices of gas and electricity. As weights we use the average share of energy spending allocated to gas and to electricity by people living in the household's local area prior to the price shock (over 2019-2020). We provide further details in Appendix A. Note our measure of energy prices is designed to capture the marginal price. Ito (2014) shows evidence that, in response to complex non-linear pricing schedules (involving several marginal rates), consumers respond to average rather than marginal prices. In the UK market, the vast majority of households face just a single marginal price for electricity and for gas.

2.5 Weather data

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

3 The distributional effects of the crisis and current policy

In this section we provide evidence on the distribution of exposure to energy price rises (based on pre-shock energy consumption), how average consumption responded to the price shock, and the distribution of consumption responses.

3.1 Exposure to energy price shocks

In Figure 3.1 we document how patterns of exposure to price rises vary across households. Specifically, we show how average monthly energy spending over 2019-2020, prior to the energy price shock, varies with average monthly income (panels (a) and (b)), and within quintiles of the average monthly income distribution (panels (c) and (d)).

Panel (a) shows that richer households, on average, spend more on energy – for instance, those in the top income decile, on average, spend roughly 50% more on energy than those in the bottom decile. However, as panel (b) shows, higher income households typically allocate a smaller fraction of their total non-durable spending to energy than poorer households. Therefore while in monetary terms higher incomes households are more exposed than less well off households to price rises (i.e., in the absence of behavioral responses their energy spending will rise by more), exposure rises less than proportionately with total household spending.

Panels (c) and (d) highlight that there is substantial heterogeneity in energy spending, *conditional* on income. For instance, although richer households spend more on energy on average, 20% of households in the bottom income quintile spend more than £100 on energy a month, or more than 15% of their total non-durable spending. Therefore a substantial fraction of relatively low income households are highly exposed to energy price rises.

Figure 3.1: *Pre-shock energy spending and budget shares by income*

(a) Average pre-shock energy spending

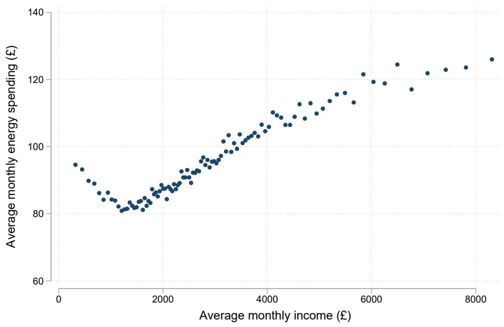

(b) Average pre-shock energy budget shares

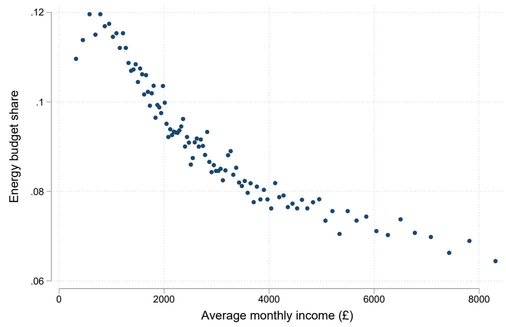

(c) Distribution of pre-shock energy spending

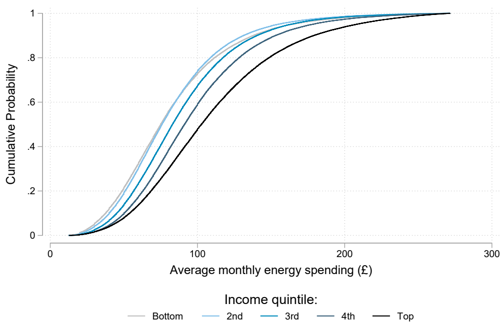

(d) Distribution of pre-shock energy budget shares

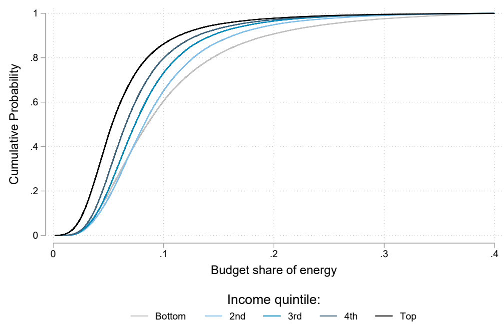

Notes: Panel (a) shows a binscatter of households' mean monthly energy spending in the pre-shock period (2019 and 2020) against mean monthly income. Panel (b) shows a binscatter of household' energy budget shares (energy spending over non-durable spending) in the pre-shock period against mean monthly income. Panels (c) and (d) show the distributions of pre-shock monthly energy spending and budget shares by quintiles of the average monthly income distribution.

3.2 Average changes in energy usage over the crisis

Figure 3.2(a) shows changes in deseasonalized log energy expenditure over the period 2020 to 2023.[6] The vertical dashed lines indicate periods when the value of the energy price cap was updated and, in April 2023, when the government ceased providing bill rebates. The shaded area highlights the period when the energy price cap was binding for the majority of households. The figure shows that average energy spending increased at each price cap rise, and decreased in April 2023 after the bill rebates stopped.

Panel (b) of Figure 3.2 focuses on the period when the price cap was binding and plots the evolution of average log energy consumption. It shows that there were sharp discontinuous drops in the quantity of energy consumed following price cap changes. The quantity response for the 27% price cap, which coincided with the introduction of energy bill rebates of £67 a month, is less pronounced that for the other cap rises. When rebates ended in April 2023 (at which point there was no price change) both energy spending and consumption exhibits drops.

Overall, the figure indicates quantity falls (and spending rises) at the 12% and 54% energy price cap rises, consistent with an energy price elasticity in between 0 and -1. The spending and quantity changes in response to the introduction and withdrawal of bill rebates indicates an economically meaningful marginal propensity to consumer energy in response to the rebates.

There is some evidence of anticipatory effects prior to the 54% cap rise. It is possible that some households topped up their pre-payment meters at the lower price, and then ran this down during the higher cap period.[7] Below we shows estimates of the energy price elasticities below, and that removing variation for the months immediately prior to and following a cap change has only a minimal impact estimates.

A potential concern with our sample, which consists of variable payment households, is that these are not representative of the UK population as a whole. To help address this concern we show that the impact of reweighting our sample (to match the UK population on the basis of age and region) on price elasticity estimate. However, it is possible that there are unobservable differences between those that pay by fixed direct debit and those who

[6]To deasonalise we first use the MoneyDashboard data to estimate calendar month effects over the period 2015-19 (controlling for the price of energy). These are shown in Appendix B. We then subtract the estimated month effects from log energy spending observed in the ClearScore data from 2020 onwards to construct our deseasonalized measure of spending. The regressions show the estimated year-month effects from January 2020 onwards, after controlling for a flexible function of local temperature and rainfall. Figure B.2 in the Appendix shows the trends in the raw data, without controlling for any seasonality.

[7]For some households using old pre-payment electricity meters this would have enabled them to pay the pre cap rise price on some post cap rise consumption. However, there was no incentive to do this for households using a smart electricity pre-pay meters, and for any type of gas pre-pay meter, as these meters allow suppliers to charge current rates for use.

Figure 3.2: *Log energy spending and quantities by month*

(a) Spending

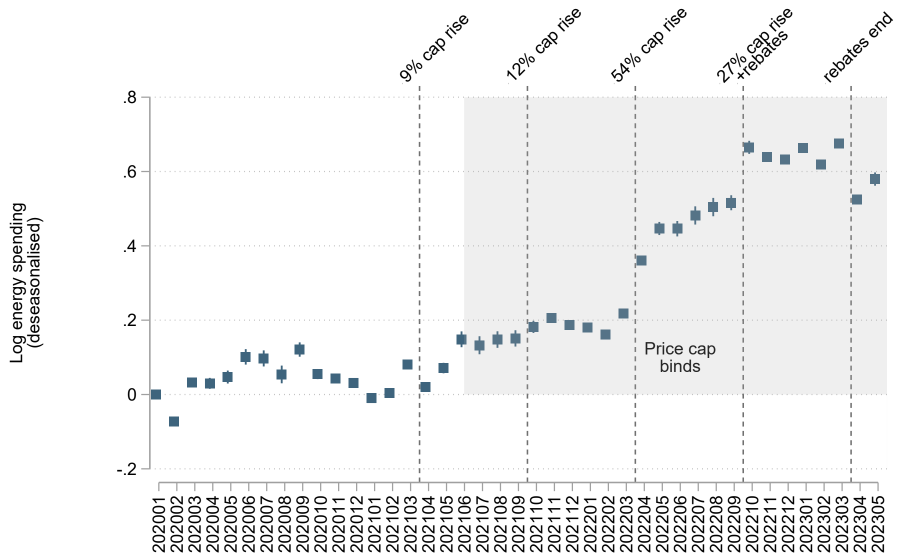

(b) Quantity

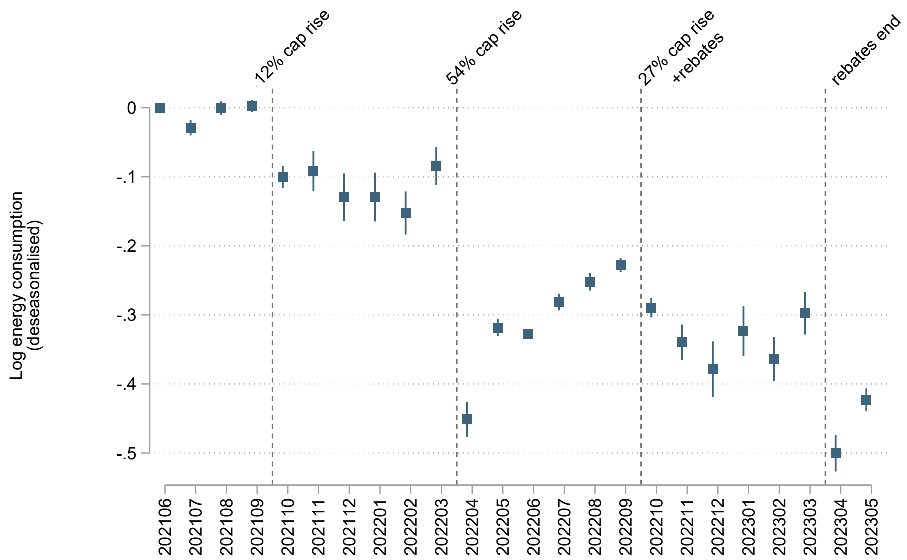

Notes: The top (bottom) figure shows total log deseasonalized energy spending (quantities) in each month.

use variable billing. Figure 2.3(b) shows that the aggregate trend in energy spending for the variable and full sample are similar. In Appendix B we compare the raw and deseasonalized trends in log energy spending and quantity for the variable and full samples. As is to be expected, there is more seasonal variation in the raw trend for the variable payment sample. However, once we strip this out, the evolution of spending over the crisis for the two samples is very similar. We show below that this also translates into similar estimate price elasticities. One potential explanation for why the fixed payment sample appears to behave similarly to the variable payment is that people were encouraged to submit regular meter readings (especially pre and post the cap change) to ensure accurate billing.

3.3 Price elasticities

To estimate energy price elasticities, we focus on the period when energy cap changes were most likely to be binding, from June 2021 to September 2022. For the moment, we exclude the period following the cap change in October 2022 as this price change coincided with the introduction of energy rebates. We return to these below.

We estimate the following equation:

$$\log q_{it} = \gamma \log p_t + g(\text{temp}_{r(i)t}, \text{rain}_{r(i)t}) + \psi X_t + \zeta_i + \epsilon_{it} \tag{3.1}$$

where q denotes energy quantity, p_t is the marginal price of a unit of energy in period t, $g(\text{temp}_{r(i)t}, \text{rain}_{r(i)t})$ is a flexible function of minimum and maximum monthly temperatures and rainfall in location r, and ζ_i is a household fixed effect. The coefficients γ is the elasticity of energy demand with respect to price. X_t includes a dummy for periods during which people were encouraged to work from home during the COVID pandemic, and, in some specifications, a set of calendar month effects. We estimate a specification where the dependent variables is deseasonalized, and one where it is not deseasonalized and we control for calendar month effects.

We report results for different specifications in Table 3.1. Column (1) shows the estimate for the specification that is based on deseasonalized energy consumption and controls for local weather (in addition to household fixed effect). This gives an energy price elasticity estimate of -0.49. In column (2), we report the estimate for the specification where the outcome variable is not deseasonalized and where instead we include calendar month dummies. This yields a similar estimate to of -0.47. In column (3), we report estimates for the same specification as in column (1), but in this case we weight the regression according to age, region and pre-shock energy expenditure. The resulting estimate of -0.48 can be interpreted as the aggregate energy price elasticity. In column (4) we report the elasticity estimate when we omit data for the month immediately prior to and following a cap change, to strip out any effect arising from households that pay by pre-payment attempting to “stock up” on their meters – this reduces the magnitude of the elasticity estimate slightly to -0.47.

Table 3.1: *Energy price elasticities (July 2021 - September 2022)*

	(1)	(2)	(3)	(4)
Log energy quantity				
$\log(p_t)$	-0.490	-0.474	-0.484	-0.470
	(0.004)	(0.005)	(0.008)	(0.004)
Household FE	Yes	Yes	Yes	Yes
Deseasonalized	Yes	No	Yes	Yes
Weather controls	Yes	Yes	Yes	Yes
Month effects	No	Yes	No	No
Anticipation effects	No	No	No	Yes
Reweighting by	No	No	Age, region, exp.	No

Notes: Standard errors in parentheses and clustered at the household level. All specifications include controls for local weather (fifth order polynomials for minimum and maximum monthly temperatures, rainfall and the squared different between the monthly minimum and maximum local temperatures) and household fixed effects. In column (2) we additionally control for a calendar month dummies. In column (3) we re-weight observations by their expenditure multiplied by inverse probability weights estimated using region and age. Column (4) includes dummies for the months immediately prior to and after a cap change.

Heterogeneity in energy price elasticities

The distributional impact of the energy price shock is, in part, determined by the *distribution* of consumer responses to the price change. In particular, if a households lowers its energy consumption in response to a price rise, and this is primarily driven by a substitution effect (switching to alternatives in response to relative price changes) rather than an income effect (lowering energy consumption due to reduced purchasing power) this will act to mitigate the welfare loss due to the price rise. Here we document heterogeneity in households' responses to price changes. In Section 4 we outline a model that we use to decompose price responses into substitution and income effects.

To measure heterogeneity in energy price elasticities we estimate:

$$\log q_{it} = \sum_{d=1}^{D} \gamma_d \log p_t + g(\text{temp}_{r(i)t}, \text{rain}_{r(i)t}) + \zeta_i + \epsilon_{it} \tag{3.2}$$

We estimate two versions of equation (3.2). In the first d indexes quintiles of pre-shock energy spending. The red line in Figure 3.3 shows that there is a downward sloping gradient in the energy price elasticity with pre-shock energy spending: on average, households with higher pre-shock energy spending have larger (in absolute terms) price elasticities.

We also specify a variant of equation (3.2) that allows us to jointly estimates heterogeneity in γ across income and pre-shock energy spending quintiles. The estimates are illustrated by the blue lines in Figure 3.3. They show that, conditional on income, the pattern of higher price responsiveness by those with higher pre-shock spending persists, and, conditional on pre-shock spending, higher income households tend to be less price sensitive.

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

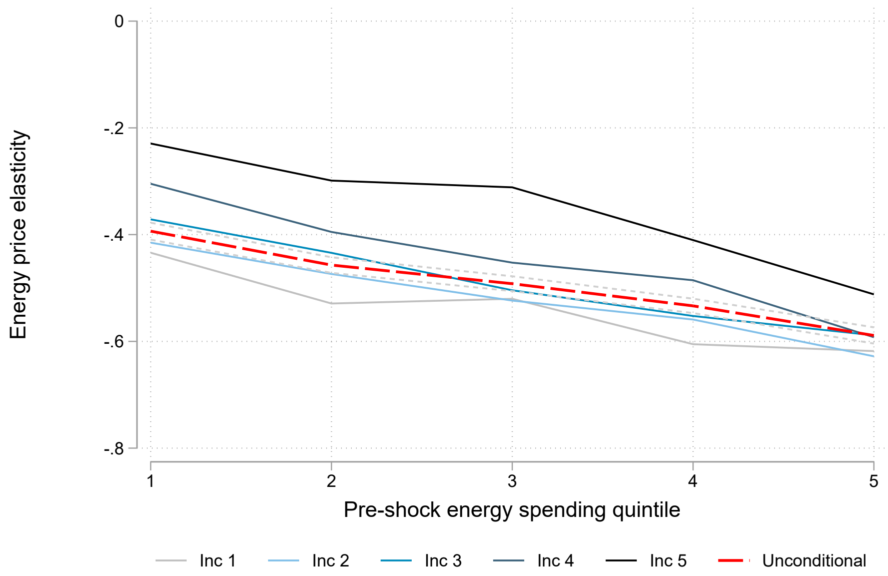

Notes: Figure shows the estimates of $\hat{\gamma}_d$ from (3.2). The red line are estimated with d indexing pre-shock energy spending quintile. The blue lines are estimated with d indexing the interaction between pre-shock energy spending quintile and income quintile. 95% confidence intervals are shown in the grey dots.

3.4 Rebates

In column 1 of Table 3.2 we report estimates based on a regression that is similar to equation (3.1), but with the exception that the outcome variable is log energy spending, we use data covering July 2021 until May 2023, and we include an indicator variable for the period when bill rebates were given. The implied energy own price elasticity is given by one minus the coefficient on log price. Reassuringly, the elasticities are very similar to those reported in Table 3.1 when we incorporate data following the October 2022 price rise.

In Column 2 of Table 3.2 we compare the impact on the log of energy spending from receiving energy rebates to the receipt of 'cost of living' (COL) payment. These were paid to those that receive means-tested benefits in July and November 2022 and April-May 2023. The payments were £326 in July, £324 in November and £301 in 2023, and were therefore considerably larger than the £66 and £67 bill rebates. Total energy spending over the period when the rebates were in effect (inclusive of the rebate), among individuals observed in all six months of this period, was £1090, and the average rebate imputed to these individuals was £321. The coefficients in column 2 of Table 3.2 imply that the rebates increased average energy spending by around 7%, or £76, which was around 22% of the rebate. By contrast, the estimates imply that they only 3% of the COL payments were spent on energy. The sample of people that received COL payments may differ to the broader sample of people that received rebates. In Column 3, we restrict the sample to those who were recorded as having received a COL payment at one point between July 2022 and May 2023. These individuals had a similar propensity to spend on energy out of the COL payments, but an even slightly higher propensity to spend out of energy rebates.

Since the fraction of those for whom the energy rebates were greater than their monthly energy spending is small (around 15% of households in September 2022), there is little reason for utility-maximizing household to treat money received through energy bill rebates differently to other income. Evidence of labeling effects has been document with respect to Winter Fuel Payments - labeled cash transfers to older UK households (see Beatty et al. (2014)) - as well in the context of SNAP benefits (see Hastings and Shapiro (2018)).

Table 3.2: *Effect of rebates and cost-of-living payments*

	(1)	**(2)**	**(3)**
Log energy spending			
$\log(p_t)$	0.562	0.556	0.553
	(0.003)	(0.003)	(0.005)
Rebates	0.0655	0.0678	0.0854
	(0.002)	(0.002)	(0.003)
COL payments		0.0274	0.0245
		(0.002)	(0.002)
Consumer FE	Yes	Yes	Yes
Deseasonalised	Yes	Yes	Yes
Weather controls	Yes	Yes	Yes
Sample	Variable	Variable	Variable, rec. COL

Notes: Standard errors in parentheses and clustered at the household level. All specifications include controls for local weather (fifth order polynomials for minimum and maximum monthly temperatures, rainfall and the squared different between the monthly minimum and maximum local temperatures) and household fixed effects. Columns (1) and (2) use data on our main analysis sample (those with variable billing); column (3) restricts the sample to only those with variable billing the received cost-of-living (COL) payments.

4 Model

In this section we outline an empirical model of household energy demand, which we use to quantity the welfare effect of observed and counterfactual policies. Two important features of our approach are, first we allow for the possibility that households, in some periods (namely when labeled rebates are administered), do not make choices that maximize true welfare. Second we use flexible to functional form to capture energy demands, as our objective is to trace out the distribution of household-level welfare effects.

4.1 Conceptual framework

Consider a household's decision over how to allocate their total period budget, x, between domestic energy consumption, e, and consumption of all other non-durables (excluding domestic energy), n. Let p_e denote the price (inclusive of any subsidy) of energy and p_n the price of other non-durables.[8]

The household's total budget is $x = \tilde{x} + t$, where $t > 0$ is a transfer given to the household and $\tilde{x}$ is its total budget in the absence of the transfer. If the transfer is given as a bill rebate the household's choice is subject to a labeling effect, which we denote by $L = 1$ (in the absence of a labeling effect, $L = 0$). The household's choice problem is:

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

Both their direct *decision* utility function, $U^L(.)$, and the associated maximum function, $\widetilde{V}^L(.)$, depend on whether there is a labeling effect, and are indexed by a parameter vector θ. The energy and non-durable demand functions take the form $e = \mathbb{e}^L(p_e, p_n, x; \theta)$ and $n = \mathbb{n}^L(p_e, p_n, x; \theta)$ for $L = \{0, 1\}$.

If the effect of labeling on the household's choice is to change their underlying normative preferences, the functions $\widetilde{V}^0(.)$ and $\widetilde{V}^1(.)$ directly tell us the level of utility attained by the household. However, in our baseline analysis we adopt the normative view that labeling distorts choices. In this case, the household's decision utility function in the absence of labeling effects, $U^0(.)$ (and hence $\widetilde{V}^0(.)$), reflects their true underlying preferences, and the function that maps the budget set into the utility level the household attains, given their choice behavior captured in equation (4.1), is given by:

$$V(p_e, p_n, L, x; \theta) = \begin{cases} U^0(\mathbb{e}^0(p_e, p_n, x; \theta), \mathbb{n}^0(p_e, p_n, x; \theta); \theta) & \text{if} \quad L = 0 \\ U^0(\mathbb{e}^1(p_e, p_n, x; \theta), \mathbb{n}^1(p_e, p_n, x; \theta); \theta) & \phantom{\text{if}} \quad L = 1. \end{cases} \tag{4.2}$$

$V(p_e, p_n, L, x; \theta)$ is similar to an indirect utility function, but with the important difference that when $L = 1$ it corresponds to choices that are sub-optimal. We refer to $V(p_e, p_n, L, x; \theta)$ as the household welfare function.

In Figure 4.1 we illustrate the difference between an undistorted choice ($L = 0$) and a distorted choice ($L = 1$), by comparing the two at the same budget set (given by 0AB). Specifically, $(\mathbb{e}^1, \mathbb{n}^1)$ is the household's (distorted) choice when given a bill rebate, while

[8] To use energy households must pay a fixed fee and marginal price. Therefore x is spending net of the fixed fee and p_e is the marginal price.

$(\mathbb{e}^0, \mathbb{n}^0)$ is their (undistorted) choice if they receive an unlabeled transfer of the same amount. The undistorted choice is at a point of tangency between the $U^0(.)$ function's indifference map and the budget constraint and yields utility level $\bar{U}^0$, while at the distorted choice the indifference curve intercepts the budget constraint (AB) and the associated utility level is $\bar{U}^1 < \bar{U}^0$.

Figure 4.1: *Distorted and undistorted choices*

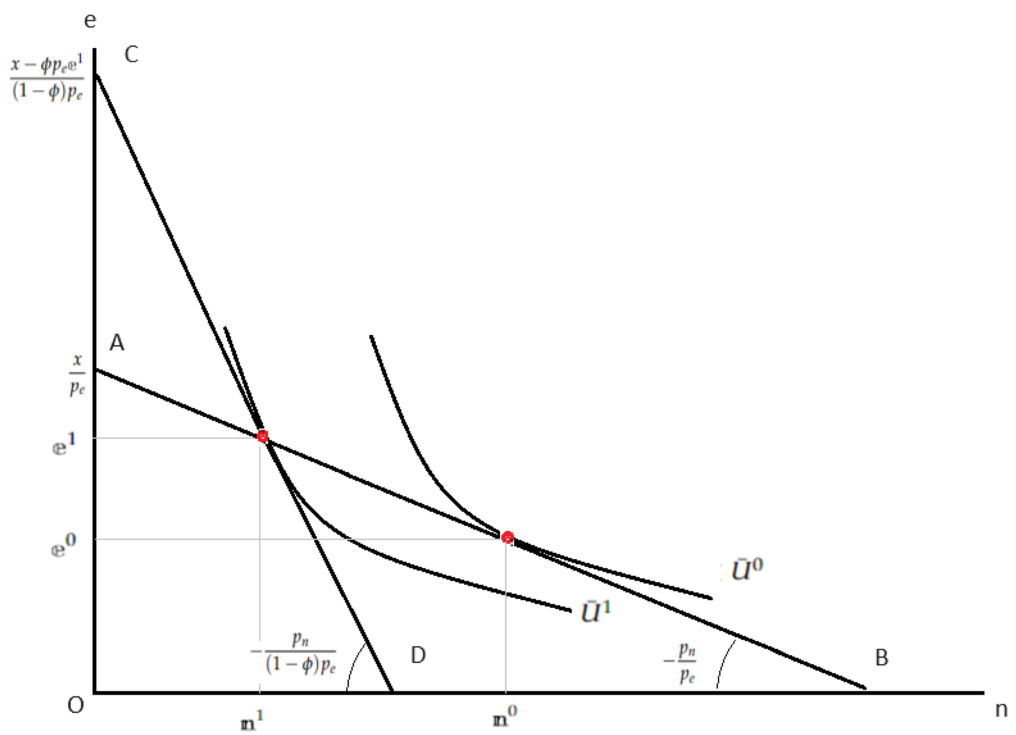

Notes: $(\mathbb{n}^1, \mathbb{z}^1)$ *and* $(\mathbb{n}^0, \mathbb{z}^0)$ *represent choices with (distorted) and without (undistorted) labeling effects, and* $\bar{U}^1$ *and* $\bar{U}^0$ *the corresponding utility levels, when the consumer faces the budget set OAB. The budget set OCD represents the rotation through the point* $(\mathbb{e}^1, \mathbb{n}^1)$ *that would result in* $(\mathbb{e}^1, \mathbb{n}^1)$ *being the undistorted choice under this new budget set. This entails changing the energy price from* p_e *to* $(1-\phi)p_e$ *and the budget from* x *to* $x - \phi p_e \mathbb{e}^1$.

Equation (4.2) defines the household welfare function in terms of the direct utility function $U^0(.)$. However, as our empirical model entails specifying a flexible form for the (inverse of) the function $\widetilde{V}^L(p_e, p_n, x; \theta)$, it is convenient to define the household welfare function in terms of $\widetilde{V}^0(.)$. In the absence of labeling both functions coincide. This is not the case when the household's choice is subject to a labeling effect. However, in this case we can find a point at which the values of the functions $V(.)$ and $\widetilde{V}^0(.)$ coincide, by considering the rotation of the budget constraint, through the point $(\mathbb{e}^1, \mathbb{n}^1)$, that results in tangency between a hypothetical budget constraint and the $U^0(.)$ indifference map (line CD in Figure 4.1). This allows us to write the household welfare function according to the following proposition.

Proposition 1. *The level of utility attained by a household when they make choices according to equation (4.1) and when their underlying preferences are reflected in their choice behavior in the absence of labeling ($L = 0$), is given by*

$$V(p_e, p_n, L, x; \theta) \;\; = \begin{cases} \widetilde{V}^0(p_e, p_n, x; \theta) & \text{if} \quad L = 0 \\ \widetilde{V}^0(p_e(1-\phi), p_n, x - \phi p_e \mathbb{e}^1; \theta) & \phantom{\text{if}} \quad L = 1 \end{cases}$$

where $\mathbb{e}^1$ and ϕ are such that:

$$\mathbb{e}^1 \in \arg\max_{e,z} U^1(e, n; \theta) \; s.t. \; p_e e + p_n n \leq x$$

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

ϕ is the compensated percent price reduction for energy that rationalizes the distorted choice $\mathbb{e}^1$ as the welfare-maximizing choice.

The welfare effects of policy

The UK government responded to the price shock by introducing an energy price subsidy, s^O (meaning the energy price households faced related to the pre-subsidy price by $p_e = (1 - s^O)P_e$) and a transfer of value t^O via a bill rebate (and hence with a labeling effect, $L = 1$). Consider a continuum of households, indexed by i, with population normalized to 1 and let $(\mathbb{s}, \mathbb{t}, \mathbb{L})$ denote an arbitrary combination of subsidy, transfer and labeling. We write a household's welfare function directly as a function of the policy parameters, $V_i(\mathbb{s}, \mathbb{t}, \mathbb{L}) \equiv V((1-\mathbb{s})P_e, p_n, \mathbb{L}, \tilde{x}_i + \mathbb{t}; \theta_i)$, where the i index reflects the household's pre-policy budget, $\tilde{x}_i$, and preference vector, θ_i. We convert $V_i(\mathbb{s}, \mathbb{t}, \mathbb{L})$ to a money-metric form using the inverse function $C(p_e, p_n, L, u; \theta) \equiv V^{-1}(p_e, p_n, L, .; \theta)$. Specifically, letting P_e^0 denote the energy price in the absence of the price shock, we define the consumer's equivalent variation as:

$$EV_i(\mathbb{s}, \mathbb{t}, \mathbb{L}) = \tilde{x}_i - C(P_e^0, p_n, 0, V_i(\mathbb{s}, \mathbb{t}, \mathbb{L}); \theta_i), \tag{4.3}$$

which gives the monetary value household i is willing to pay in the absence of the price shock (and hence any policy intervention), to avoid the shock with policy response $(\mathbb{s}, \mathbb{t}, \mathbb{L})$.

To understand the role played by each element of the government's policy response, we compare observed policy (s^O, t^O, L^O) to the alternative policies; no intervention $(0, 0, 0)$, subsidy only $(s^O, 0, 0)$, and subsidy and unlabeled transfer $(s, t', 0)$. In the case of an unlabeled transfer, we consider both $t' = t^O$ and when the transfer value is adjusted to achieve budget balance. In the latter case we use the government's budget constraint;

$$\mathbb{s} P^E \int_i e_i(\mathbb{s}, \mathbb{t}, \mathbb{L}) di + \mathbb{t} = \bar{R}, \tag{4.4}$$

where $e_i(\mathbb{s}, \mathbb{t}, \mathbb{L}) \equiv \mathbb{e}^{\mathbb{L}}((1-\mathbb{s})P_e, p_n, \tilde{x} + \mathbb{t}; \theta)$ is household energy demand and $\bar{R}$ is the resources the government uses to fund observed policy (s^O, t^O, L^O).

Optimal policy

To consider the optimal balance between subsidies and transfers, it is necessary to specify a social welfare function, which entails taking a stand on inter-household welfare comparisons. We consider a social welfare function where a household's contribution depends both on the percentile of the income distribution to which they belong, $y_{p(i)}$, and their equivalent variation:

$$W(\mathbb{s}, \mathbb{t}, \mathbb{L}) = \int_i f(y_{p(i)}) G(EV_i(\mathbb{s}, \mathbb{t}, \mathbb{L})) di, \tag{4.5}$$

where $f(.) > 0$, $f'(.) \leq 0$, $G'(.) > 0$ and $G''(.) \geq 0$. The function $f(.)$ controls the government's concern for vertical equity. If $f'(.) = 0$ the government values a given loss equally across households, regardless of their income level; if $f'(.) < 0$ the government places higher weight on a loss the lower is a household's income. The function $G(.)$ controls the government's concern for horizontal equity. If $G''(.) = 0$ the government places the same weight on a marginal cash lose to a household regardless of whether it is the first £ lost of the 1000^{th}. If $G''(.) > 0$ the government places more weight on avoiding large monetary losses than small loses. The government's problem is to choose the policy parameters $(\mathbb{s}, \mathbb{t}, \mathbb{L})$ to minimize equation (4.5) subject to its budget constraint (equation (4.4)). Note, as labeling distorts households' choices and leads to higher consumption of a subsidized good, it is clear that optimal policy entails setting $\mathbb{L} = 0$. In addition, note that when the government has no horizontal equity concerns, so $G''(.) = 0$, minimizing equation (4.5) is equivalent to the standard optimal tax problem of maximizing the (vertical social welfare weighted) sum of household-level utilities.

4.2 Empirical demand model

We specify a flexible form for the choice model in equation (4.1). Using information on our sample of households and year months (indexed τ) on energy consumption, prices ($\mathbf{p}_\tau \equiv (p_{e\tau}, p_{n\tau})$), total budgets, an indicator for whether there is a labeled rebate and conditioning variables, $(e_{i\tau}, \mathbf{p}_\tau, x_{i\tau}, L_\tau, \mathbf{z}_{i\tau})$, we estimate the parameters governing the energy demand equation and hence the maximum function $\tilde{V}^L(.)$. We use this to construct the household welfare function based on proposition 1.

We specify an energy demand model in the Exact Affine Stone Index (EASI) class, developed in Lewbel and Pendakur (2009). This class of demand models provides a way of capturing rich heterogeneity in behavior across households, while observing the behavioral restrictions implied by consumer theory. We specifying a EASI form for Hicksian budget share demand for energy, $w_{i\tau} = \omega(\mathbf{p}_\tau, L_\tau, u_{i\tau}, \mathbf{z}_{i\tau}, \epsilon_{i\tau}; \mathbf{\Psi})$, where $w_{i\tau} \equiv \frac{p_{e\tau}e_{i\tau}}{x_{i\tau}}$ is the energy budget share,[9]

$u_{i\tau}$ is attained level of *decision* utility,[10] $\epsilon_{i\tau}$ is a preference shock and $\mathbf{\Psi}$ a parameter vector.[11] A defining feature of EASI demand systems is that they give rise to *implicit* functions for the Marshallian budget share demand and indirect (decision) utility, of the form $w_{i\tau} = \omega(\mathbf{p}_\tau, L_\tau, y_{i\tau}, \mathbf{z}_{i\tau}, \epsilon_{i\tau}; \mathbf{\Psi})$ and $y_{i\tau} = v(w_{i\tau}, \mathbf{p}_\tau, L_\tau, x_{i\tau}, \mathbf{z}_{i\tau}; \mathbf{\Psi})$, where by construction $y_{i\tau} = u_{i\tau}$.

We specify the implicit Marshallian budget share demand:

$$\begin{aligned} w_{i\tau} = \big(A + \sum_{l\in\mathcal{Z}_1} A_l z_{i\tau l}\big) + \big(B + \sum_{l\in\mathcal{Z}_2} B_l z_{i\tau l}\big) \times (\log p_{e\tau} - \log p_{n\tau}) + \big(C_1 + \sum_{l\in\mathcal{Z}_1} C_{1l}\big) y_{i\tau} \\ C_2 y_{i\tau}^2 + D\,(\log p_{e\tau} - \log p_{n\tau}) \times y_{i\tau} + \big(\delta + \sum_{l\in\mathcal{Z}_2} \delta_l z_{i\tau l}\big) L_\tau + \epsilon_{i\tau}, \end{aligned} \tag{4.6}$$

The A parameters determines the budget share intercept, the B parameters govern the demand response to a (compensated) price change, the C parameters capture the energy Engel curve and D allows for an interaction effect between price responses and the Engel curve. By including the log difference in prices in equation (4.6), we ensure preferences satisfy adding-up, demand homogeneity, and Slutsky symmetric.[12] The δ parameters captures the behavioral effect of labeling. We allow for indicator variables for the pre-shock energy spending decile the household belongs to (collected in the set $\mathcal{Z}_2$) to shift the intercept, price response, Engel curve and labeling effect. In addition, we include seasonal month dummies and detailed weather controls in the set of intercept shifters ($\mathcal{Z}_1 \supset \mathcal{Z}_2$).

The form for implicit decision utility consistent with equation (4.6) is:

$$y_{i\tau} = \frac{\log x_{i\tau} - (w_{i\tau}\log p_{e\tau} + (1-w_{i\tau})\log p_{n\tau}) + \frac{1}{2}\big(B + \sum_{l\in\mathcal{Z}_2} B_l z_{i\tau l}\big) \times (\log p_{e\tau} - \log p_{n\tau})^2}{1 - \frac{1}{2}D \times (\log p_{e\tau} - \log p_{n\tau})^2} \tag{4.7}$$

Estimation

Let $\Psi = (A, \{A_l\}_{l\in\mathcal{Z}_1}, B, \{B_l\}_{l\in\mathcal{Z}_2}, C_1, \{C_{1l}\}_{l\in\mathcal{Z}_2}, C_2, D, \delta, \{\delta_l\}_{l\in\mathcal{Z}_2})$ denote the model parameters. We obtain estimates for these using an iterated two-stage least squares estimator (see Blundell and Robin (1999)). This entails fixing an initial guess of implicit decision utility

[9]As households pay a fixed fee for using energy, $p_{e\tau}e_{i\tau}$ is variable spending on energy, $x_{i\tau}$ is total non-durable spending net of the fixed fee, and $w_{i\tau}$ is the share of net non-durable spending the household allocates to variable energy spending.

[10]Specifically u is given by the function $V^L(.)$ in equation (4.1).

[11]The budget of other non-durables is equal to $1 - w_{i\tau}$, by the adding-up restriction.

[12]Consumer theory also requires that the inverse of the decision utility function $\widetilde{E}^L(p_e, p_n, u; \theta) \equiv \widetilde{V}^{L\,-1}(p_e, p_n, .; \theta)$ is concave in prices and increasing in u. This implies inequality constraints that we check are satisfied ex post. See Appendix C for a detailed discussion of the demand model and the regularity conditions imposed by consumer theory.

$y_{i\tau}$, estimating the energy budget share demand (equation (4.6)), updating implicit decision utility (equation (4.7)) and continuing the procedure until convergence. At each stage of the iteration we estimate the budget share demand using two-stage least squares, instrumenting for implicit decision utility. We do this for two reasons. First, implicit decision utility itself depends on the budget shares, which leads to a mechanical correlation between implicit decision utility and the unobserved shock, $\epsilon_{i\tau}$. Second, it may be that total expenditure, $x_{i\tau}$, is correlated with $\epsilon_{i\tau}$ – a period of unexpectedly high energy needs could lead the household to raise total period expenditure. Our instrument for $y_{i\tau}$ is:

$$\tilde{y}_{i\tau} = \frac{\log m_{i\tau} - (\bar{w}\log p_{e\tau} + (1-\bar{w})\log p_{n\tau}) + \frac{1}{2}\left(B + \sum_{l\in\mathcal{Z}_2} B_l z_{i\tau l}\right) \times (\log p_{e\tau} - \log p_{n\tau})^2}{1 - \frac{1}{2}D \times (\log p_{e\tau} - \log p_{n\tau})^2}, \quad (4.8)$$

where $m_{i\tau}$ is the monthly period τ income and $\bar{w}$ is the average energy budget shares across all household and time periods.[13]

Measuring household welfare

With the estimated parameters, we can solve equations (4.6) and (4.7), for any set of prices, expenditures and/or labeling, to compute the corresponding Marshallian energy demands and maximized decision utility. We denote these functions by $\mathbb{e}(p_e, p_n, L, x, \mathbf{z}_{i\tau}, \epsilon_{i\tau}; \mathbf{\Upsilon})$ and $\tilde{V}^L(p_e, p_n, x, \mathbf{z}_{i\tau}, \epsilon_{i\tau}; \mathbf{\Upsilon})$. When $L = 0$, the household's welfare function is given by: $\tilde{V}^0(p_e, p_n, x, \mathbf{z}_{i\tau}, \epsilon_{i\tau}; \mathbf{\Upsilon})$. When $L = 1$, we use proposition 1, which entails finding the rotation of the budget constraint (captured by the parameter ϕ) through the point of distorted choice that rationalizes this choice as an undistorted choice, which allows us to use $\tilde{V}^0(.)$ to measure the household's true utility level at the distorted choice. Specifically, let $\mathbb{e}^1 = \mathbb{e}(p_e, p_n, 1, \mathbf{z}_{i\tau}, \epsilon_{i\tau}; \mathbf{\Upsilon})$ denote the distorted choice. We solve for the budget rotation ϕ using the iterative algorithm:

$$\phi^{(i)} = \phi^{(i-1)} + \log \mathbb{e}^1 - \log \mathbb{e}(p_e(1-\phi^{(i-1)}), p_n, 0, x - \phi^{(i-1)} p_e \mathbb{e}^1, \mathbf{z}_{i\tau}, \epsilon_{i\tau}; \mathbf{\Upsilon}),$$

until $||\phi^{(i)} - \phi^{(i-1)}|| < 10e^{-3}$.

5 Results

5.1 Estimates

In Table 5.1 we report the demand model parameter estimates. The first column reports the baseline parameter estimates, and the remaining columns the pre-shock spending decile

[13] In practice, when iterating between equations (4.6) and (4.7) we do note update the parameters of the instrument. Rather, we initially iterate between the two equations using $\log m_{i\tau} - (\bar{w}\log p_{e\tau} + (1-\bar{w})\log p_{n\tau})$ as the instrument. Using the converged parameters we construct the $\tilde{y}_{i\tau}$ according to equation (4.8). We then undertake the iterative estimation procedure a second time with this new instrument. The reported estimates in Table 5.1 are the converged parameters from this second stage.

interaction effects. The table shows that model estimates are statistically significant at conventional levels. Energy consumption responses to price or total expenditure changes are determined jointly by all the parameters, meaning individual parameter estimates are difficult to interpret. We therefore report elasticities implied by the model in Figure 5.1.

Table 5.1: *Parameter estimates*

		×Pre-shock spending decile								
		2	3	4	5	6	7	8	9	10
Constant										
(A)	1.5712	0.0876	0.1650	0.2245	0.2754	0.3125	0.3463	0.4133	0.4630	0.5290
	(0.0501)	(0.0118)	(0.0120)	(0.0123)	(0.0123)	(0.0126)	(0.0129)	(0.0139)	(0.0153)	(0.0186)
Price										
(B)	0.1322	-0.0041	-0.0045	-0.0036	-0.0040	-0.0032	-0.0055	-0.0065	-0.0050	-0.0195
	(0.0079)	(0.0014)	(0.0014)	(0.0014)	(0.0014)	(0.0014)	(0.0014)	(0.0015)	(0.0016)	(0.0018)
Implicit utility										
(C_1)	-0.4186	-0.0113	-0.0214	-0.0290	-0.0354	-0.0398	-0.0435	-0.0519	-0.0572	-0.0631
	(0.0135)	(0.0017)	(0.0017)	(0.0018)	(0.0018)	(0.0018)	(0.0018)	(0.0020)	(0.0022)	(0.0026)
(C_2)	0.0286	-	-	-	-	-	-	-	-	-
	(0.0010)									
Price× Implicit utility										
(D)	-0.0134	-	-	-	-	-	-	-	-	-
	(0.0012)									
Labeling										
(δ)	0.0143	0.0004	-0.0010	-0.0006	-0.0017	-0.0032	-0.0033	-0.0039	-0.0075	-0.0158
	(0.0007)	(0.0009)	(0.0009)	(0.0009)	(0.0009)	(0.0009)	(0.0009)	(0.0010)	(0.0010)	(0.0011)

Notes: Estimates are for budget share energy demand (equations (4.6) and (4.7)) and are estimated by iterated 2SLS. We use monthly income to instrument for total expenditure. The model includes month dummies, 5^{th}-order polynomials in local monthly minimum and maximum temperature, the squared difference between maximum and minimum temperature, local monthly rainfall and an indicator for Dec 2021 and Jan 2022, when a work-from-home order was in place.

In Figure 5.1 we report simulated elasticities based on the period on the October 2021-September 2022. We compare consumption during this period with counterfactual consumption had the price of energy relative to other non-durables been at its "pre-shock" April-September 2021 level. This corresponds to a 9% price increase over October 2021-March 2022 and a 62% price rise over April 2022-September 2022.[14] In panel (a) we report mean Marshallian price elasticities, computed based on the percent change in household energy consumption divided by the percent change in the energy price, associated with the increase in the (relative) price of energy from its pre-shock level. We report means by quintiles of the pre-shock energy spending and income distributions. The pattern of elasticities increasing with pre-shock energy spending, and conditional on this, exhibiting little variation by income, shown in the figure matches the evidence in Figure 3.3.

The Marshallian price elasticities reflect both substitution responses associated with the relative increase in energy prices, and any income effects resulting from the reduction in purchasing power due to the price rise. In panel (b) we isolate the former, by reporting

[14]These are the same price increases we use to estimate price elasticities in Section 3.3.

compensated prices elasticities, which we simulate adjusting each household's total expenditure so that their original consumption bundle remains affordable.[15] For all groups the mean compensated price elasticities are smaller in magnitude than Marshallian ones, implying that energy is a normal good. This can also be seen in panel (c), which shows mean expenditure elasticities. For all income and pre-shock energy consumption quintiles energy is a necessity, with lower income and higher pre-shock consumption groups having lower expenditure elasticities.[16]

Figure 5.1: *Slutsky decomposition*

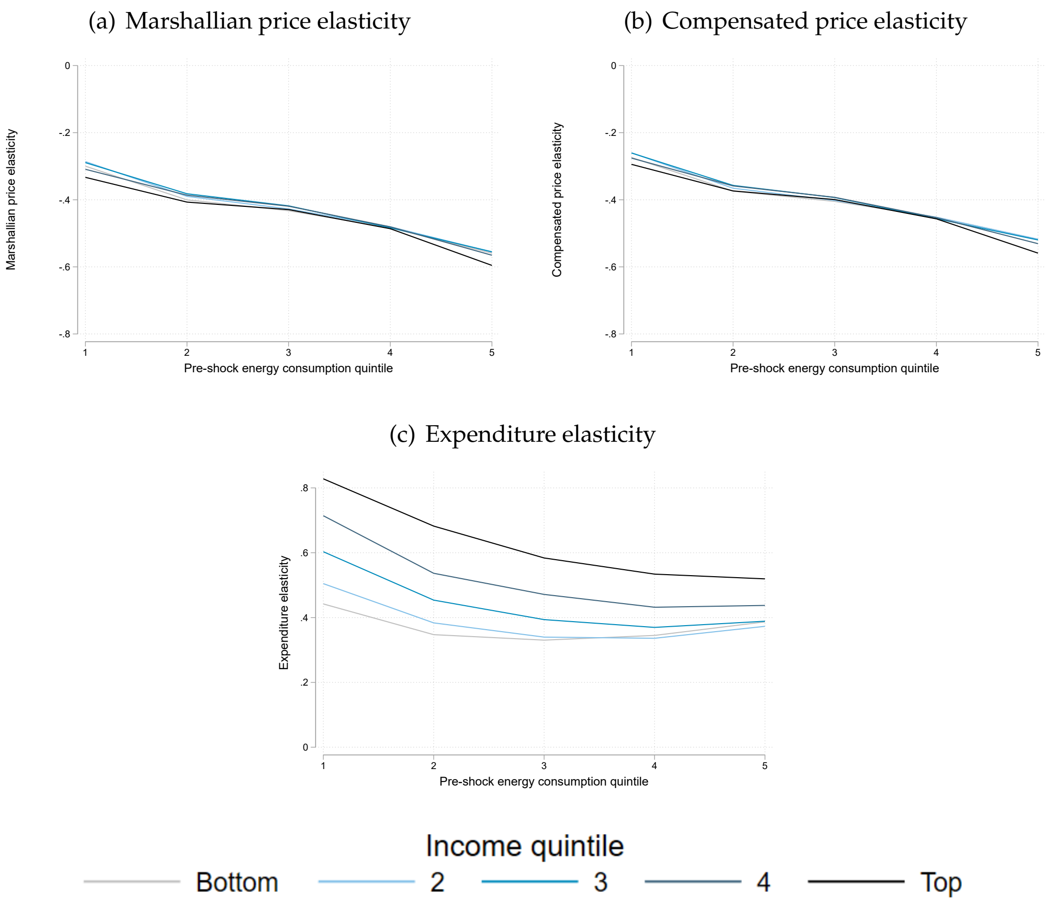

Notes: Figure reports price and expenditure elasticities simulated based on the demand estimates. Specifically, for observations (i, τ) during the period Oct. 2021-Sept. 2022 we compute the percent change in energy consumption associated with the price change from the pre-shock Apr.-Sep. 2021 period and divide it by the percent price change. Panel (a) reports the mean by pre-shock energy consumption and income quintiles. For panel (b) we repeat this exercise but adjust household level expenditure so that the pre-shock consumption bundle remains affordable. For panel (c) we construct the percent change in consumption due to this expenditure adjustment, divided by the percent expenditure change, and report means by pre-shock energy consumption and income quintiles.

[15] These are therefore Slutsky-compensated elasticities. The Hicksian-compensated elasticities (where a household's total expenditure is adjusted to keep their original utility level affordable) are very similar.

[16] The Marshallian, compensated and expenditure elasticities are related through the (discrete) Slutksy equation. Let ϵ^M, ϵ^C and ϵ^X denote the arc elasticities for a given consumer and w the post-shock budget shares, then: $\epsilon^M = \epsilon^C - w\epsilon^X$.

5.2 Incidence

Price shock and policy response

In October 2022 the UK energy regulator raised the energy price cap to a level 310% higher than its level in Summer 2021. The UK government responded with the Energy Price Guarantee, which entailed the introduction of a price subsidy (of 38.5% of the energy price) and a £67 per month bill rebate (which lasted until March 2023). In Figure 5.2 we summarize the effect of the price shock and policy intervention, relative to had prices remained at their pre-shock Summer 2021 level, across households. We consider three scenarios; had the government not introduced a subsidy or rebates ("No interverion"), had it introduced only the price subsidy ("Subsidy"), and the case of the observed policy response of a subsidy and rebate ("Subsidy+rebate").

Panels (a) and (b) report average effects by quintile of the income distribution; panel (a) shows the proportional impact on energy consumption and panel (b) reports the impact on household welfare (measured in £s). In the absence of government intervention energy consumption would have fallen, on average, by between 51%, for households in the bottom quintile of the income distribution, and 55%, for those in the top quintile. This gradient is driven by higher price elasticities among those with high (pre-shock) consumption, who, on average, tend to have higher incomes. The average monthly welfare losses from the price shock (in the absence of a policy response) range from £143 for the bottom income quintile to £223 for the top. This gradient is driven by those with higher incomes tending to have higher levels of energy consumption.

The price subsidy (on its own) reduces both the magnitude of energy consumption responses (which range from a fall of 33% to 37% moving from the bottom to top income quintiles) and welfare losses (which range from £78 to £125). The proportionate reduction in monetary welfare losses across income groups is approximately constant and around 55%, however, higher income groups gain more in £ terms (the subsidy lowers the average loss of those in top income quintile by £98, whereas it lowers it by £64 for those in the bottom quintile). The effect of adding the bill rebate to the subsidy is to further reduce the fall in energy consumption (to between 23% for the bottom income quintile and 29% for the top one) and the monetary welfare cost of the price shock. The rebate reduces each income groups' monetary loss by around £63 (meaning it has larger proportional effect for lower income households). There are two reasons why the £67 rebate lowers welfare losses by the lesser amount of £62. This first is because we measure losses based on equivalent variation defined at pre-shock prices (see equation (4.3)) and the purchasing power of a marginal £

is less after the energy price rise. The second is that the rebate distorts behavior through a labeling effect. We turn to the magnitude of this effect in the next section.

Figure 5.2: *Impact of policy on welfare and energy consumption*

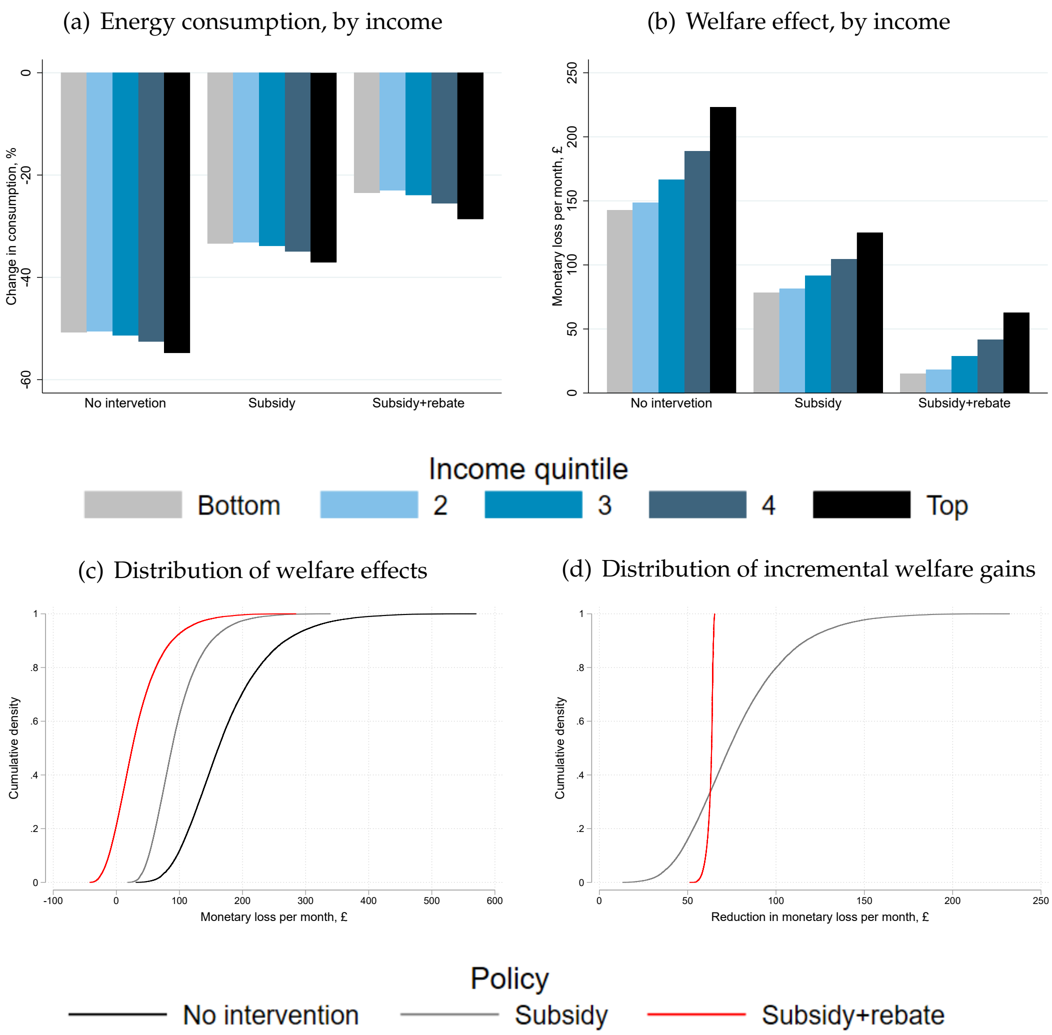

Notes: For observations (i, τ) during the period Oct. 2022-Mar. 2023 we compare energy consumption and equivalent variation (see equation (4.3)), relative to had the energy price been at its Apr.-Sep. 2021 level, under different policy responses. "No intervention" corresponds to no government policy response; "Subsidy" corresponds to a price subsidy alone; "Subsidy+rebate" corresponds to the implemented response of a price subsidy and energy bill rebate. Panels (a) and (b) show average change in energy consumption and equivalent variation by income quintiles. Panel (c) shows the empirical cumulative distribution function of equivalent variations. Panel (d) shows the empirical cumulative distribution function of incremental falls in equivalent variation from the subsidy and rebate.

Panels (a) and (b) of Figure 5.2 mask a great tell of heterogeneity across households that is not captured averages by income groups. In panel (c) we illustrate this by plotting the empirical cumulative density function across households in monetary welfare losses. In panel (d) we show the empirical cumulative density of the reduction in losses due to the subsidy and the incremental reductions in losses through adding the rebate (on top of the subsidy).

In the absence of government intervention the bottom and top percentiles of the loss distribution are £61 and £400. Under the subsidy and rebate, these percentiles fall to a welfare *gain* of £30 and a loss of £171. While there is wide heterogeneity in the extent to which the price subsidy reduces welfare losses, the incremental impact of the rebate on monetary losses is similar across all households. This reason for this is that the subsidy is tied to energy usage (which is highly heterogeneous), whereas the rebate is constant in value across households (though there is some variation in the extent of behavioral distortions due to the labeling effect). As we show in Section 5.3, the relative value the government places on limiting large losses, versus supporting low income households, has an important bearing on the optimal policy prescription.

Labeling effect

As we document in Section 3.4, the marginal propensity to consume energy out of rebates is substantially higher than that out of cash. In Figure 5.3 we quantity the strength of this labeling effect across quintiles of the income distribution on energy consumption responses (panel (a)) and welfare losses (panel (b)). In doing this we separate out two distinct channels. The first captures the distortionary impact on behavior of the labeling effect. We quantify this by replacing the labeled rebate with an unlabeled cash transfer of the same amount. The comparison of responses under "observed policy" and under this alternative policy ("-behavioral distortion") reflects the influence of labeling based behavioral distortions. The second channel captures the fiscal externality induced by the labeling effect through encouraging consumption of a subsidized good. We quantify this channel by replacing the rebate with an unlabeled transfer that is adjusted in value to expend the same resources as observed policy. The removal of the labeling induced fiscal externality allows for the transfer to be raised to £77, without the government expending more resources. Comparison of policy where the transfer equals the value of the rebate ("-behavioral distortion") and where it equals £77 ("-fiscal distortion") reflects the influence of this fiscal distortion.

Figure 5.3 shows that removing the behavioral distortion due to labeling leads to considerably larger energy consumption falls (33% larger, on average for the bottom income quintile, falling to 25% larger for those in the top quintile). The reduction welfare losses due to eliminating the labeling based choice distortion is relatively modest (around £2 per month on average). Conversely, the incremental effect of removing the labeling based fiscal distortion on energy consumption is relatively modest (as the rise in the transfer value from £67 to £77 exerts a relatively small income effect on energy consumption), yet it has a bigger impact on welfare, reducing welfare losses by around £10 per month on average).

Figure 5.3: *Impact of labeling distortion on welfare and energy consumption, by income*

(a) Energy consumption (b) Welfare effect

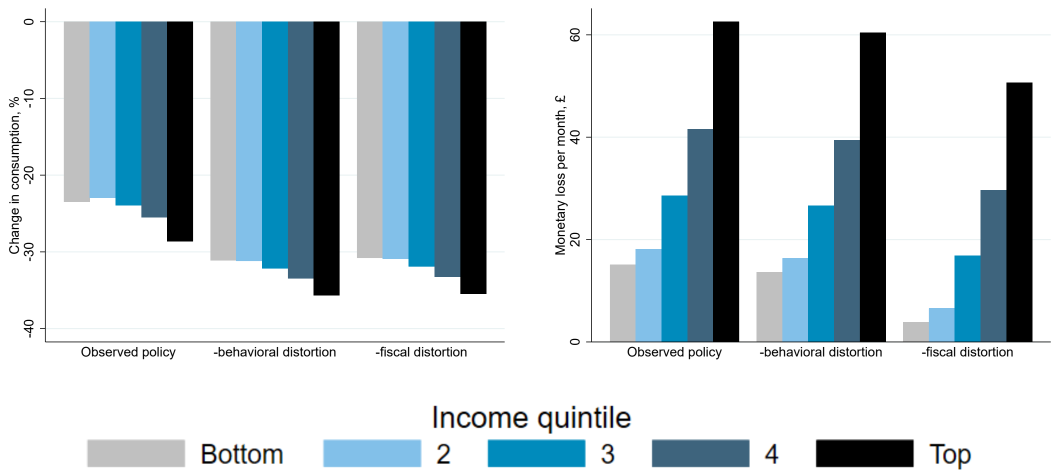

Notes: For observations (i, τ) during the period Oct. 2022-Mar. 2023 we compare energy consumption and equivalent variation (see equation (4.3)), relative to had the energy price been at its Apr.-Sep. 2021 level, under different policy responses. "Observed policy" corresponds to the implemented response of a price subsidy and energy bill rebate, "-behavioral distortion" corresponds to replacing the bill rebate with an unlabeled transfer of the same amount, "-fiscal disortion" corresponds to raising the unlabeled transfer value until total government resource expended are the same as under observed policy. We show average effects by income quintiles.

Behavioral responses

Our model allows us to quantity the welfare effects of observed policy (the simultaneous introduction of a subsidy and labeled transfer) and of alternative policies. In this section we use the period of elevated prices over April-September 2022 (prior to the introduction of the government's support package) to compare our model's welfare predictions with welfare approximations that do not entail specifying the form of preferences.

A first-order welfare approximation uses the envelope condition to approximate a household's equivalent variation with $EV^{1^{st}} \equiv \mathbb{e}_0 \Delta P_e$, where $\mathbb{e}_0$ is pre-shock energy consumption and ΔP_e is the energy price change. This approximation is exact for a marginal price change and has the advantage of using only information on pre-shock consumption. However, for finite price changes it ignores any substitution effects, which results in an upwards bias. A trapezoid approximation to the welfare cost, under the assumption of homothetic preferences, is given by $EV^T \equiv \frac{1}{2}(\mathbb{e}_0 + \mathbb{e}_1)\Delta P_e$. Unlike the first-order approximation, this requires data on post-shock energy consumption, $\mathbb{e}_1$, but this comes with the benefit of allowing for substitution responses.[17] The assumption of homothetic preferences has the advantage that

[17]This is made clear by re-writing the approximation $EV^T \equiv \mathbb{e}_0 \Delta P_e - \frac{1}{2}\Delta\mathbb{e}\Delta P_e$, where $\Delta\mathbb{e} \equiv \mathbb{e}_0 - \mathbb{e}_1$ is the reduction in energy consumption. Note this can also be re-written in terms of the Marshallian price elasticity, ϵ, using $\epsilon \approx -\frac{\Delta\mathbb{e}}{\Delta p}$ to obtain: $\mathbb{e}_0 \Delta P_e + \frac{1}{2}\epsilon\,(\Delta P_e)^2$.

the approximation depends on observed (this time pre- and post-shock consumption), but it comes at the cost of a non-homotheticity bias, which will be positive for a normal good.[18]

Figure 5.4: *Impact of subsidy, model compared to approximations*

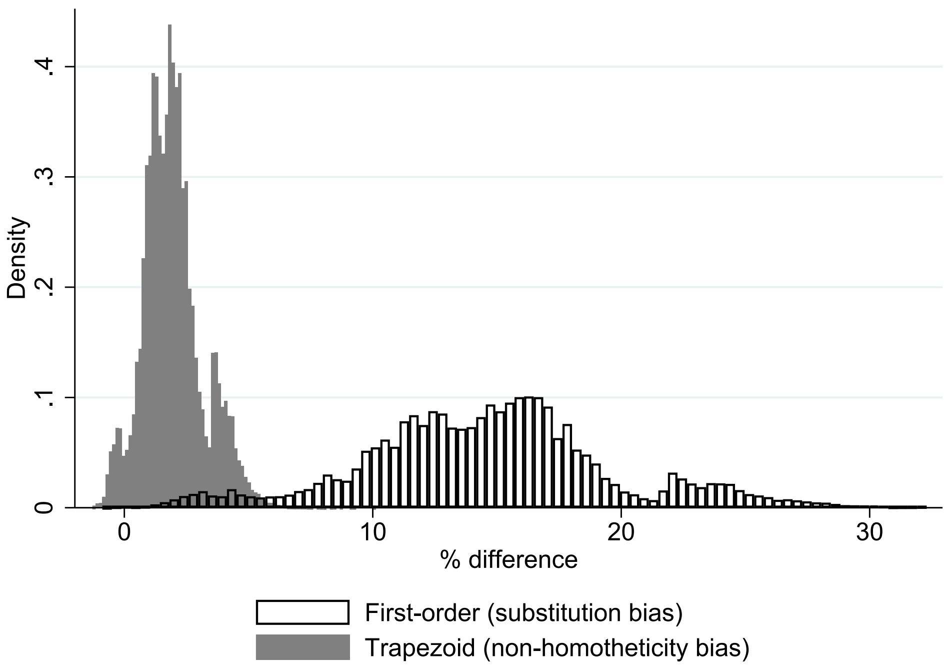

Notes: For observations (i, τ) during the period Apr. 2022-Sept. 2022 we compute the welfare loss associated with the energy price rise relative to its Apr.-Sept. 2021 level. We do this using the model based equivalent variation (see equation (4.3)) and a first-order and Trapezoid approximation. The figure shows histograms of the percent difference between the approximations and the model based estimate.

In Figure 5.4 we plot a histogram of the percent difference between both the first- and Trapezoid welfare approximations and the model generated equivalent variations. As expected both approximations over-predict welfare losses, relative to the model. The average difference for the first-order approximation is 15%. The relatively large differences are unsurprising, given the large price increases (and incentive to lower energy consumption) households faced. The Trapezoid approximation over-predicts welfare losses, relative to the model estimates, for over 95% of the sample, though the average difference is much small (under the 2%). The degree of over-prediction is strongly related to households' energy budget share – households with a pre-shock budget share below 5% have an average difference of 1% while this with a share above 15% have an average difference of 4%. This is consistent with energy being a normal good; energy price rises lead to a fall in energy de-

[18]A trapezoid approximation to equivalent variation that does not rely on homothetic preferences is $EV^{T-NH} \equiv \frac{1}{2}(\mathbb{e}_0^c + \mathbb{e}_1)\Delta P_e$ where $\mathbb{e}_0^c$ is the compensated pre-shock energy demand at the post-shock utility level. However, this approximation depends on an unobserved compensated demand. Note, for a price change rise, which means the pre-shock utility level is higher than the post-shock level, and when the good is normal good, $\mathbb{e}_0^c < \mathbb{e}_0$ and hence $EV^{T-NH} < EV^T$. This illustrates the upward bias in EV^T.

mand due to an income effect (which is larger for those with high budget shares) and this translates to an upwards non-homotheticity bias in the Trapezoid approximation (which is not present in our model, which does account for income effects). We view the fact that our model generated equivalent variations are similar to those based on the Trapezoid approximation, and that the differences between them are consistent with the bias embedded in the approximation, as further reassurance of the performance of our model.

5.3 Optimal compensation

During October 2022-March 2023 the UK government implement a specific combination of energy price subsidy and bill rebates. In this section we explore the optimal balance between the price subsidies and an unconditional transfer. We focus on an unlabeled transfer, since, as we have shown, labeling lowers welfare both by distorting choices and inducing a fiscal externality. However, we show in the appendix that our main conclusions are unaffected if we instead consider a labeled transfer.

The transfer and subsidy values are linked through the government's budget constraint (equation (4.4)). We can therefore treat the subsidy rate as the policy parameter, with the government's budget constraint determining the associated transfer value. We consider non-negative subsidy and transfer values, meaning the policy space in bounded by a subsidy of 0, which implies a transfer of £201 and an energy price increase of over 245%, and a subsidy of 57%, which entails no transfer and an energy price increase of 50%.

In Figure 5.5 we summarize how household-level monetary welfare losses vary with the subsidy value. In panel (a) we show the average loss in each percentile of the income distribution, and in panel (b) we report average losses within each percentile of the loss distribution. On average, households towards the bottom of the income distribution prefer a low subsidy and high transfer (with zero subsidy leading, on average, to welfare gains for the bottom half of the income distribution). On average, households at top of the distribution prefer are relatively large subsidy, for instance households in the top percentile of the income distribution see the lowest average welfare loss when the subsidy rate is 30%. However, panel (b) makes clear that the distribution of losses is mush more dispersed than the average losses by income percentile, and that lower subsidy values in particular, are associated with some households experiencing particularly large monetary losses.

Figure 5.5: *Household-level welfare effects, by subsidy value*

(a) By income (b) Distribution

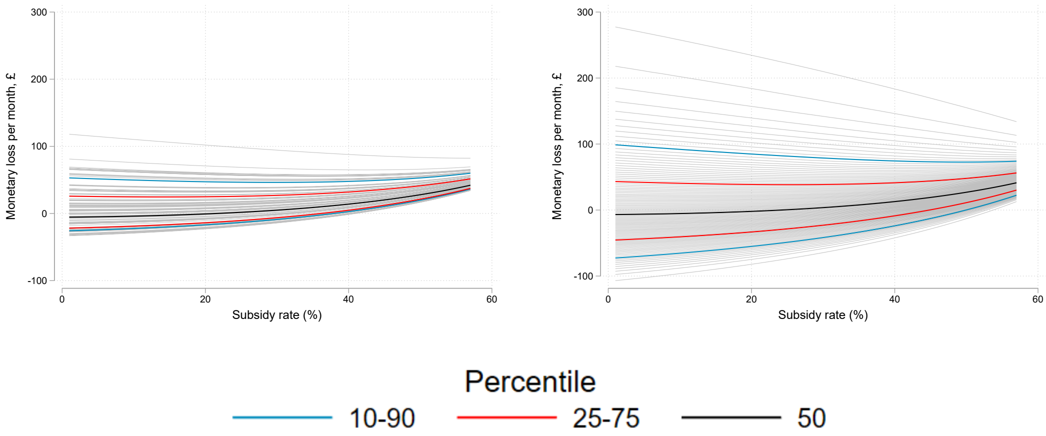

Notes: For observations (i, τ) during the period Oct. 2022-Mar. 2023 we simulate balanced-budget combinations of energy price subsidies and unlabeled transfers. Panel (a) shows how the average equivalent variation (see equation (4.3)) in each percentile of the income distribution varies with the subsidy rate. Panel (b) shows how average equivalent variation within each percentile of the equivalent variation distribution varies with the subsidy rate.

When setting policy the government must weight both the relative value of a £ loss across households of different incomes (vertical equity), and how the value of a £ loss varies with the overall loss size (which we refer to as horizontal equity). We encode these into the social welfare function (equation (4.5)) through vertical social welfare weights, $f(y_{p(i)})$, which depend on the income percentile the household belongs to, $y_{p(i)}$, and by allowing household level equivalent variation to enter through an increasing (weakly) convex function, $G(EV_i(\mathbb{s}, \mathbb{t}, 0))$. We specify these functions:

$$f(y_{p(i)}; \alpha) = \exp(-\alpha y_{p(i)}) \quad \alpha \geq 0$$

$$G(EV_i(\mathbb{s}, \mathbb{t}, 0); \psi) = \begin{cases} EV_i(\mathbb{s}, \mathbb{t}, 0) & \text{if } \psi = 0 \\ \frac{1}{\psi}\left[\exp(\psi EV_i(\mathbb{s}, \mathbb{t}, 0) - 1\right] & \psi > 0, \end{cases}$$

α captures the government's degree of concern for vertical equity (targeting policy at mitigating losses among the poor) and ψ captures the government's concern for horizontal equity (targeting policy at mitigating large losses). Figure 5.6 illustrates how these parameters determine the government's marginal valuation of a £ loss.[19] Panel (a) plots how the by-income social welfare weights, $f(y_{p(i)}; \alpha)$ vary across the income distribution for different values of α. When $\alpha = 0$, the government places equal weight on a given equivalent variation loss, regardless of the individual's income. For $\alpha = 0.2$, the government places 1.2 times

[19]For ease of presentation we reparameterize $f(y_{p(i)}; \alpha) = f(y_{p(i)}; \tilde{\alpha}(\alpha))$ and $G(EV_i(.); \psi) = G(EV_i(.); \tilde{\psi}(\psi))$ such that the values of α and ψ we consider are on the $[0, 1]$ interval. The reparameterization is: $\tilde{\alpha} = 10^{3\alpha - 1}$ and $\tilde{\psi} = \frac{1}{3.3} * 10^{0.3\psi - 1}$.

the weight on an equivalent variation loss for someone in the bottom percentile of the income distribution relative to the same monetary loss for an individual with median income. When α is 0.4 and 0.6 the relative weight is 2.2 and 22. When α is 0.8 only the bottom 10 percentiles have a relative weight bigger than 1/10 that of the bottom percentile, and when α is 1 this is true only of the bottom five percentiles. Panel (b) shows how the government's valuation of an additional £ loss varies with the overall size of the loss (conditional on income). This vertical marginal social welfare weight is given by $g(EV;\psi) \equiv G'(EV;\psi) = \exp(\psi EV)$ and is endogenous in the sense that the marginal welfare weight assigned to an individual will vary with government policy through its effect on equivalent variation. When $\psi = 0$ the government values the loss of the first £ equal to all subsequent £ losses. As ψ approaches 1 the government place a very large weight on avoiding large losses.

Figure 5.6: *Welfare weights*

(a) Vertical weights (b) Horizontal weights

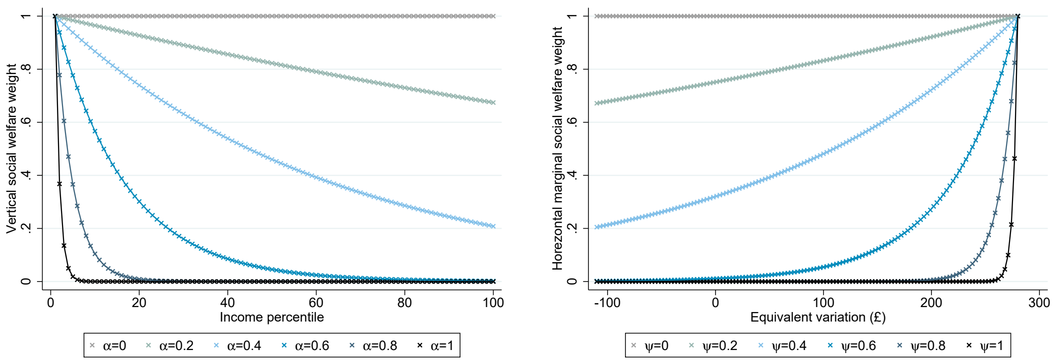

In Figure 5.7 we show how the optimal subsidy rate (panel (a)) and associated transfer value (panel (b)) vary with the parameters controlling vertical and horizontal equity concerns. The figures shows that the degree of horizontal equity concerns is key to driving optimal policy. For any degree of vertical equity concern, if concerns over horizontal equity are sufficiently low, optimal policy entails a relatively low subsidy and high transfer. The reason for this can be seen in Figure 5.5(b), which shows that for most of the equivalent variation loss distribution, even as high as the 75^{th} percentile, losses are increases in the subsidy value. It is only for the top of the loss distribution that that losses are decreasing in the subsidy level. As the government's concern for horizontal equity becomes stronger, its desire to limit losses of those most exposed to the price shocks becomes more influential in determining policy, and the optimal subsidy rate rises. It is for intermediate values of horizontal equity concern, that vertical equity concerns exert the most influence on optimal policy. For instance, when $\psi = 0.5$, the optimal subsidy rate ranges from X when the government weights households of different incomes equally ($\alpha = 0$), to Y, when the government places almost all vertical weight on the lowest income households ($\alpha = 1$). This pattern is driven by lower income households, benefiting most on average from as a high transfer as possible and higher income households, on average, prefering a moderate to high subsidy rate. In reality the UK government opted for a relatively high subsidy rate of 0.4, which is rationalizable in our framework only with relative strong concerns for horizontal equity.

Figure 5.7: *Optimal policy, by social preference parameters*

(a) Subsidy rate (b) Transfer

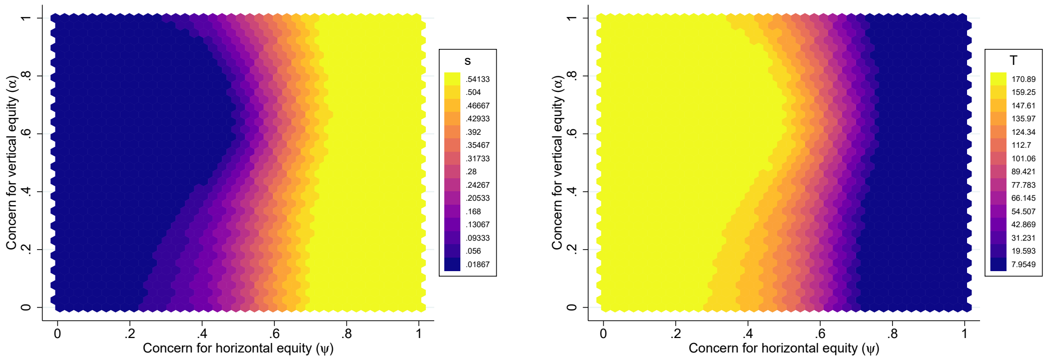

Notes: For each combination of α and ψ we compute the subsidy value (and associated transfer value, using the government budget constraint, equation (4.4) *that maximizes the weighted convex sum of equivalent variation (the social welfare function, equation (4.5)). Panel (a) describes how the optimal subsidy varies with α and ψ and panel (b) describes how the associated optimal transfer varies with α and ψ.*

6 Conclusion

To add

References

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

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*.

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

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.

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

Borenstein, S. and L. W. Davis (2012). The Equity and Efficiency of Two-Part Tariffs in U.S. Natural Gas Markets. *Journal of Law and Economics 55*(1), 75–128.

Dao, M., A. Dizioli, C. Jackson, D. Leigh, and P.-O. Gourinchas (2023). Unconventional Fiscal Policy in Times of High Inflation. *IMF Working Paper No. 2023/178*.

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

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

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*.

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

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

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

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.

Office for Budget Responsibility (2023). An international comparison of the cost of energy support packages. https://obr.uk/box/an-international-comparison-of-the-cost-of-energy-support-packages/.

Ofgem (2018). Record number of customers with small and medium sized suppliers. https://www.ofgem.gov.uk/publications/record-number-customers-small-and-medium-sized-suppliers.

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.

Online Appendix

A Setting and data

A.1 Further details on prices and payment methods

Figure A.1: *Energy price cap, energy price guarantee and cheapest available tariff for pre-payment consumers, 2019-2023*

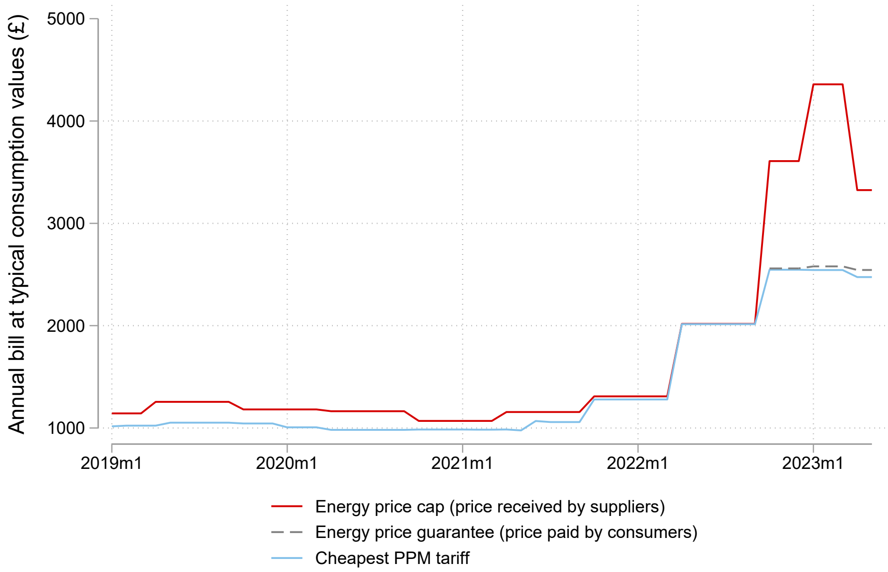

Notes: Data from Ofgem (2023). Figures are costs of an annual bill at 'typical' consumption values of gas and electricity (12,000kWh of gas and 2,900kWh of electricity) for a consumer paying for energy using prepayment. The figure shows the cheapest available prepayment tariff on the market (unlike Figure 2.2 which shows an average of the cheapest tariffs), as Ofgem doesn't publish a similar average for prepayment users. Only tariffs that are generally available to consumers are included when selecting the cheapest. The costs of the energy price cap are an average of costs for pre-payment consumers across regions in Great Britain.

A.2 Bank account data

Sample construction

Our focus is on consumers who are responsible for paying the households' energy bills. We start by selecting consumers with at least one account that records spending on energy. We keep all periods that longer than 6 months and contain an energy purchase at least every 150 days. We aggregate the data to the consumer-year-month level, summing spending and income across linked accounts. We construct a measure of non-durable spending, which comprises expenditures on: energy, groceries and other fast-moving consumer goods, vehicle fuel, childcare, recreation, personal services, transport services, other bills. Income is measured as money flowing into the account, excluding transactions tagged as bank transfers.

Table A.2 shows the number of consumers and observations in each year of the data. Over the full 2019-23 period, we have more than 11 million observations covering 464,000 consumers. 2019 is the first year of the ClearScore data so there are fewer users and observations in that year. In our analysis, we either control for consumer fixed effects to remove any sample compositional changes, or use a balanced panel that consists only of consumers present over the whole time period. This balanced sample contains 2.8 million observations for 68,000 consumers.

Table A.1: *Number of consumers and observations in the ClearScore data*

	Number of	
	Consumers	Observations
2019	146979	648085
2020	282799	2628777
2021	351567	3173922
2022	397791	3716342
2023	242768	963610
Full sample	464885	11130736
Balanced sample	68230	2817357

Notes: The first column shows the number of consumers and the second column the number of observations (consumer-year-months) present in each year of the sample and in total. The balanced sample is defined as consumers who are present between 2019 and 2023 (inclusive).

Representativeness

Figure A.2 compares the geographic and age composition of the ClearScore sample with the UK population. The ClearScore data overrepresents younger individuals and those living in the North West, and underrepresents those in London and the South East. We use the population data to construct weights based on region and 5-year age bands that we use to reweight the data to match the UK population.

Figure A.3 compares the distributions of income, non-durable spending, and energy spending in the ClearScore data (reweighted along age bands and region) and the Living Costs and Food Survey. The distributions align closely.

Figure A.2: *Age and geographic sample composition*

(a) Region

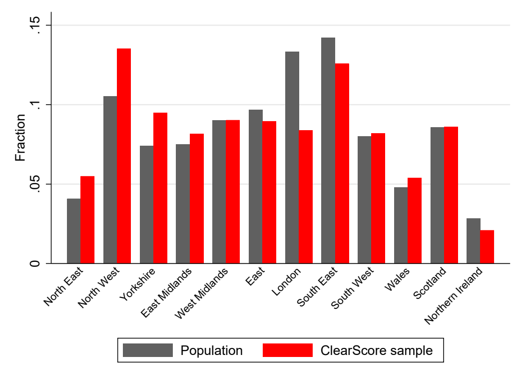

(b) Age

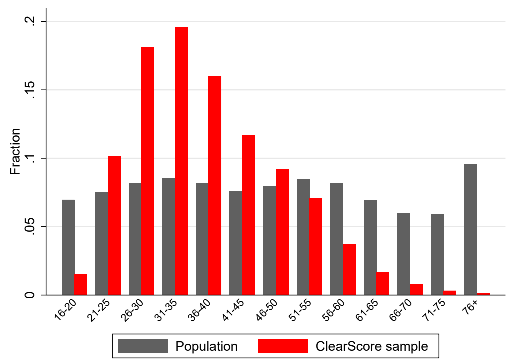

Notes: The left-hand panel shows the fraction of (i) users in the ClearScore sample and (ii) the UK population in each region. The right hand panel shows the fraction of (i) users in the ClearScore sample and (ii) the UK population in 5-year age bands (measured in 2021).

Figure A.3: *Comparison of spending and income with the Living Costs and Food Survey*

(a) Income

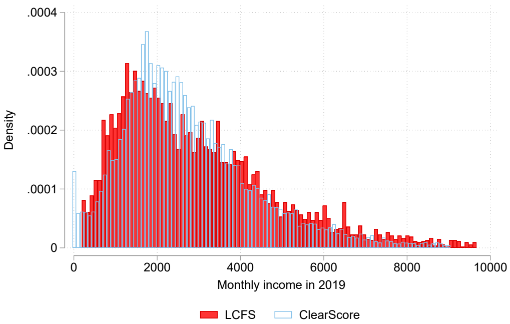

(b) Non-durable spending

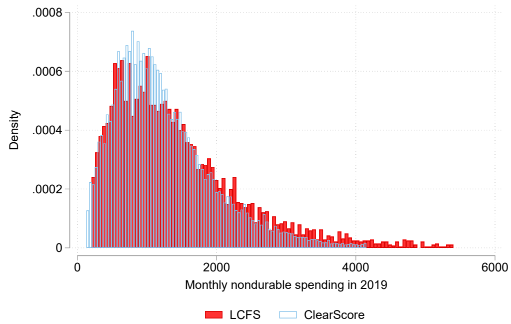

(c) Energy spending

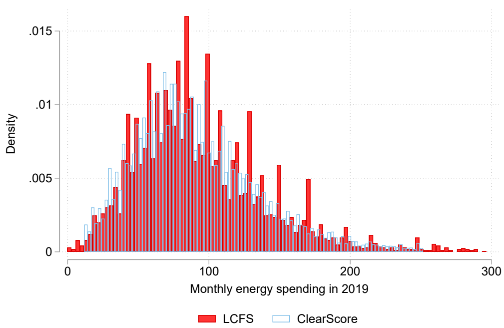

Notes: The figure compares the distribution of monthly income (top left), non-durable spending (top right), and energy spending (bottom) in the ClearScore data with that in the Living Costs and Food Survey. Distributions drawn for 2019; for the ClearScore data we take the average of each variable across the months they are present in the sample. ClearScore data are reweighted along age and region to match the UK population.

Identifying energy payment method and constructing weights

Consumers pay for their energy consumption in different ways. The most common method is to pay via direct debit with a fixed payment amount each month. We refer to consumers paying via this method as belonging to the "fixed payment sample". The remaining consumers pay in such a way that their monthly energy spending corresponds to their monthly energy usage (the spending of those in the fixed payment sample corresponds to their average (over the year) estimated energy use). We refer to these consumers as the "variable payment sample". This consists of people who pay by variable direct debit, pre-pay for their energy, or pay via standard credit (i.e. pay the bill they receive each period using card payment). We identify each of these payment methods in our data as follows.

The data contain information on the payment type of each transaction, such as: direct debit, card payment, or transfer (this is extracted from consumers' bank statements). We first identify all direct debit payments using this variable. We identify consumers on "variable" direct debits as those for whom the payment amount changes at least every other month, on average. Pre-payment consumers are identified as those who do not pay by direct debit, and whose payment amount is both less than £100 and ending in either a 5 or a 0 (i.e. a round number). Finally, standard credit consumers are those that do not pay by direct debit and do not satisfy the requirements for pre-pay described above.

Table A.2 shows the number observations (consumer-year-months) for our fixed and variable payment samples, where we split the variable payment sample into variable direct debit, pre-payment and standard credit. When reweighted to match the UK population along age and region, 62% of observations are fixed direct debit payments. The remaining 37% are variable payments, of which the majority are those that pre-pay for their energy.

Figure A.4 compares the distributions of the variable payment sample with the full sample. The distributions are similar, although the average income is a bit lower in the variable payment sample. This is to be expected since pre-payment tends to be more common among low-income consumers. We construct weights that adjust the variable sample to match the full sample on the basis of income decile.

Figure A.5 provides a validation of our sample definition by looking at variability of energy spending across the calendar year. The figure plots the average monthly energy spend in the pre-crisis period (2019-2020) for the fixed payment sample, the variable payment sample, and the subset of the variable payment sample that are direct debit consumers. It shows that, as we would expect, the fixed payment sample sees essentially no change in payment amount over the calendar year. The variable payment samples see higher payments in the winter and lower ones in the summer. This gives us confidence that our variable payment sample consists of consumers whose energy spending closely corresponds to their energy usage.

Table A.2: *Number of observations by payment type*

	Observations		
	No.	Unweighted %	Weighted %
Fixed payment	6380450	57.3	62.2
Variable payment	4750286	42.7	37.8
which consists of:			
Variable direct debit	395508	3.6	4.1
Pre-payment	3883193	34.9	30.1
Standard credit	471585	4.2	3.6

Notes: The first column shows the number of observations in the fixed and variable samples, with the latter further split into the three different payment types. The second column shows the unweighted share of observations in each sample and the third column shows the weighted share using age and region weights to match the UK population.

Figure A.4: *Comparison of spending and income between the full and variable payment samples*

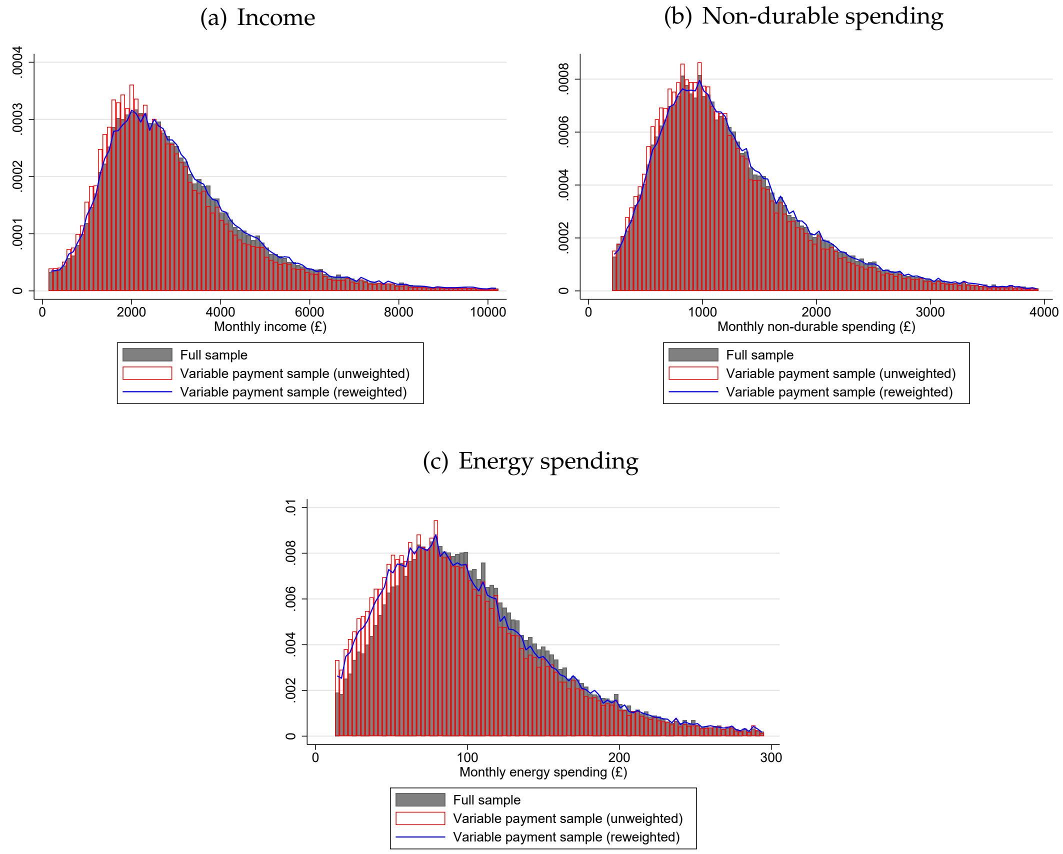

Notes: The figure compares the distribution of monthly income (top left), non-durable spending (top right), and energy spending (bottom) in the full sample with that of the variable payment sample. Distributions are drawn over the consumer's average of each variable over the period they are present in the sample. The grey and red bars show the distributions weighted by age and region to match the UK population. The blue line reweights the variable payment sample to match the income distribution of the full sample.

Figure A.5: *Seasonality of energy spending by payment type*

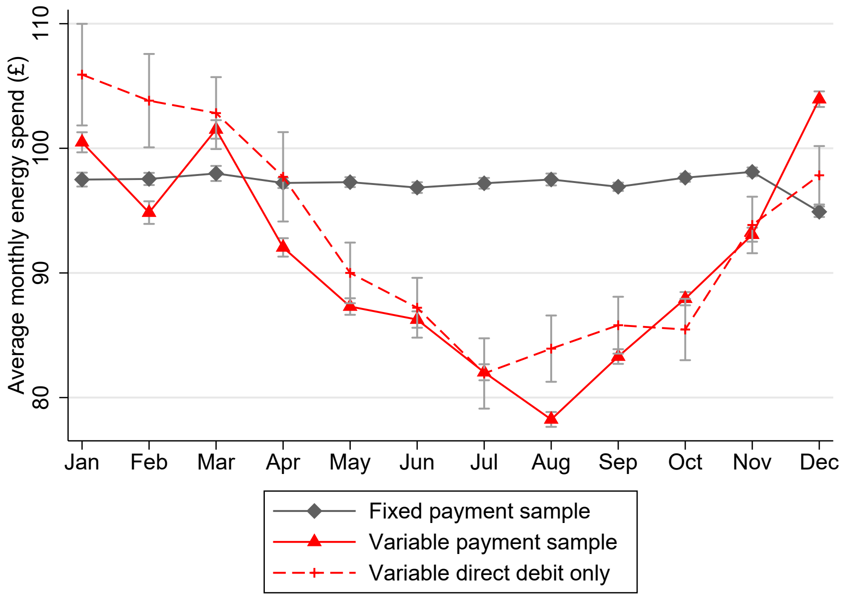

Notes: We estimate the average energy spending in each calendar month in 2019 and 2020 for three samples: the fixed payment sample, the variable payment sample and the subset of the variable payment sample that pay via direct debit. We control for user fixed effects. The regressions are weighted to match the age and region composition of the UK, and the variable sample is also weighted to match the full sample on the basis of income decile. 95% confidence intervals are shown.

B Additional event study results

Figure B.1: *Estimated seasonality in energy spending*

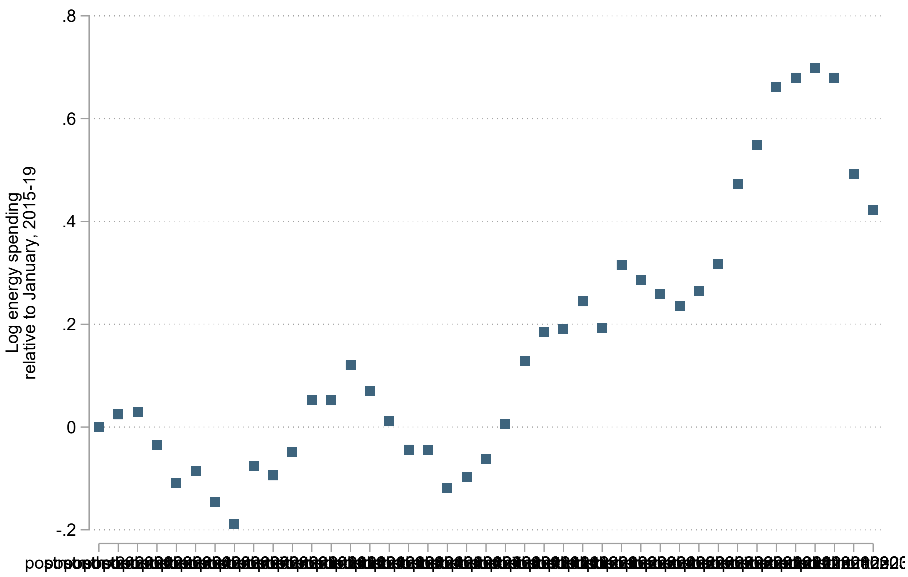

Notes: The figure shows the month effects in log energy spending estimated using the Moneydashboard data over the period 2015-19, controlling for the log CPI for energy.

Figure B.2: *Log energy spending and quantities by month, seasonal*

(a) Spending

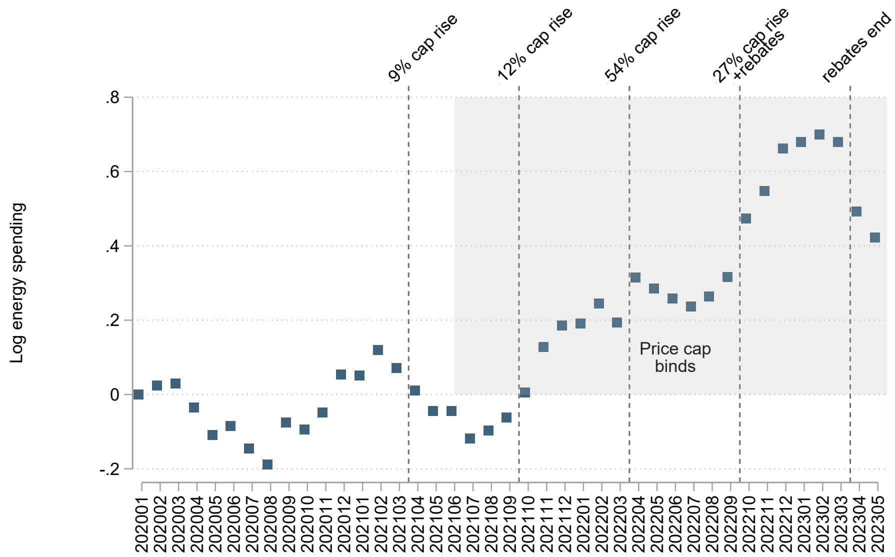

(b) Quantity

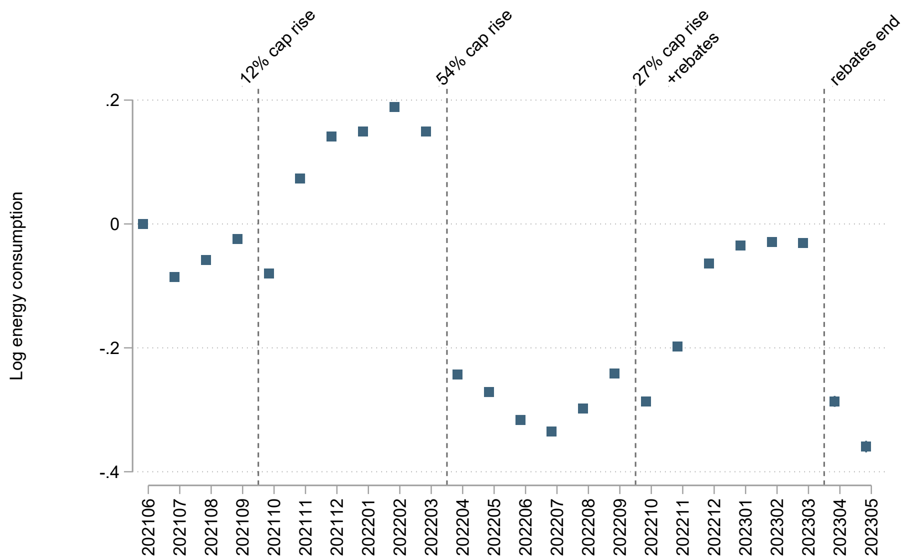

Notes: The top (bottom) figure shows total log energy spending (quantities) in each month, without controlling for seasonality or temperature fluctuations.

Figure B.3: *Log energy spending and quantities by month, comparison of full and variable sample*

(a) Spending (raw)

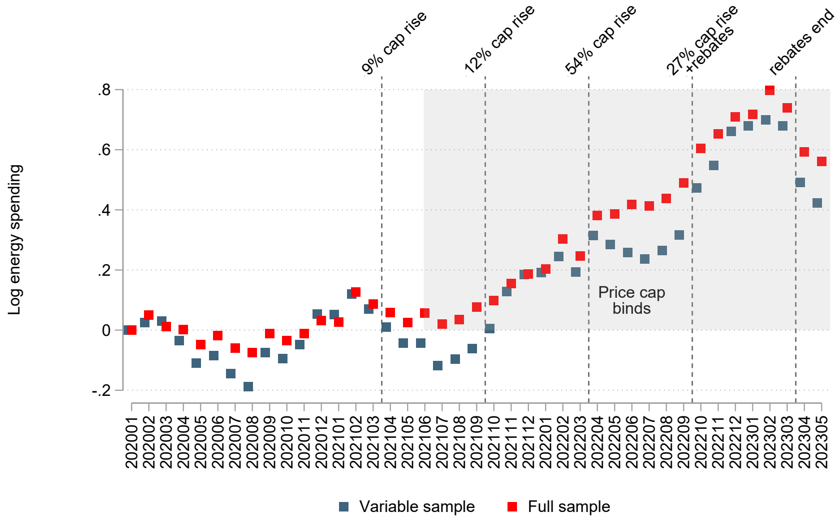

(b) Spending (deseasonalized)

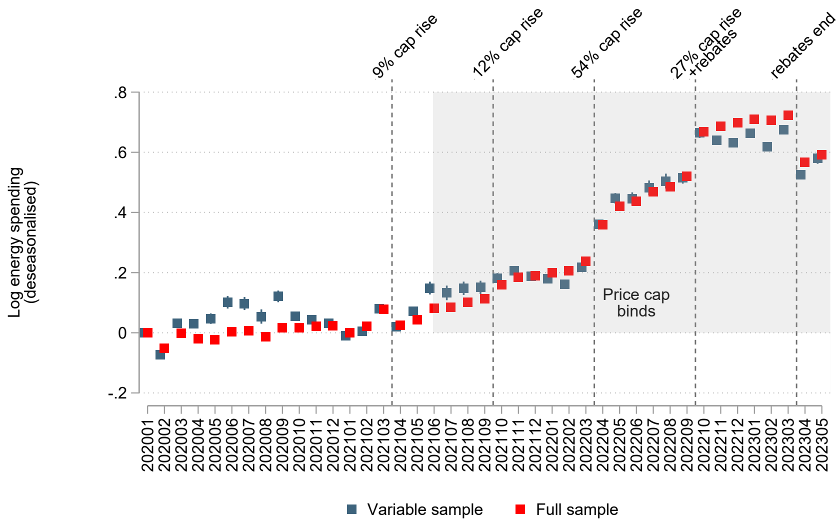

Notes: The top (bottom) left figure shows log energy spending (quantities) in each month before any seasonal adjustment, comparing the variable sample with estimates for the full Clearscore sample. The top (bottom) right figure shows deseasonalized log energy spending (quantities) in each month, comparing the variable sample with estimates for the full Clearscore sample.

C Details of energy demand model

Here we give further details on the consumer preference model underlying our demand equation (equation (4.6)).

C.1 Exact Affine Stone Index demand

Consider an individual's problem of allocating their total expenditure – denoted in logs by x – across J goods, where $\mathbf{p}$ and $\mathbf{w}$ are $J \times 1$ vectors of log prices and budget shares. The individual is characterized by an $L \times 1$ vector of observable demographic variables (we denote the set of L demographics $\mathcal{Z}$). $\boldsymbol{\epsilon}$ is a $J \times 1$ vector capturing unobserved heterogeneity.

The EASI demand system, proposed by Lewbel and Pendakur (2009), has a log-expenditure function of the form:

$$C(\mathbf{p}, u, \mathbf{z}, \boldsymbol{\epsilon}) = u + \mathbf{p}'\mathbf{m}(u, \mathbf{z}) + T(\mathbf{p}, \mathbf{z}) + S(\mathbf{p}, \mathbf{z})u + \mathbf{p}'\boldsymbol{\epsilon},$$

where $\mathbf{m}(u, \mathbf{z})$ is a J-vector valued function and $T(\mathbf{p}, \mathbf{z})$ and $S(\mathbf{p}, \mathbf{z})$ are single valued functions.

By Shephard's Lemma the Hicksian budget shares are:

$$\mathbf{w} = \mathbf{m}(u, \mathbf{z}) + \nabla_{\mathbf{p}} T(\mathbf{p}, \mathbf{z}) + \nabla_{\mathbf{p}} S(\mathbf{p}, \mathbf{z})u + \boldsymbol{\epsilon}$$

Substituting into the log-expenditure function and rearranging yields the implicit Marshallian budget shares,

$$\mathbf{w} = \mathbf{m}(y, \mathbf{z}) + \nabla_{\mathbf{p}} T(\mathbf{p}, \mathbf{z}) + \nabla_{\mathbf{p}} S(\mathbf{p}, \mathbf{z})y + \boldsymbol{\epsilon},$$

where implicit utility takes the form:

$$y = \frac{x - \mathbf{p}'\mathbf{w} - T(\mathbf{p}, \mathbf{z}) + \mathbf{p}'\nabla_{\mathbf{p}} T(\mathbf{p}, \mathbf{z})}{1 + S(\mathbf{p}, \mathbf{z}) - \mathbf{p}'\nabla_{\mathbf{p}} S(\mathbf{p}, \mathbf{z})},$$

and where by construction $y = u$. Implicit utility is an affine transformation of the Stone price index, $x - \mathbf{p}'\mathbf{w}$.

C.2 Our specification

We use the following functional forms:

$$\mathbf{m}(u,\mathbf{z}) = \mathbf{A}_0 + \mathbf{A}\mathbf{z}_0 + \sum_{r=1}^{R} \left(c_r + \mathbf{C}_r \mathbf{z}_2\right) u^r$$
$$T(\mathbf{p},\mathbf{z}) = \frac{1}{2}\mathbf{p}' \left(\mathbf{B}_0 + \sum_{l \in \mathcal{Z}_1} \mathbf{B}_l z_l\right) \mathbf{p}$$
$$S(\mathbf{p},\mathbf{z}) = \frac{1}{2}\mathbf{p}'\mathbf{D}\mathbf{p},$$

where $\mathbf{z}_1$, $\mathbf{z}_2$ and $\mathbf{z}_3$ are weak subsets of the demographic variables, $\mathcal{Z}_0, \mathcal{Z}_1, \mathcal{Z}_2 \subseteq \mathcal{Z}$.

The leads to Marshallian implicit budget shares of the form:

$$\mathbf{w} = \mathbf{A}_0 + \mathbf{A}\mathbf{z}_0 + \left(\mathbf{B}_0 + \sum_{l \in \mathcal{Z}_1} \mathbf{B}_l z_l\right) \mathbf{p} + \sum_{r=1}^{R} \left(c_r + \mathbf{C}_r \mathbf{z}_2\right) y^r + \mathbf{D}\mathbf{p}y + \boldsymbol{\epsilon},$$

where implicit utility is

$$y = \frac{x - \mathbf{p}'\mathbf{w} + \frac{1}{2}\mathbf{p}' \left(\mathbf{B}_0 + \sum_{l \in \mathcal{Z}_1} \mathbf{B}_l z_l\right) \mathbf{p}}{1 - \frac{1}{2}\mathbf{p}'\mathbf{D}\mathbf{p}}.$$

Consumer theory requires that the demand functions satisfy adding up and the expenditure function is homogeneous in prices; these properties are satisfied as long as $\mathbf{1}_J'\mathbf{A}_0 = 1$, $\mathbf{1}_J'\mathbf{c}_r = 0$ and $\mathbf{1}_J'\mathbf{C}_r = \mathbf{0}_{n(\mathcal{Z}_2)}$ for $r = 1, \ldots, R$, $\mathbf{1}_J'\mathbf{B}_0 = \mathbf{1}_J'\mathbf{B}_1 = \mathbf{1}_J'\mathbf{D} = \mathbf{0}_J$, $\mathbf{1}_J'\mathbf{A} = \mathbf{0}_{n(\mathcal{Z}_0)}$ and $\mathbf{1}_J'\boldsymbol{\epsilon} = 0$. It also requires symmetry of the Slutksy matrix of substitution effects, which is satisfied if $\mathbf{B}_0$, $\mathbf{B}_l$ and $\mathbf{D}$ are symmetric.

As our focus is on energy demand, we specify a two good demand system over energy and other nondurable consumption. In this case we can impose the adding up, homogeneity and Slutsky symmetry regularity conditions in the demand equation for one of the two goods (the energy equation) by specify the demand equation:

$$w^E = \left(A_0 + \sum_{l \in \mathcal{Z}_0} A_l z_l\right) + \left(B_0 + \sum_{l \in \mathcal{Z}_1} B_l z_l\right)(p^E - p^0) + \sum_{r=1}^{R} \left(c_r + \sum_{l \in \mathcal{Z}_2} C_{lr} z_l\right) y^r + D(p^E - p^0)y + \epsilon, \tag{C.1}$$

which is same as equation (4.6). It is not necessary to estimate demand for other nondurables as the only unknown parameter in the second equation is the constant term, which can be obtained as $A_0^O = 1 - A_0$.

Two other regularity conditions implied by consumer theory are that the expenditure function (which is given by $c(\mathbf{p}, u, \mathbf{z}, \boldsymbol{\epsilon}) = \exp\left(C(\mathbf{p}, u, \mathbf{z}, \boldsymbol{\epsilon})\right)$) is concave in prices and strictly increasing in utility. In our two good case, this implies the inequality restrictions:

$$(w^E)^2 - w^E + B_0 + \sum_{l \in \mathcal{Z}_2} B_l z_l + Dy < 0 \tag{C.2}$$

$$(p^E - p^O)\left(\sum_{r=1}^{R}\left(c_r + \sum_{l \in \mathcal{Z}_1} C_{lr} z_l\right) r y^{r-1} + \frac{1}{2} D(p^E - p^O)\right) > -1, \tag{C.3}$$

which we check are satisfied post estimation.