

How large is the energy savings potential in the UK?

Thiemo Fetzer, Ludovica Gazze, Menna Bishop

November 2022

No: 1437

Warwick Economics Research Papers

ISSN 2059-4283 (online)

ISSN 0083-7350 (print)

How large is the energy savings potential in the UK?

Thiemo Fetzer

Ludovica Gasse

Menna Bishop *

November 8, 2022

Abstract

Which households will be most affected by the energy price shock? How large are the energy, financial, and environmental benefits of improved energy efficiency of the British residential building stock? How do policies or interventions in price setting in energy markets affect these incentives? We develop a measurement and *ex-ante* modelling approach using granular property-level micro data representing around 50% of the English and Welsh building stock. This allows us to quantify the likely impact of recent energy price shocks on energy bills and how these bills would look like if energy savings measures were implemented. We find, on average, that the energy price shock acts as a form of progressive taxation hitting better-off regions more than poorer ones, in absolute terms. We estimate that on aggregate, 30% of energy consumption could be saved if buildings were upgraded to their highest energy efficiency standard. At market prices, these savings range between GBP 10 to 20 billion pounds per year with the highest energy savings largely concentrated in the wealthiest parts of the UK. However, current policies weaken incentives for households to invest in energy efficiency upgrades. Current policies, such as the energy price cap, appears to be very regressive. Alternative, more targeted policies, are cheaper, easily implementable and could align incentives better.

Keywords: ENERGY CRISIS, ECONOMIC HARDSHIP, POPULISM

*This working paper will develop over the coming months as the economic crisis unfolds. All three authors are based at the University of Warwick. We would like to thank Arun Advani, Jane Snape, Stephanie Seavers, Sheila Kiggins, Ben Moll, Carlo Perroni, Ben Lockwood, Mirko Draca, Jakob Schneebacher, Pedro CL Souza, Lily Shevchenko, Carlo Schwarz, and many others for helpful comments.

1 Introduction

Energy prices in the UK are projected to grow by over 600% between the first quarter of 2021 and the start of 2023, propelled by the post-pandemic economic recovery first, and Russia’s invasion of Ukraine later (IEA, 2022). Natural gas prices in particular saw drastic increases. In the absence of government intervention, shocks to wholesale energy prices will pass through to households’ heating bills, meaning price hikes have the propensity to create significant welfare costs. These costs will vary according to the energy requirements each household faces, which in turn are driven in part by the energy efficiency of their properties. For example, the change in the cost of heating a well-insulated home to a subsistence level during winter under higher prices is relatively small. The opposite is true for homes equipped with only single glazing and poor insulation. Thus, wholesale price shocks have the potential to create differential incentives to invest in insulation and energy efficiency measures (Houde and Myers, 2021). As such, the current crisis can inform on the potential implications of a carbon tax.

In this paper, we develop a measure of exposure to the energy crisis for properties in England and Wales. We do so by harnessing the Energy Performance Certificate (EPC) database, which includes over 22 million certificates detailing estimates of energy expenditure along with detailed and granular other public data on energy consumption. The underlying set of unique properties – around 15 million – represents a large share – at least 50% – of the English and Welsh residential building stock. Each EPC includes model-based energy consumption estimates that uses the physical characteristics of the building, a thermodynamic modelling approach, and a set of assumptions on occupation to derive energy consumption for space heating, hot water generation, and electrical light consumption. The publicly available EPCs report these estimates in GBP. We reconstruct estimates of energy consumption measured in kWh at the property level. We anchor this derived, theoretical measure of energy consumption with anonymized individual-level meter reading data along with granular spatial energy consumption data. This moment-matching based rescaling approach ensures that, at least to some extent, we capture local demographics and socio-economics as an important unobservable factor that may drive energy consumption above and beyond what can be accounted for in the

EPC data.

Equipped with these estimates of energy consumption by property, we can assess the impacts different energy price scenarios or policy interventions. We construct measures of the (likely) energy bills that households may face based on these consumption estimates that are specific to property characteristics and, in particular, its energy efficiency. For example, we model how energy bills change for different households as a result of changes in the UK's energy price cap.¹ In other words, we use this modelled consumption to define an *ex-ante* impact measure of likely policy shocks. In this sense, ours are true *intention-to-treat* estimates which abstract away from variation in energy expenditure driven by differences in behaviour across households, for example.

A desirable feature of the EPC data is that it contains recommendations for energy efficiency improvements in each property and estimates of the impact that these improvements would have on energy use. Thus, alongside our estimate of actual energy costs under various price scenarios, we are able to develop a measure of energy savings potential and price this savings potential under different policy scenarios. The difference between these actual and potential energy use is the measure of exposure to the price shock that we use throughout our analysis. It captures the extent of the energy price burden relative to the hypothetical scenario in which the residents upgraded their homes. This exercise allows us to quantify the hidden cost of underinvestment in energy efficiency in both monetary units and physical units of energy consumption. Such hidden cost is now very salient given recent dislocations in energy markets.

In addition to the *intention-to-treat* estimate that relies on physical attributes and assumed household characteristics, such as occupancy, we also rescale our estimates using both anonymized individual meter reading data that include a range of property characteristics, and granular subnational electricity and gas consumption data. This rescaling is primarily used to derive more accurate local area *ex-ante* impact measures that take into account not only the physical attributes of properties but also the actual level of demand along with average behaviours of residents under normal circumstances. These rescaled measures may, in a statistical sense, have

¹The energy cap sets the maximum price that energy suppliers are allowed to charge customers, and is chosen by regulator Ofgem.

superior out-of-sample predictive performance when thinking about actual energy consumption vis-a-vis hypothetical energy consumption.

We validate our estimates using data that we do not otherwise use to construct our measures – a form of out-of-sample validation. Further, we also attempt to bound statistically the degree to which we can explain variation in energy consumption across and within households over time. We find that our modelled energy consumption, at least at the level at which we have data for this comparison, comes close in terms of goodness-of-fit to what household-level empirical models achieve. Indeed, models of time-varying energy consumption at the property level, with property fixed-effects and a saturated set of time fixed effects achieve, at most, a goodness-of-fit of around 75%, the same goodness-of-fit we achieve with our data. Thus, even with ideal data that reflects area-level property characteristics and socio-economic characteristics of residents, our model may struggle to do a better job than that.

We then carry out four sets of interconnected descriptive analyses. Since we do not have granular socio-economic data matched to the individual properties, we have to carry out this descriptive analysis at a coarser spatial level – the Middle Layer Super Output Area – at which we can observe some socio-economic data from public sources. We begin by characterising which statistical areas in the UK would be, on average, more exposed from the energy price-shock. We use a best-subset selection approach to carry out this analysis in a way that is hands-off. We note that, in absolute term, regions that are quite privileged tend to be more exposed to the energy price shock. This is simply driven by the observation that well-off households tend to live in bigger, older, and more energy inefficient properties, while lower income households tend to live in flats or social housing which is more modern and inherently more energy efficient.

This implies that the current government intervention in the price setting mechanism in energy markets, through, the Energy Price Guarantee (EPG), is particularly regressive. As highlighted in Fetzner (2022), we note that subsidizing energy consumption by artificially lowering the energy price relative to market prices stands to benefit well-off households the most. We document the degree to which the energy price shock is progressive by focusing on two area-level characteristics: the median property price in an area, as well as a measure of total household income.

We note that the exposure to the energy price shock is strongly positively correlated with these measures highlighting that an energy price shock in essence, is a form of progressive taxation.

We further document in depth how, by subsidizing energy consumption, the UK's current EPG makes the energy price shock much less progressive. Using measures of income and house prices in an area, we estimate the degree of progressiveness of the energy price shock by evaluating the energy price shock at market prices, under EPG-moderated shock, as well as under an alternative two-tier energy tariff such that the lower tier of energy consumption is subsidized and the upper tier has each marginal unit priced at market or even higher rates. We note that, while the EPG weakens energy saving incentives for higher earners living in well-off areas, the two-tier tariff implies a similar degree of progressiveness as the market prices would imply with a more targeted relief towards lower income households.

In the third step we document where the highest energy savings potential lies in England. This highlights that the highest energy savings potential is found particularly in areas that are quite affluent. This surprising observation complements the earlier findings: not only is the energy price shock progressive, it especially hits households that, on the margin, have a lot of energy savings potential. We can only speculate on the reason why such energy efficiency upgrades have not taken place. A possible factor is the degree to which these investments would pay off: prior to the energy crisis, as energy prices were very low, the annual savings were projected to be quite little compared with the financial and non-financial burden, for example, in terms of opportunity cost.

Understanding the patterns of exposure to the ongoing energy crisis is an area of critical public interest and policy concern. The Ofgem figures often quoted in the press as estimates of household bills under successive price caps represent average values under an assumed level of consumption. In practice, the cap differs according to region, payment method, and meter type. Further variation in bills is introduced by different consumption needs of different households, for example due to the efficiency of their homes. We seek to tackle this oversimplification by deriving a measure of exposure to energy costs which is specific to the heating system and energy efficiency levels of an individual property. Even though it does not capture differences in consumption arising from differences in behaviours across

household, it is arguably more informative for forecasting the impact of bills on households.

In analysing which groups are the most exposed to energy shocks, we hope to identify potential candidates for the targeting of interventions that incentivise energy efficiency improvements. This is critical to tackling fuel poverty, which currently stands at 13.2% in England (Department for Business, Energy & Industrial Strategy, 2022). Furthermore, energy emissions due to home heating currently account for 14% of UK emissions (Institute for Government, 2021), and reducing these must be a first-order priority for a government committed to net-zero.

This work will contribute to several strands of literature. First, it will further our understanding of how interventions in energy markets affect the distributional implication of the energy crisis (see e.g. Harari et al., 2022; Bhattacharjee et al., 2022; Bachmann et al., 2022; Fetzer, 2022). A key unknown is to what extent households can adjust their energy consumption. Most literature in this domain finds relatively small short-term elasticities but potentially larger longer term elasticities. Labandeira et al. (2017) carry out a meta-analysis finding a short-term elasticity of -0.21 and a long term elasticity of -0.61. Moreover, changes in price were found to have the largest impact on gasoline consumption and the weakest impact on heating oil consumption. Second, we hope to contribute to the research investigating whether there is an energy efficiency gap, and if so what are its determinants Allcott and Greenstone, 2012; Gerarden et al., 2017; Christensen et al., 2021.

Third, several papers have analysed the appropriate targeting of policies to increase energy efficiency more in details. Zhang et al. (2012) reviews the literature on residential consumption demand and develop a model for archotyping UK energy consumers based on behaviour and property characteristics. They demonstrate how energy interventions must target the appropriate types of energy consumers to be effective. Similarly, Ahlrichs et al. (2022) document the spatial distribution of different building archetypes and detect a strong correlation between energy efficiency and socioeconomic factors. Gregório and Seixas (2017) focus on energy efficiency in historic town centres. They develop an index that characterises the energy renovation capacity of a community based on socioeconomic variables, property characteristics, and energy savings potential. Attari et al. (2010) conduct a survey of energy consumption and savings in the US. Their findings point to important defi-

ciencies in public understanding of energy conservation, including under- and over-estimation of the energy use and savings associated with different activities, underlining the potential benefits of interventions that educate households in this area. On the other hand, Myers (2019) finds that homebuyers are attentive to changes in fuel prices. In their study of the UK residential mortgage market, Guin et al. (2022) document that mortgages for properties with a higher level of energy efficiency are less often in payment arrears. Dalton and Fuerst (2018) conduct a meta-analysis of studies investigating the existence of a ‘green premium’, i.e. higher prices and rents for more energy efficient homes. Myers (2020) finds that in the US, landlords are not able to charge higher rents for energy-efficient units as tenants are not fully informed about energy savings, which depend on idiosyncratic behavior, for example. However, Myers et al. (2022) finds that mandatory disclosure of energy efficiency at sale increases investments and premia for energy-efficient homes.

In the following section, we describe how we arrive at a measure of the (likely) exposure to the energy price shock in England.

2 Developing an energy-price shock exposure measure

To model the likely exposure of household i to the energy price shock, we need to gain an understanding of baseline energy consumption. Energy consumption of household i in house p is driven by at least three factors:

$$E_{i,p} = f(\text{What}_p, \text{Who}_{i,p}, \text{How}_{i,p})$$

The What_p captures the type of property or building in which energy is consumed. The predominant sources of domestic energy use are space heating, hot water generation, room lighting, and appliances. Certain types of property, all else equal, will consume more energy across these uses because of their physical characteristics. For example, poorly insulated and draughty properties experience more heat loss. The second main factor, $\text{Who}_{i,p}$, represents the type of residents, for example in terms household size and composition. Two individuals living in a one-bedroom flat are likely to consume less energy in heating their property relative to two individuals living in separate one bedroom flats. Different socio-economic

backgrounds may also imply different levels of energy demand. The third factor, $\text{How}_{i,p}$, represents people's preferences. For example, people have different perceptions as to what constitutes a comfortable indoor temperature. In addition, whereas some households run dishwashers, others do dishes by hand. Moreover, these factors may interact nonlinearly: energy demand may be structurally higher in a poorly insulated property, but even more so if its individuals prefer a relatively high indoor temperature.

We start our work with a measure of energy consumption based on the What_p , i.e. the underlying characteristics of a property. We augment this exogenous measure with anonymised data on actual energy consumption at the individual property level, along with energy consumption aggregates at spatially granular levels using a moment-matching approach. In doing so, we are also able to incorporate the $\text{Who}_{i,p}$ and the $\text{How}_{i,p}$ into our measure of energy consumption, i.e. the patterns of energy consumption behaviour that exist in reality across households.

The data generation sequence is visually described in Appendix Figure A1.

In the next sections, we describe the underlying data and the generation of energy consumption estimates.

2.1 Deriving proxy measures for energy consumption

The first step in our data construction involves deriving energy consumption measures from energy performance certificate (EPC) data. EPCs provide buyers and tenants with information on the energy efficiency rating of residential properties as well as estimates of likely energy costs. EPCs also contain recommendations of measures to improve the properties' energy efficiency, including estimates of the costs and impact on energy demand of these measures. This information allows us to calculate a measure of *actual* and *potential* energy consumption by property. The potential measure captures an estimate of how much energy would be consumed, all else equal, if all recommended energy savings measures were implemented.

The requirement for properties to have an EPC was introduced in 2007 following the EU Directive on the energy performance of buildings (Department for Levelling Up, Housing & Communities, 2017). This requirement was initially applied just to homes for sale, but has since been extended to all domestic and commercial prop-

erties being sold, constructed, or rented (Department for Levelling Up, Housing & Communities, 2021). EPCs for all domestic and commercial buildings are available to download online from the national database of all registered EPCs.² In total, the database includes 22,179,913 current certificates for more than 15,621,668 unique properties across England and Wales. While we derive energy consumption measures for all certificates and the underlying properties, we focus in most exercises on slightly smaller subsets of the data that include only properties that use electricity and/or gas for space-heating and hot water generation. This amounts to 13,462,394 properties or around 51% of the English and Welsh residential building stock, as council tax data estimates the total number of residential properties to be 26,328,530.

A limitation of the EPC data is that certificates are valid for 10 years, meaning properties may have undergone changes, for example via the addition of an extension or insulation, that are not reflected in their most recent certificate. A second potential concern is that the EPC data may not be representative of the building stock that is not included. A comparison by the ONS of the EPC data representing 51% the building stock vis-a-vis the population of properties from the Valuation Office Agency (VOA) data, built for council tax purposes, suggests that the data are very similar on observables.³⁴ In fact, very naive reweighting of empirical moments from the EPC sample (in essence, just multiplying aggregate metrics by a factor of two) produces aggregate energy demand values that are very comparable with aggregate data. In terms of the potential energy savings, there are good reasons to believe that the properties that do not have an EPC rating may, in fact, have, on average higher improvement potential.⁵

The estimates of actual annual energy costs and potential annual energy costs

²Data are available here <https://epc.opendatacommunities.org/>.

³See Office of National Statistics, Energy efficiency of housing in England and Wales: 2021, <https://www.ons.gov.uk/peoplepopulationandcommunity/housing/articles/energyefficiencyofhousinginenglandandwales/2021>.

⁴Still, Department for Business, Energy & Industrial Strategy (2020) suggests that that the EPC database under-represents medium-sized properties and bungalows and over-represents smaller properties and flats.

⁵This assertion is based on the observation that, empirically, energy efficiency measures produce a larger energy savings effect among properties without an EPC certificate. For this work, see BEIS National Energy Efficiency Data-Framework (NEED): impact of measures data tables 2021, <https://www.gov.uk/government/statistics/national-energy-efficiency-data-framework-need-impact-of-measures-data-tables-2021>.

included in the public EPC data are expressed in terms of GBP and not in energy units (kWh). They are provided separately for space heating, water heating, and lighting. The Standard Assessment Procedure (SAP) sets out the methodology used to produce these estimates (BRE, 2014). We combine these estimates with price data to back out estimated energy consumption in kWh for space heating, water heating, and lighting, in effect reverse-engineering the SAP calculations.

We detail the technical approach in Appendix A. The end result, which is the start of further refinements that we outline in the next section, is a two vectors of energy demand proxies measured in kWh for each property p :

$$E_{p,act}^{EPC} = \{S_{p,act}, W_{p,act}, V_{p,act}\}$$

$$E_{p,pot}^{EPC} = \{S_{p,pot}, W_{p,pot}, V_{p,pot}\}$$

These capture the actual modelled energy demand and potential demand if properties were upgraded to their highest energy efficiency potential, respectively. The three main energy use functions that are modelled are for Space heating, hot Water generation, and electricity use for lighting indicated by V .

These breakdowns allow us to model energy bills as a function of which fuels are used for each energy use type. For example, homes heated via electricity will face different bills than those heated via gas. Note that these forms of energy use exclude the running of appliances like TVs, computers, cookers, washing machines, or dishwashers. The predominant driver of combined energy consumption modelled is space-heating.

We next describe how we refine and rescale the $E_{p,act}^{EPC}$ measure to match with other observed data on energy consumption.

2.2 Percentile-based rescaling

We further refine the EPC-derived measures using two percentile matching-based rescaling approaches. We leverage two sources of energy consumption data which are derived from meter readings. By doing so, we are able to anchor $E_{p,act}^{EPC}$ and $E_{p,pot}^{EPC}$ in data which reflects *emph*who lives in property p and *how* they live, both

of which affect the level of energy demand but are missing from our EPC-derived hypothetical consumption measures.

Anonimized individual property level consumption data. The first approach leverages anonymized energy data collected through the UK’s National Energy Efficiency Data Framework (NEED). This dataset includes gas and electricity meter reading data for 4 million properties. The sample is designed to be representative of domestic properties in England and Wales.⁶ The NEED data also include a range of property and area-level characteristics, such as property age and region, which can also be found in the EPC data, allowing for matching.

We rescale the EPC-derived energy consumption measures using these meter reading-based energy consumption data based on the distribution of energy consumption in each source. In particular, consumption estimates for properties in the EPC data which are in a given percentile of EPC-derived energy consumption will be rescaled using the consumption estimates for properties that same percentile of NEED-derived energy consumption. We do this separately for properties with different combinations of characteristics. While this first rescaling allows us to account for variation in real consumption behaviour driven by property characteristics, it may still exclude variation driven by local demographics. To incorporate the latter information, we employ a second rescaling method.

Local area consumption data. BEIS publishes energy consumption data down to the postcode level, excluding only postcodes that include fewer than five readings. The data include both mean and median consumption for electricity and gas.⁷ We then repeat the moment-matching approach described above, rescaling both the EPC and EPC-NEED augmented measures using the mean and median energy consumption values that correspond to a property’s postcode.

Appendix B describes this process in more detail. Importantly, we rescale both the EPC-derived actual modelled energy demand and potential energy demand.

⁶The data are available on <https://www.gov.uk/government/statistics/national-energy-efficiency-data-framework-need-anonymised-data-2021>.

⁷The data are available for electricity at <https://www.gov.uk/government/collections/sub-national-electricity-consumption-data> and for natural gas at <https://www.gov.uk/government/collections/sub-national-gas-consumption-data>.

This process leaves us with four measures proxying energy demand at the individual property p level:

$$\mathbf{E}_{p,\text{act}} = (E_{p,\text{act}}^{\text{EPC}}, E_{p,\text{act}}^{\text{NEED}}, E_{p,\text{act}}^{\text{Local}}, E_{p,\text{act}}^{\text{EPC,NEED,Local}})$$

Each measure is broken down into space heating, water heating, and lighting. Again, it is worth reiterating that space heating is the dominant factor in domestic energy use. For most of the analysis, we will leverage a simple *ensemble* average measure, $E_{p,\text{act}}^{\text{ensemble}}$, which is the unweighted average of each of these four measures.

We also have this same vector for potential consumption estimates. In the next section, we provide evidence on how these different rescaling methods impact the goodness-of-fit of our estimates with respect to real consumption data.

2.3 Illustrating the goodness-of-fit

We next describe how our derived property-level consumption measures fit actual energy consumption available at four levels of spatial aggregation: local authority district, MSOA, LSOA, and postcode level. Note that the postcode-level data is the same source used in our second rescaling approach. In Figure 2 we plot our four EPC-derived median energy consumption measures $\mathbf{E}_{p,\text{act}}$ against MSOA-level medians. Panel A highlights that, at the MSOA-level, the crude $E_{p,\text{act}}^{\text{EPC}}$ measure does a decent job at fitting the data but the relationship between the EPC hypothetical energy consumption and the actual reported median consumption is quite weak.

In Panel B we present the same scatterplot but now using EPC data rescaled using NEED microdata. This rescaling appears to produce a tighter fit. Panel C takes the measure in Panel A and applies the postcode-level rescaling. Not surprisingly, harnessing data on energy consumption from roughly one million postcodes improves the goodness-of-fit substantially. Panel D then presents the combined postcode and NEED rescaled measures.

Appendix Figures A2 and A3 present corresponding scatterplots for estimates of average and total energy consumption at the MSOA-level respectively. While the goodness-of-fit for mean and median values is very good, we note a mechanic underestimation of total energy consumption by our EPC-derived estimates. This underestimation can be explained by the fact that the EPC data covers only around

50-60% of properties.

These figures illustrate how our refinements help improve the goodness-of-fit of our energy consumption estimates. While the MSOA-level real energy consumption data has not been used for training, we nevertheless would expect a good fit because the data is a more aggregated view of the individual- and postcode-level data that we did use. Section 3 presents a proper validation exercise.

Figure 3 provides a more dense visual summary of the information contained in the scatterplots. It displays the R^2 of a set of regressions using MSOA-level energy consumption data: mean, median and total energy consumption against corresponding moments from our four derived measures $E_{p,act}$. This figure further sheds light on how the goodness-of-fit varies as a function of the coverage of the EPC data, i.e. the number of properties from a given MSOA that appear in the EPC data relative to the population of properties in that MSOA based on council tax records which captures an estimate of the stock of residential properties. As we move from left to right across each panel, the estimating sample expands. Not surprisingly, the goodness-of-fit is quite low when we restrict our sample to areas where the EPC data capture only a small fraction of the total residential stock. The fit improves rapidly when we employ more representative data. Figure 3 further emphasises that rescaling improves goodness-of-fit across each of the three moments that we consider.

Interestingly, the goodness-of-fit appears to peak at around 75% across our four energy consumption measures. In Section 3.1 we find a similar maximal goodness-of-fit measure when studying individual property-level actual energy consumption data. The result using property-level data highlights that, possibly, without more granular detailed time-varying data on who and how people live in a property, it may be difficult to achieve a higher goodness-of-fit than the 75% reported here. We next describe how we use the energy consumption estimates to arrive at estimates of energy bills under different policy and price scenarios.

2.4 Estimating energy bills

With the above vectors of energy demand proxies broken down by respective energy use functions, along with information on which fuels are used to heat properties

and the appropriate energy tariff, we can derive estimates of household energy bills under different price and policy scenarios.

We consider the following scenarios:

1. Energy price cap.

The energy cap sets the maximum price that energy suppliers are allowed to charge customers, and is chosen by regulator Ofgem for gas and electricity prices to reflect the costs of supplying energy and to allow modest profits (Ofgem, 2022a). The cap has been updated every 6 months since its introduction in January 2019, but from October 2022 will be updated on a quarterly basis. The price cap was originally conceived to protect inattentive consumers from being charged unfair rates. In its early years, some energy contracts on the market were cheaper than the cap, but since the summer of 2021 the cap has been the cheapest rate available. This phenomenon is due to price increases between the time at which the price cap is set and the time at which it comes into effect (as of October 2022, this gap has been shortened from two months to 25 working days) (Ofgem, 2022b). As such, the cap has been a more accurate reflection of the prices faced by households in recent months than in previous years. Our study incorporates price cap values from October 2021 and October 2022.

2. Energy Price Guarantee.

In September 2022, the UK government announced the Energy Price Guarantee programme as a response to the ongoing energy crisis. The EPG reduces the maximum per unit rate below the level of the October 2022 price cap in an attempt to limit the average household energy bill to around £2,500. As discussed in Fetzer (2022), the standing charge is maintained at the level of the October 2022 price cap.

3. Historical average energy prices.

The Department for Business, Energy and Industrial Strategy (BEIS) publishes data on average gas and electricity prices for 2010-2021.⁸ These data are par-

⁸Data are available here <https://www.gov.uk/government/statistical-data-sets/annual-domestic-energy-price-statistics>

ticularly valuable for estimating energy bills pre-2019, when the energy price cap had not yet been introduced.

4. Two-tier tariff.

This is an alternative policy proposal to the energy-price guarantee that is discussed in more detail in Fetzer (2022). It consists of a two-tier tariff wherein the standing charge would be fixed at the level of the October 2021 price cap, as would unit prices for the first 9,500 kWh of natural gas consumption and the first 2,500 kWh of electricity consumption. As 50% of UK households consume less than 12,100 kWh of natural gas and 2,900 kWh of electricity, this would drastically limit energy price increases for the bulk of households.⁹ The second tier of the energy tariff would be set at steeper levels which could be aligned with the EPG. For example, a second tier unit price of 20 pence per kWh for natural gas and 60 pence per kWh for electricity, together with the first tier described above, would have a similar cost to the government as the EPG. This would offer much more targeted support without undermining the incentive to save energy created by higher unit prices.

For vectors of our property-level energy consumption estimates $E_{p,act}$ and $E_{p,pot}$, we are therefore able to produce vectors of spending estimates. For example, for the preferred ensemble average energy consumption estimate $E_{p,act}^{ensemble}$, we produce the following four spending estimates:

$$C_{p,act}^{ensemble} = (C_{p,act,21-10}^{ensemble}, C_{p,act,22-10}^{ensemble}, C_{p,act,EPG}^{ensemble}, C_{p,act,Two-tier}^{ensemble})$$

This allows us to measure changes in the energy bills under different price scenarios and policy interventions at the individual property-level. We will study statistical patterns in these changes in Section 4 to characterise how these measures affect households with different characteristics. We next present some validation and bounding exercises.

⁹See <https://www.gov.uk/government/statistics/national-energy-efficiency-data-framework-need-consumption-data-tables-2021>.

3 Validating and bounding the measures

Our attempt to match and rescale the EPC-derived data using actual consumption data aims to ensure that, at the property level, we achieve a measure of baseline actual consumption that incorporates, at least to some extent, the likely distribution of the resident population. This matching ensures that we are more likely to achieve a good simulated actual exposure measure to the energy price shock.

A natural mismatch between energy consumption across properties and the EPC-derived measures of consumption can arise because properties are not inhabited by the number of people that are assumed in the model used to produce the underlying EPC data. For example, a two bedroom house is typically assumed, in the EPC model, to be inhabited by more than one person. If, in fact, the property is only inhabited by one person instead of the assumed number, this may imply a lower level in energy consumption, which would naturally yield a poorer match with the actual consumption data.

3.1 Attempting to bound the what, who, and how

It is inherently challenging to separate the drivers of energy consumption. Naturally, there is an interaction across at least three factors:

$$E_{i,p} = f(\text{What}_p, \text{Who}_{i,p}, \text{How}_{i,p})$$

We leverage anonymized meter-reading data from England and Wales at the property-level to bound the extent to which we can explain variation in energy use between the What_p . In the NEED anonymized microdata we observe a range of property characteristics that could drive variation in energy demand.¹⁰ We characterise the extent to which we can capture variation in the observed energy consumption data across properties (or households) saturating simple linear regression specifications of the form

¹⁰The data are a stratified random sample from the population of properties. Unfortunately, BEIS does not make the sampling weights available for each strata, which means we can not correct for the respective under- and oversampling. We have requested this information but are still awaiting a response.

$$E_{i,p,t} = x_{i,p,t} \times \beta + \epsilon_{i,p,t}$$

The features in $x_{i,p,t}$ include:

- Property characteristics: property type (six categories) such as detached, semi-detached, or flat; property age band (four bands) capturing the date range when a property was built; an indicator for whether gas is the main heating fuel; floor area bins (five categories) ranging from less than 50 square meters to over 200 square meters. Further, we also have measures capturing whether a property has had some energy efficiency measures such as cavity wall insulation or loft insulation installed.
- (Mild) socio-economics: quintiles of the English- and Welsh indices of Multiple Deprivation from 2019. That is, for every property we know the location in terms of the region (there are 10 regions that make up England and Wales), along with whether a property falls into a region that is in a specific quintile of the English- or Welsh deprivation ranking.

In addition, we have a property identifier which will serve as a *property fixed effect* in some specification as the most demanding, but also least informative, way of trying to absorb property- and time-invariant resident-specific observable and unobservable characteristics.

To allow for potential non-linear interactions between different property characteristics driving energy consumption such as an interaction between floor area and property age, we construct a measure that captures the unique combinations of each of the property characteristics. That is, each unique combination of the features that are called property characteristics is identifying an own *group* which we refer to as Property. There are 9,846 unique combinations in the data capturing different combinations of these characteristics.

We do the same for the (mild) socioeconomic indices capturing the council tax band and the IMD measure. These are the only measures that provide some information about the area or the property that can characterise who may live there due to typical patterns of socio-economic segregation that are common in residential

choice. As with the property characteristics, we combine these into a group variable that captures all unique combinations that exist in the microdata and we refer to this as *Socioeconomics*.

Lastly, we further interact each of these variables with year fixed effects to allow for non-linear interactions of property characteristics and possibly year-on-year unobservable shocks.

Results. We present the results from this characterisation exercise by plotting the estimated *adjusted R²* in Figure 4 showing both combined gas and electricity, along with gas and electricity consumption separately. We note that property characteristics and or the mild socio-economic characteristics can, at most, capture 50% of the variation in energy consumption. In particular, electricity consumption appears much more idiosyncratic compared to natural gas consumption. This finding is not surprising given that demand for natural gas is predominantly driven by space-heating and hot-water generation which do not vary much with household composition and tastes compared to electricity consumption. We note that the adjusted *R²* can reach up to around 75% in the specifications with property fixed-effects.

Interpretation. The results of this characterisation exercise suggest that property characteristics alone cannot explain much of the variation in energy demand. At most, characteristics can explain around 50% of the variation in residential energy use. Moreover, the maximal goodness-of-fit attainable appears to be bounded around 75%, obtained when we control for property fixed-effects, which may capture some of the underlying unobservable socio-economics factors (who lives there) along with behavioural factors (how do they live).

Naturally property fixed effects absorb a lot of meaningful variation, but, the estimates of the fixed effects are not informative by themselves. We deploy them here in a statistical sense as a most demanding way to measure property-specific energy demand – as a substitute for the factors that we do not have in the data – specifically on the *what, who* and *how*.

Interestingly, our validation exercise of the energy demand measure that we constructed at the property level produces a goodness-of-fit vis-a-vis statistical moments such as the mean and in particular, the median, that also achieves a similar

adjusted R^2 of around 75% – coinciding with the upper bound derived here from microdata. This renders us with some further confidence that our energy demand measures can do a good job at picking up variation in the data.

We next consider additional validation exercises comparing empirical moments that were not used in the training step.

3.2 Out-of-sample validation comparing empirical moments

For some local authority level, we have data that provide pairwise measures of both the mean and median electricity and gas consumption by district and by property-type and floor area band. We did not use these data in the rescaling, but, if our rescaling approaches were successful, we should do a good job at matching these empirical moments. Indeed, our rescaling attempts to preserve both the variation within an area across the different property characteristics through the NEED rescaling and the variation in energy demand that is location-specific due to the local resident population through the rescaling based on granular local area consumption data. Thus, this can be seen as an out-of-sample validation approach.

We leverage district-level gas and electricity consumption data measuring the mean and median consumption in a district by floor area group and property type.¹¹ Note that we could have used these data to rescale our measures as well, but since we attempt to create more granular view, this approach seemed too coarse.

We do so by estimating

$$E_{d,c}^{BEIS} = \alpha + \beta \times E_{d,c,act}^j + x_{d,p} \times \nu + \epsilon_d$$

where $E_{d,c}^{BEIS}$ stands for the median or mean energy consumption of a property with a characteristic c in district d . That is, these are derived moments in the actual energy consumption data. We construct the corresponding moment in aggregated form, either the median or mean, at the district by property characteristic c based on the property-level microdata described in the previous section.

Our attention will be on the estimated coefficient β . In the regressions that exclude other control variables or shifters, this coefficient should be close to one if

¹¹Local authority table, England and Wales, <https://www.gov.uk/government/statistics/national-energy-efficiency-data-framework-need-consumption-data-tables-2021>.

there was a near one-to-one mapping of the EPC-derived consumption measures and the actually observed consumption data. A second focus will be on the combined R^2 of these regressions. We would hope this R^2 is close to one which would indicate that, on average, our approach to measure hypothetical consumption captures the variation in actual consumption quite well.

Lastly, we also are interested in whether, after absorbing district fixed effects and property characteristics, our EPC-derived consumption measure $E_{d,c,act}^j$ carries signal over and above area and property characteristic specific idiosyncracies. In other words, this exercise tests whether our two-way rescaling approach achieves its goal.

Unconditional fit. In Figure 5 we present the simple unconditional scatterplot of the two datasets. On the horizontal axis we plot the EPC-derived median energy ensemble predicted energy consumption at the district by floor area combination level in Panel A, and the district by property type level in Panel B. The vertical axis plots the actual observed median consumption for 2019.

We observe a tight fit even in the unconditional regressions. We next explore this validation more systematically.

Conditional fit. We first compare the BEIS empirical moments of the median and the mean electricity and gas consumption with the five measures we construct based on the EPC measures. These are presented in Table 1.

Across the panels we vary whether we add additional control variables. Technically, if the data is a good fit of actual energy consumption and, in particular, of the variation in energy consumption both within and between districts and property types, we should expect to see a regression producing an R^2 that is close to one. Further, in the regressions that do not control for either floor area band or district fixed effects, we would expect the coefficient to be close to one. Statistically, such a finding suggests a near one-to-one mapping in terms of the underlying variation. A similar picture emerges when studying the district-by-property-type empirical moments presented in Table 2.

We conclude that our empirical approach calibrates the EPC-derived data to actual consumption data well, which allows us to provide a richer view of the likely

impact of the energy price shock. We next carry out the distributional analysis under different price scenarios using data aggregated to the MSOA level.

4 Empirical analysis

We next describe the empirical analysis aimed at answering four questions:

- Where will the energy price shock hit the hardest?
- Which places stand to benefit most from the two policy alternatives considered (EPG vis-a-vis two-tier tariff)?
- Where is the highest energy savings potential?
- How does present policy affect energy savings investment incentives?

Naturally, we would want carry out such analysis at the individual property level, but we do not have matched household data for the individual properties.¹² As a result, we are left with aggregated cuts of our data. In this exercise, we focus on England only.¹³

4.1 Where will the energy price shock hit the hardest?

We begin by statistically characterising which areas in England would be most hit by the energy price shock absent policy intervention. In this counterfactual world, Ofgem would have implemented the price cap it had announced for October 2022. We compute the overall increase in the energy bills between the October 2021 to the October 2022 price cap as:

$$\Delta C_{p,act,22-21}^{ensemble} = C_{p,act,22-10}^{ensemble} - C_{p,act,21-10}^{ensemble} \quad (1)$$

¹²We are in discussion with the Office of National Statistics (ONS) and the Understanding Society (USOC) household panel data owners to carry out a merge of our data.

¹³We focus on England because decentralisation efforts in the UK create barriers to access data from Wales and Scotland in particular, as the nations have devolved part of the statistical functions away from the Office of National Statistics. The Welsh and Scottish websites do not use the same access tools and less data are available.

We then consider a vector of socio-economic variables and perform best-subset selection (BSS) as a simple machine learning method to characterise which parts of England would be most hit by the energy price shock in the absence of government intervention. In uncovering which combinations of attributes have the largest explanatory power for the energy shock, BSS enables us to understand the patterns of vulnerability to said shocks across England. This is important for understanding how to optimally design and target policies to provide relief to households.

We study the following specification at the MSOA level

$$\Delta C_{m,act,22-21}^{ensemble} = \alpha_{d(m)} + x_m \times \beta + \epsilon_d$$

where we use the BSS algorithm to include ever more sets of control variables. We identify the optimal model as that which minimises the Akaike Information Criterion (AIC). The AIC measures the quality of a model by weighing up its goodness-of-fit against its simplicity, i.e. the number of features that are included in the statistical model.

Throughout, we absorb district-level fixed effects as a set of features not included in the BSS selection approach. This ensures that we only focus on within district variation across MSOAs. The features in variable x_m have been built from various public data sources for England such as the Census or labor market and other statistics collected by the Office of National Statistics (ONS) and are described in more detail in Appendix E.

Results. We present the results in Table 3. As we move across columns, the BSS method adds more variables. The order by which variables are added is informative about the underlying signal they carry. We note that the first variable that gets added and which stays in the model throughout is the median house price of an MSOA: the higher the house prices in an area, the higher is the incidence of the shock in that area.

The following picture emerges: places with higher house prices, an older age structure, higher educational attainment of the resident population that live in households with more than two members (typically, families), and where there is higher shares of people living in fuel poverty (pre crisis) are more exposed to the energy price shock.

These results suggest that the energy price shock is more likely to hit areas stronger that are typically considered to be either quite well off or places that already stand out with having a resident population that struggles economically with fuel poverty. We also note that the shock will hit less hard areas that stand out by having relatively high shares of households living in either social or other private rented accommodation. This finding might be due to the fact that social or privately rented homes are typically flats or apartments that are more energy efficient, on average.

The observation that the energy price shock is expected to hit more strongly areas that are either well off and areas in which a high share struggles already socio-economically from fuel poverty highlights the importance that interventions in energy markets are targeted in the appropriate way to ensure they reach the right set of households. We next show and contrast how untargeted and regressive the energy price guarantee is in this regard.

4.2 Illustration of untargeted and regressive nature of the EPG

We next show and characterise statistically how untargeted the EPG is and how it distorts incentives to save energy. A more succinct characterisation of the regressive nature of the EPG can be found in Fetzer (2022). The EPG puts a wedge between the consumer-facing prices and the prices set by the Ofgem regulator that aim to proxy market prices.

We denote the increase in bills that would have arisen if energy prices had been set as per the Ofgem price cap announced in October 2022 relative to the October 2021 price cap. We can consider this to be the passed-through regulated price that consumers would face without government intervention.

$$\Delta C_{p,act,22-21}^{ensemble} = C_{p,act,22-10}^{ensemble} - C_{p,act,21-10}^{ensemble} \quad (2)$$

With the Energy Price Guarantee (EPG) this increase can be decomposed into two components. For each property p , the first component represents the increase from October 2021 bills to the energy bills that *consumers* face under the EPG:

$$\Delta C_{p,act,EPG-21}^{ensemble} = C_{p,act,EPG}^{ensemble} - C_{p,act,21-10}^{ensemble} \quad (3)$$

The second component represents the implicit subsidy that the government pays, that is the wedge between the Ofgem market price and the EPG price:

$$\Delta C_{p,act,EPG-22}^{ensemble} = C_{p,act,22-10}^{ensemble} - C_{p,act,EPG}^{ensemble} \quad (4)$$

We can construct the same decomposition also for the two-tier tariff. We create the MSOA-level equivalent, computing the MSOA-level average of these metrics and then proceed to statistically compare what MSOA-level socio-economic characteristics drives the underlying variation. We next characterise to what extent, under the different policies, the energy price shock appears more or less progressive.

Empirical specification. We estimate the following linear specification

$$\Delta C_{m,act,j}^{ensemble} = \alpha_{d(m)} + x_m \times \beta + \epsilon_d$$

where x_m measures socio-economic characteristics of the area or the resident population. We are particularly interested in measures capturing the relative deprivation of an MSOA as this can provide some insights into the incidence of the energy price shock and, in turn, the relative regressiveness of progressiveness of the intervention in the energy price mechanism. We also absorb district fixed effects, $\alpha_{d(m)}$.

Results. The results are presented in Table 4. We study three measures of relative deprivation: the median property price in an area, average household income, and a rank measure of income of the resident population from the Indices of Multiple Deprivation.

Column 1 documents that the energy price shock is very progressive: the exposure of the average property in our data is notably higher in areas with higher property prices, higher average income, and higher income rank. This is not surprising as the energy price shock, in essence, hits households that consume a lot of energy – which tend to be better off – more. As a result, the energy price shock, in essence, has a similar effect of a carbon taxation.¹⁴

¹⁴The empirical evidence on whether a carbon tax is progressive or may be regressive is not clear with very little empirical work apart from modelling work on the subject to date. A notable exception is Andersson and Atkinson (2020).

Columns 2 and 3 look at what happens under the EPG. The measures 3 and 4 separate the energy price shock into a component absorbed by consumers vis-a-vis a component that is absorbed by the government providing an implicit energy consumption subsidy. The results suggest that, not surprisingly, the consumer-facing increase in the energy bill becomes much less progressive: the coefficient in Column 2 is about half the size of the coefficient in Column 1. On the other hand, the subsidy is extremely regressive: the government supports households that live in areas that are economically (much) better off. The EPG thereby can increase inequality across regions further.

Columns 4 and 5 examine the two-tier tariff for comparison on the same two metrics. The income gradient of the consumer-facing energy price shock is similarly progressive as under market prices. Yet, the subsidy is actually progressive, suggesting that households in regions that are economically well-off do not benefit from the subsidy to the same extent. This pattern simply arises from the fact that the two-tier tariff is much more targeted.

Visual illustration. We can illustrate this pattern quite nicely using a binned scatter plot along with a linear fit to gain more insight in the underlying relationships. In Figure 6, we visualize columns 1, 2, and 4 from Table 4.

Panel A presents a binned scatterplot with the linear regression fit considering house prices as measure of an area's relative wealth. We note that, without any government support, average bills would have increased drastically between 2021 and 2022 (navy diamonds). The energy price guarantee shifts this line downwards as it provides a relief for all consumers based on their consumption. But, owing to the fact that wealthier households consume more energy and due to the fact that, on average, they tend to live in particularly energy-inefficient properties, as we will show, this implies that the subsidy benefits better-off areas much more than regions that are worse-off. More importantly, the fact that the price gets distorted downwards implies that the energy savings incentives get weakened.

In contrast, a two-tier tariff keeps the house-price and energy bill gradient nearly the same as under market prices. In other words, a two-tier tariff resembles a lump sum transfer to households that have relatively low energy consumption. The marginal price signal remains intact and in fact, is slightly steepened. As a result,

energy savings incentives remain intact and this intervention appears much less regressive.

We can illustrate this finding with another measure of an area’s relative economic standing based on local area income statistics. We note a similar pattern. The importance of preserving the price signal becomes even more relevant when we document that, in fact, the highest energy savings potential is locked up in building stock in some of the most affluent parts of the UK.

4.3 Where is the highest energy savings potential in England?

We next study where the highest energy savings potential is in England and Wales. To do so, we construct the difference in expected energy bills, holding energy prices constant, between the actual- and potential energy consumption estimate derived from the EPC data. That is, we compute:

$$\Delta C_{p,act-pot,22}^{ensemble} = C_{p,act,22-10}^{ensemble} - C_{p,pot,22-10}^{ensemble} \quad (5)$$

Empirical specification As in the previous descriptive piece of work, we perform a best-subset selection (BSS) exercise to characterise the statistical patterns that drive the variation in energy savings potential. The focus here is on the components that we cannot measure in the property-level data: the socio-economic makeup of the resident population. While the previous exercise shed light on the likely incidence of the energy cost increase relative to 2021 *without government intervention*, we noted that the energy price shock - while striking in the heart of society - it appears to disproportionately hit two groups: households in rather well-off areas and, in particular, households that already find themselves living in fuel poverty.

We estimate the following linear specification at the MSOA level:

$$\Delta C_{m,act-pot,22}^{ensemble} = \alpha_{d(m)} + x_m \times \beta + \epsilon_d$$

where x_m measures socio-economic characteristics of the area or the resident population. We also absorb district fixed effects, $\alpha_{d(m)}$. We next want to characterise where the energy savings potential is distributed spatially. This characterisation may cast some light on potential policy interventions to encourage energy savings

that can alleviate the likely hardship induced by the energy price shock.

Results. The results are presented in Table 5. The table is to be read in the same way as the previous exercise except that the interpretation of the coefficients is now a different one as the dependent variable now measures the energy savings potential.

The results suggest that the energy savings potential increases with house prices in an area. This is not surprising: prior to 2021, energy was very cheap and hence, it hardly made any difference for households final bills with the energy bills of high consuming households being higher than that of low energy consuming households, but not by much. As a result households had very little financial incentive to save energy. With the present energy price shock, the personal incentive to save energy is now aligned which creates a unique opportunity to devise targeted interventions to either encourage households that are well-off to improve the energy efficiency of their properties, or to help households in fuel poverty that cannot finance their energy savings measures to boost these.

Similar to the pattern we saw for the energy price shock incidence, the energy savings potential appears more concentrated in areas with higher house prices, with an older and more educated demographic as one cluster. A second cluster highlights areas that have high existing degrees of fuel poverty, while a third cluster characterises MSOAs with little savings which is typically more densely populated MSOAs with a high share of people living in rented accommodation

These patterns highlight that it is imperative that well-off households should be given all incentives to upgrade their properties and make them more energy efficient and to ensure that the market price signals point in that direction. At the same time, households that are fuel poor should receive assistance or grants to enable them to stem their upgrade investments. We next document that this is not the case with the current policy framework that is in place.

We also note a strong, negative correlation between the fraction of households renting their properties from their council and exposure to the energy crisis. This is consistent with the socially rented sector being England's most energy efficient tenure (ONS, 2021). Social housing was the focus of the Decent Homes Programme, which sought to bring properties to minimum efficiency standards by 2010 (Leices-

ter and Stoye, 2017).

4.4 How does present policy affect energy savings investment incentives?

We documented three facts. First, the energy price shock will hit the most affluent parts of England more than the least affluent parts. Second, as a result these households stand to benefit most from the energy price guarantee, making this particular policy instrument so regressive. And, as we highlight this is a policy choice: policy alternatives such as a variable price cap and other designs that preserve the incentives do exist (see Bhattacharjee et al., 2022; Bachmann et al., 2022). Third, these patterns are due to decades of inaction: the more affluent parts of England and Wales boast a building stock that is quite energy inefficient and is home to a lot of people who consume, on average, more energy because they are economically better off.

We next characterise the disincentive effect of the EPG, that is how it disincentivizes energy efficiency investments, and the regressiveness of this disincentive effect. To do so, we construct the measure of the value of the energy savings measures under the October 2022 price cap

$$\Delta C_{p,act-pot,22}^{ensemble} = C_{p,act,22-10}^{ensemble} - C_{p,pot,22-10}^{ensemble} \quad (6)$$

as well as under the EPG

$$\Delta C_{p,act-pot,EPG}^{ensemble} = C_{p,act,EPG}^{ensemble} - C_{p,pot,EPG}^{ensemble} \quad (7)$$

and the alternative two-tier tariff

$$\Delta C_{p,act-pot,Two\ tier}^{ensemble} = C_{p,act,Two\ tier}^{ensemble} - C_{p,pot,Two\ tier}^{ensemble} \quad (8)$$

We construct the MSOA-specific averages of these three measures and then estimate a simple linear regression that provides the gradient of these measures with respect to proxy measures capturing the relative economic affluence of said areas:

$$\Delta C_{m,act-pot,j}^{ensemble} = \alpha_{d(m)} + x_m \times \beta + \epsilon_d$$

Results. The results are presented in Table 6. Column 1 documents that the value of energy savings for the average property in our dat increases with an area’s house prices, the average income of the households, and the overall income rank based on the IMD data.

In column 2, we perform the same exercise but now we value the energy savings at the prices implied by the EPG. We note that the gradient becomes notably weaker: energy savings measures, under the EPG, provide a weaker monetary incentive for households in well-off areas as compared to a scenario with prices that are closer to market prices in Column 1. Column 3 documents how the energy savings incentive is maintained and even strengthened by a two-tier tariff compared to what market prices would imply in terms of the income or house price gradient in Column 1.

5 Conclusion

This paper develops a measurement framework that enables us to model the likely impact of the energy price shock in a forward-looking way. We use this framework to characterise this impact across different areas of England and Wales. This method gives us a window into measuring the energy savings potential and quantifies the cost and benefit analysis of energy savings investments. Fetzer (2022) presents more detailed estimates.

We observe that the present UK policy path is quite incoherent. The EPG disproportionately benefits well-off households as the reduction in the unit rate relative to market prices disproportionately benefits households with higher levels of energy consumption. Because energy consumption strongly increases with household income, this disproportionately benefits households on higher incomes. We estimate that the financial benefit of the EPG is skewed even among high-earners: Among an estimated 280,000 households with an annual income above £150,000, a small number of around 14,000 households, representing the 95th percentile of energy consumption in this group, consumes more than twice as much energy compared to the nearly 50% of all other households in this high-income group. These super-

consumers will benefit most from the EPG.

Our analysis highlights that the UK has a large and untapped energy savings potential. We estimate that England and Wales alone could save up to 29% of primary energy consumption in the residential sector through reduced electricity and natural gas consumption used for space heating and hot water generation if residential properties were upgraded to their highest energy efficiency standard – primarily through improved insulation measures. The EPG weakens incentives to invest in energy efficiency upgrades by around 30%: Energy efficiency upgrade investments would save households between £10 to 16 billion per year if they had to pay current market prices for energy. The EPG lowers the prices that consumers face, as a result, energy efficiency upgrade investments under the EPG would only save households between £7 to 11 billion rendering them less economical.

Energy efficiency upgrades would provide permanent financial relief to households and offer large environmental benefits. Energy efficiency investments could permanently lower CO₂ emissions which itself has a large present and future monetary value as carbon taxes are set to increase energy prices going forward. We estimate that the energy efficiency upgrades may save between 25 to 40 million tons of CO₂ per year which, with a carbon price of £75 per ton would provide further savings between £1.5 to £3 billion per year.

Households with the highest energy consumption levels are least encouraged to make energy saving measures. A household in the top 5% of the energy consumption distribution is estimated to save more than 7,500 kWh per year if the property is upgraded to its highest energy performance standards. But the savings these consumers stand to make through the EPG significantly weakens their financial incentive to make such investments in boosting the energy efficiency.

Energy savings investments could pay for themselves within a relatively short period of time. With an interest rate of 3%, projected savings of £10 billion, and an investment volume of £60 billion for the properties for which we have EPC-based recommendations (around 50% of the building stock), we estimate that energy efficiency upgrades, in particular insulation and boiler replacement, would pay for themselves within six to seven years.

What is vital to highlight is that, the UK government, unlike governments in other countries actually has, at its disposal, much of the data needed to ensure

timely, targeted, and cost-effective interventions.

With these data, there are broad sets of alternative policies that could be considered. For example, instead of a price guarantee, the government could propose a two-tier tariff providing more generous targeted support without eroding energy savings incentives. Alternatives that provide even more targeted support with better incentive preservation do also exist and may also be implementable (see Bhattacharjee et al., 2022; Bachmann et al., 2022). The two tier-tariff could be designed to have a similar costing as the governments EPG but could be even more targeted.

We estimate that boiler replacements for the properties we have EPC data for would cost around £10 billion. The insulation program would cost around £50billion. A targeted means-tested insulation and boiler replacement grants program could be devised to help low-income households that do not have the means to afford the upgrades. It is imperative to build more evidence to understand why existing schemes have not been successful in achieving scale. For higher income households, a homeowner energy savings upgrade incentive program operating via tax-credits could be devised to encourage homeowners to take up energy efficiency measures. A national energy savings lottery may encourage behavioral changes around energy use.

Finally, further research is needed to study the role that planning laws play in facilitating or hindering energy saving building upgrades as well as the role of local initiatives to promote economies of scale and wider take-up.

References

- Ahlrichs, Jakob, Simon Wenninger, Christian Wiethe, and Björn Häckel**, “Impact of socio-economic factors on local energetic retrofitting needs - A data analytics approach,” *Energy Policy*, 2022, 160 (February 2021).
- Allcott, Hunt and Michael Greenstone**, “Is there an energy efficiency gap?,” *Journal of Economic perspectives*, 2012, 26 (1), 3–28.
- Andersson, Julius and Giles Atkinson**, “The Distributional Effects of a Carbon Tax: the Role of Income Inequality,” *Centre for Climate Change Economics and Policy Working Paper 378/Grantham Research Institute on Climate Change and the Environment Working Paper 349*, 2020, 5709 (378), 1–35.
- Attari, Shahzeen Z., Michael L. DeKay, Cliff I. Davidson, and Wändi Bruin De Bruin**, “Public perceptions of energy consumption and savings,” *Proceedings of the National Academy of Sciences of the United States of America*, 2010, 107 (37), 16054–16059.
- Bachmann, Rüdiger, Moritz Kuhn, Rüdiger Bachmann, David Baqaee, Christian Bayer, Moritz Kuhn, Andreas Löschel, Ben McWilliams, Benjamin Moll, Andreas Peichl, Karen Pittel, Moritz Schularick, and Georg Zachmann**, “How it can be done,” *EconTribute Working Paper*, 2022.
- Bhattacharjee, Arnab, Max Mosley, and Adrian Pabst**, “A ‘ Variable Energy Price Cap ’ to Help Solve the Cost-of-Living Crisis,” *NIESR Working Paper*, 2022, (September).
- BRE**, “The Government’s Standard Assessment Procedure for Energy Rating of Dwellings,” Technical Report 2014.
- Christensen, Peter, Paul Francisco, Erica Myers, and Mateus Souza**, “Decomposing the wedge between projected and realized returns in energy efficiency programs,” *The Review of Economics and Statistics*, 2021, pp. 1–46.
- Dalton, Ben and Franz Fuerst**, “The ‘green value’ proposition in real estate: A meta-analysis,” *Routledge Handbook of Sustainable Real Estate*, 2018, pp. 177–200.
- Department for Business, Energy & Industrial Strategy**, “Evaluation of the Domestic Private Rented Sector Minimum Energy Efficiency Standard Regulations,” Technical Report 2020.
- , “Annual Fuel Poverty Statistics in England, 2022 (2020 data),” Technical Report

2022.

Department for Levelling Up, Housing & Communities, “A guide to energy performance certificates for the marketing, sale and let of dwellings,” Technical Report 2017.

– , “Energy Performance of Buildings Certificates: notes and definitions,” Technical Report 2021.

Fetzer, Thiemo, “The Energy Price Guarantee: what next?,” *CAGE Policy Briefing*, 2022, pp. 1–24.

Gerarden, Todd D, Richard G Newell, and Robert N Stavins, “Assessing the energy-efficiency gap,” *Journal of Economic Literature*, 2017, 55 (4), 1486–1525.

Gregório, Vera and Júlia Seixas, “Energy savings potential in urban rehabilitation: A spatial-based methodology applied to historic centres,” *Energy and Buildings*, 2017, 152, 11–23.

Guin, Benjamin, Perttu Korhonen, and Sidharth Moktan, “Risk differentials between green and brown assets?,” *Economics Letters*, 2022, 213, 110320.

Harari, Daniel, Brigid Francis-Devine, Paul Bolton, and Matthew Keep, “Research Briefing - Rising cost of living in the UK,” *House of Commons Library*, 2022, pp. 4–5.

Houde, Sébastien and Erica Myers, “Are consumers attentive to local energy costs? Evidence from the appliance market,” *Journal of Public Economics*, 2021, 201, 104480.

IEA, “Electricity Market Report - July 2022,” Technical Report 2022.

Institute for Government, “Decarbonising heating at home - Learning from past successes and failures to improve energy policy making ,” Technical Report 2021.

Labandeira, Xavier, José M Labeaga, and Xiral López-Otero, “A meta-analysis on the price elasticity of energy demand,” *Energy Policy*, mar 2017, 102, 549–568.

Leicester, Andrew and George Stoye, “Factors Associated with the Presence of Domestic Energy Efficiency Measures in England,” *Fiscal Studies*, 2017, 38 (2), 331–356.

Myers, Erica, “Are home buyers inattentive? Evidence from capitalization of energy costs,” *American Economic Journal: Economic Policy*, 2019, 11 (2), 165–88.

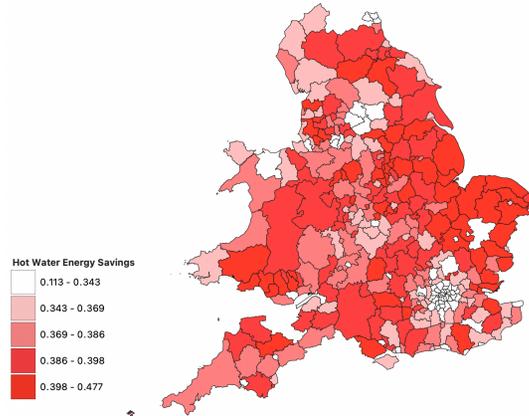
– , “Asymmetric information in residential rental markets: Implications for the

- energy efficiency gap,” *Journal of Public Economics*, 2020, 190, 104251.
- , **Steven Puller, Jeremy West et al.**, “Mandatory Energy Efficiency Disclosure in Housing Markets,” *American Economic Journal: Economic Policy*, 2022, 14 (4), 453–87.
- Ofgem**, “Direct debits - What you need to know,” Technical Report 2019.
- , “Ofgem updates price cap level and tightens up rules on suppliers,” Technical Report 2022.
- , “Price cap – Decision on changes to the wholesale methodology,” Technical Report 2022.
- ONS**, “Home People, population and community Housing Energy efficiency of housing in England and Wales Energy efficiency of housing in England and Wales: 2021,” Technical Report 2021.
- Zhang, Tao, Peer Olaf Siebers, and Uwe Aickelin**, “A three-dimensional model of residential energy consumer archetypes for local energy policy design in the UK,” *Energy Policy*, 2012, 47, 102–110.

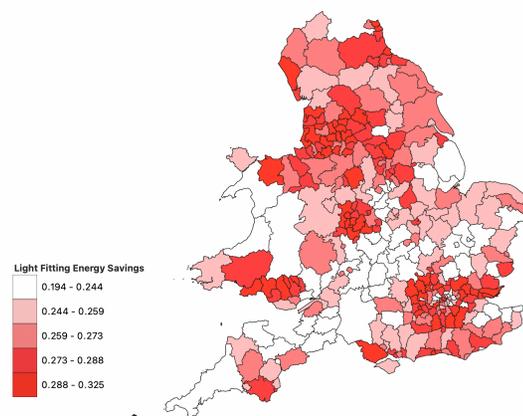
Figures and tables

Figure 1: Energy saving potential measured in real quantities

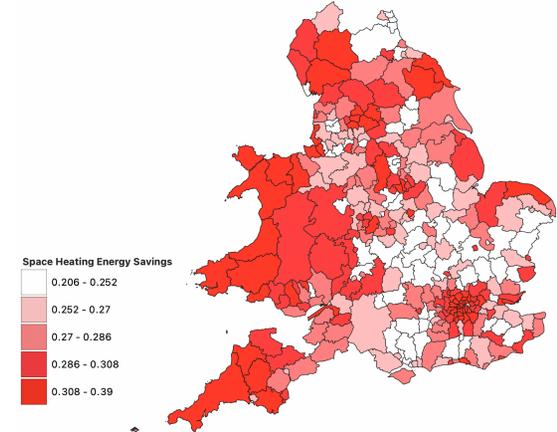
Panel A: Hot Water



Panel B: Light Fittings

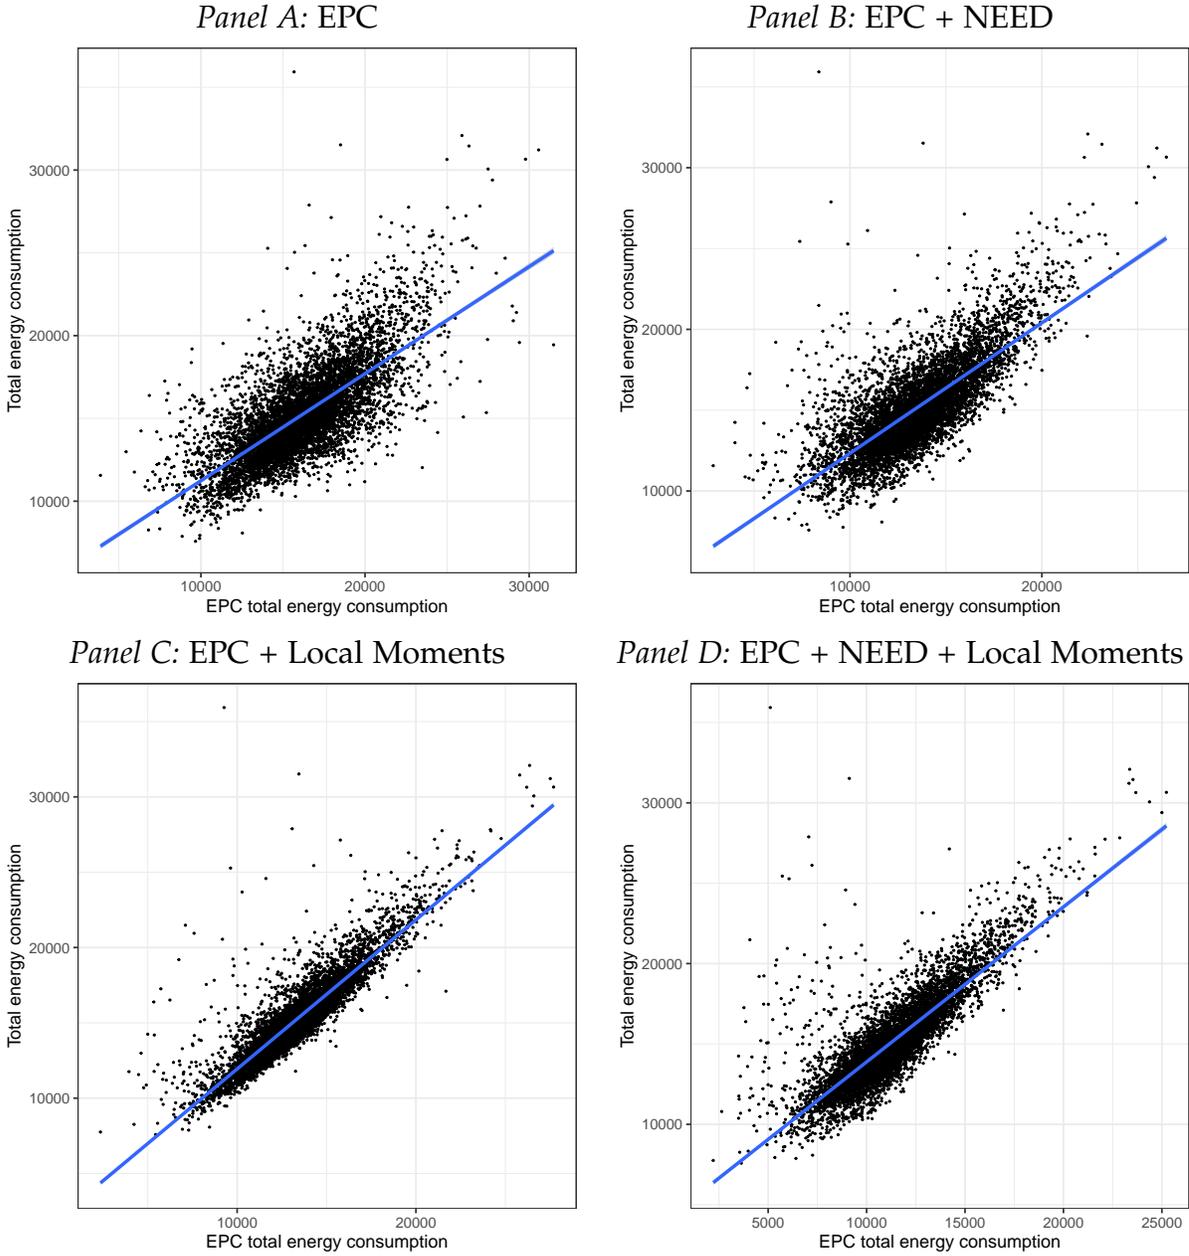


Panel C: Space Heating



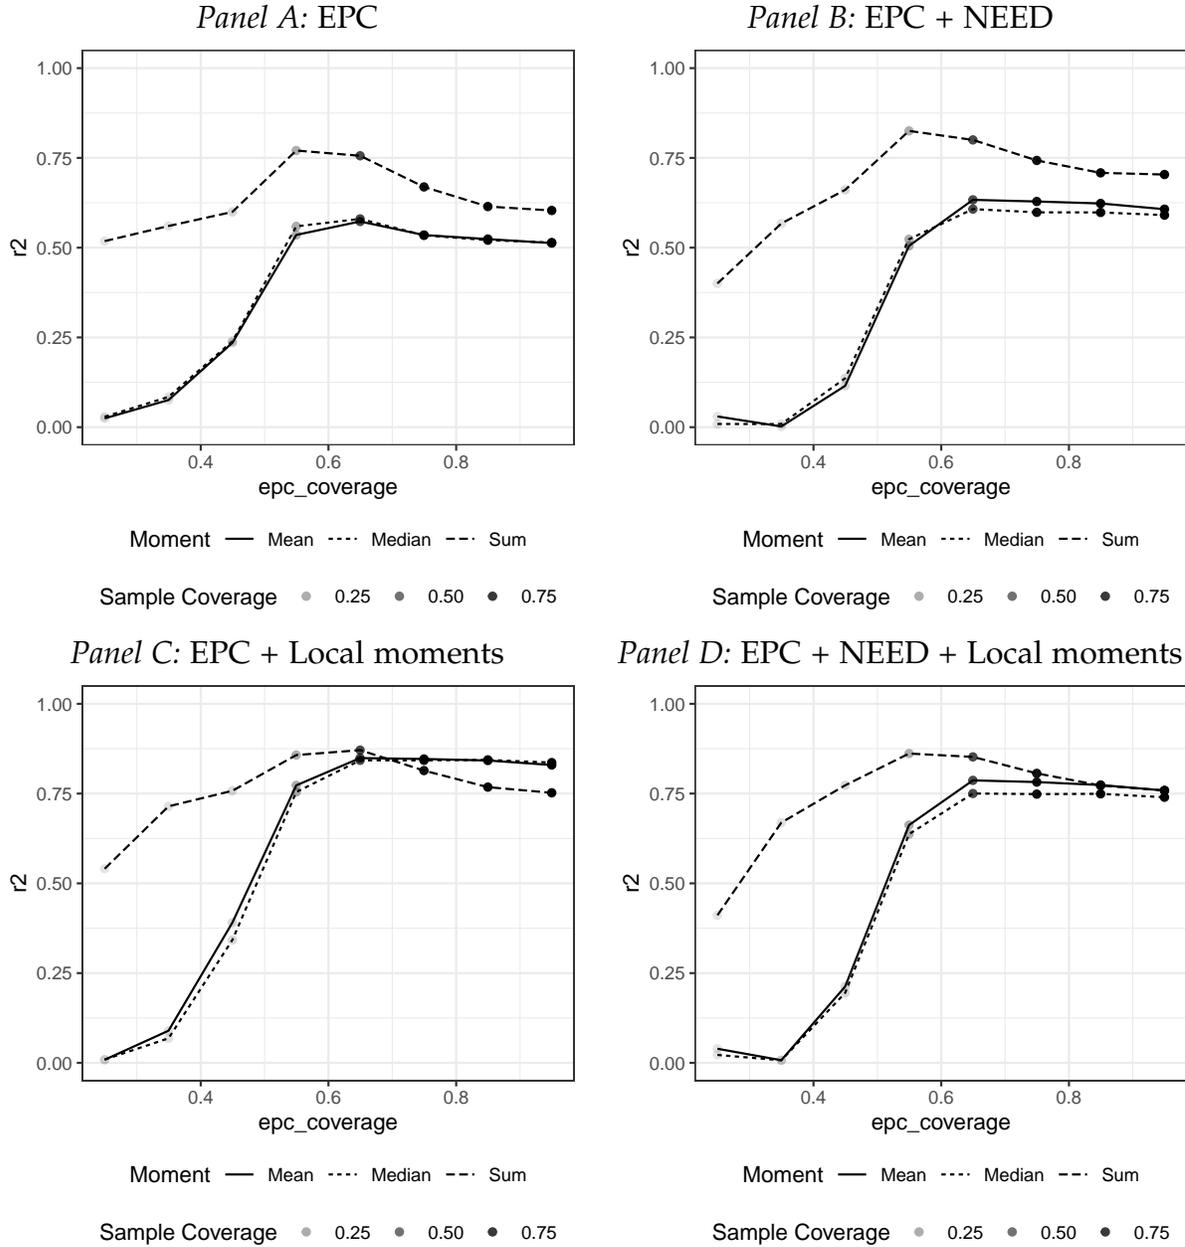
Notes: Figures present aggregate energy savings potential, measured in % of real units (kwh), that is the ratio of potential energy consumption over actual energy consumption. The higher the % the higher the gap between actual and potential consumption in relative terms.

Figure 2: Median property-level energy consumption measures at the MSOA-level compared with median imputed energy consumption measures from EPC-NEED data



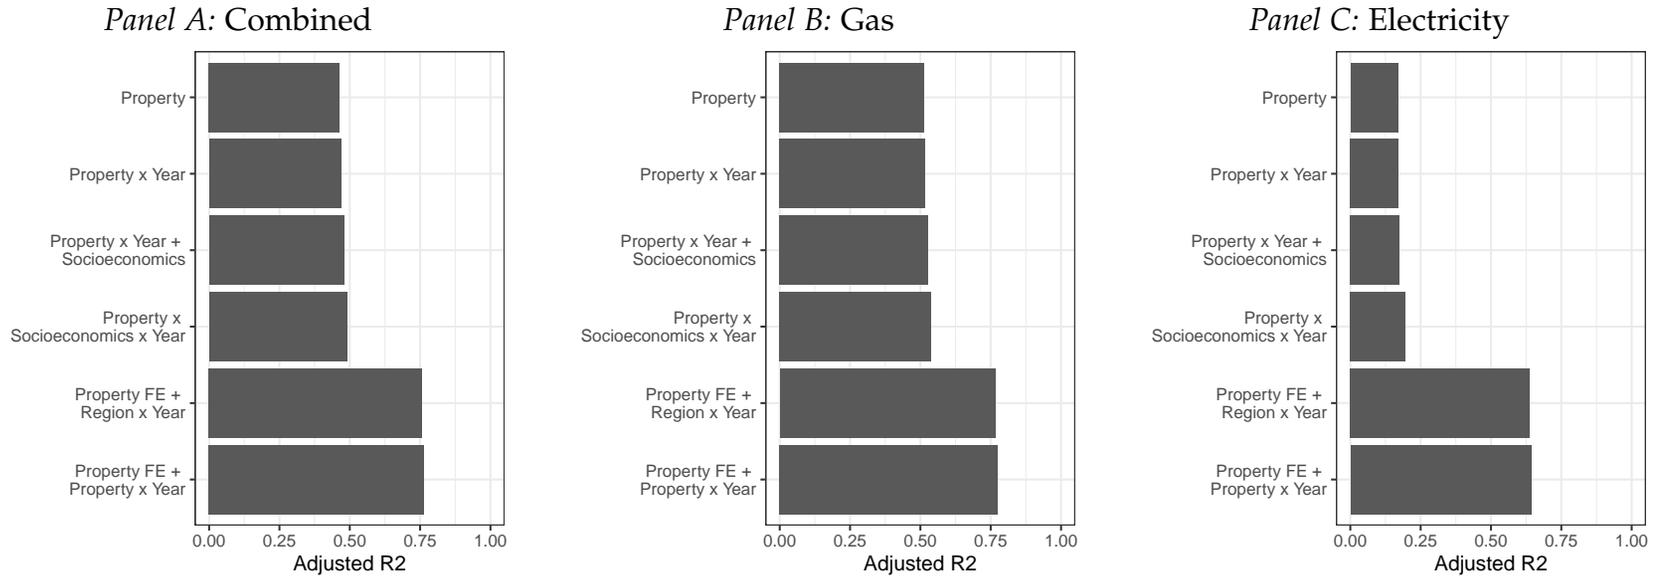
Notes: Figures provide a scatterplot of estimates of the median energy consumption per meter from published data at the MSOA-level (for metered electricity and gas only) on the vertical axis and the median of various imputed energy consumption measures that leverage different data on the horizontal axis. Panel A provides the implied consumption estimates from the EPC data as is. Panel B augments the EPC data with a matching-of-moments approach based on anonymized individual level meter reading data collected under the NEED framework. Panel C uses the EPC raw energy consumption estimates and augments it with matched granular area-specific moments. Panel D is the final measure that combines the EPC raw data, the property-specific moment-matching and the local area specific moment matching.

Figure 3: Correlation between moments of derived consumption proxy measures and moments from actual consumption data



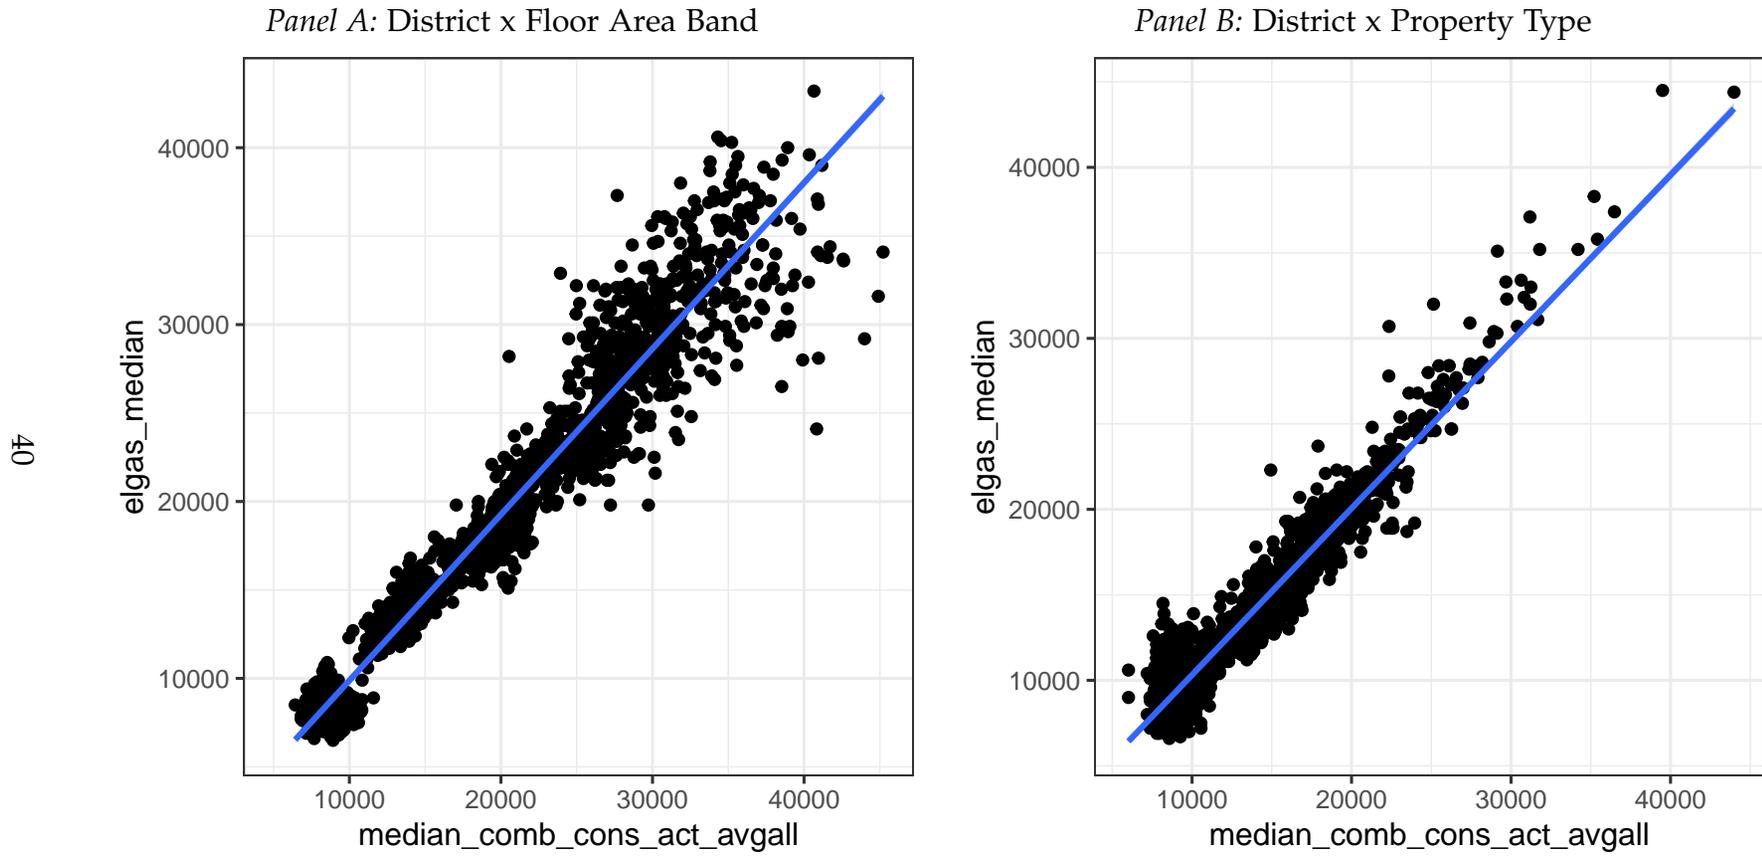
Notes: Figures plot the R^2 that is obtained from validating the derived implied consumption measures and three moments: the total consumption, the mean, and median consumption against actual consumption data that is published from gas and electricity meters across the country. For the four different derived measures, we compare the goodness-of-fit of the three moments against the corresponding moment from subnational statistics. The horizontal axis captures the ratio of the number of EPC properties against the population of properties in an area based on council tax data. A value of 0.4 on the axis implies that the estimating sample includes data from all MSOAs that have at most 40% of their building stock captured in the EPC data. We note that the goodness-of-fit remains stable across each of the moments when the estimating sample includes MSOAs with an EPC coverage of up to 60%.

Figure 4: Decomposition of variance in the anonymized individual property-level energy consumption data documenting to what extent different features can characterise the variation in energy consumption



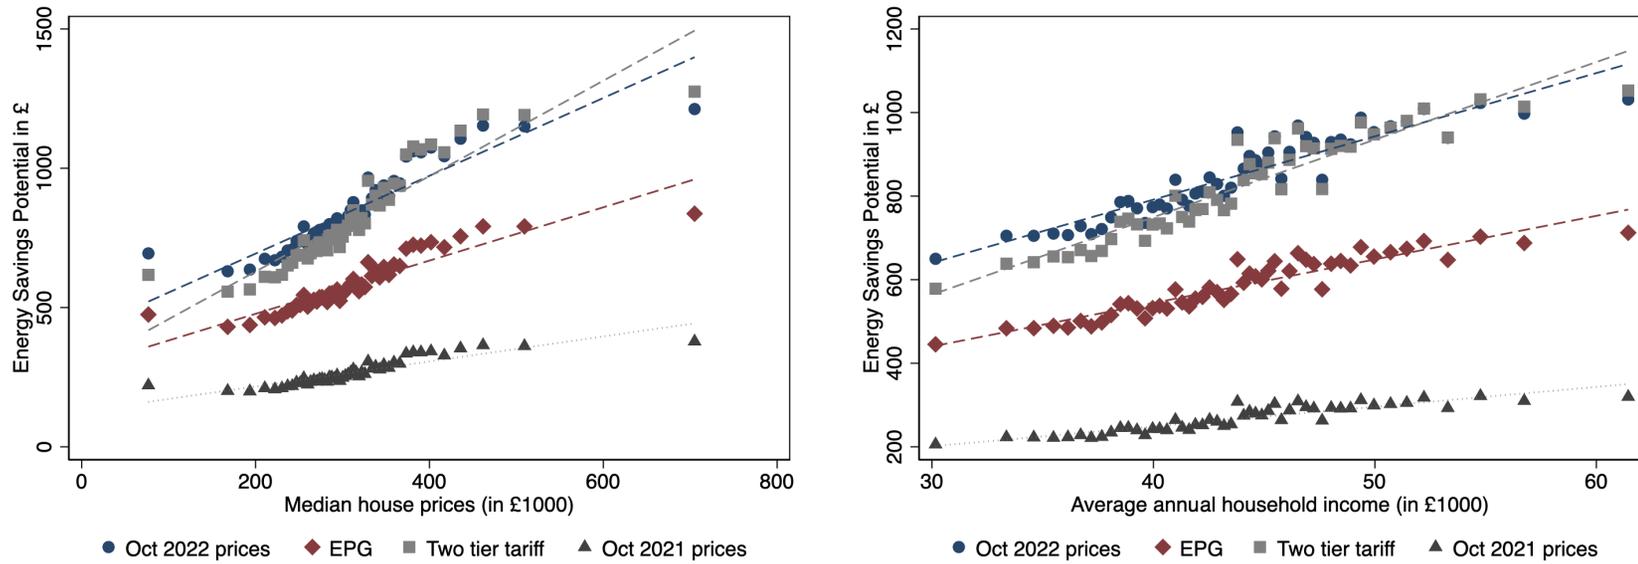
Notes: Figures plot out the adjusted R^2 obtained from regressing combined, gas, and electricity anonymized property-level consumption data against a set of features.

Figure 5: Unconditional raw scatter plot district-by-floor-area or district-by-property-type median energy consumption data vis-a-vis our EPC-derived ensemble consumption estimate



Notes: Figures plot a raw scatterplot of the median district-level energy consumption by floor area in Panel A or district-level energy consumption by property type in Panel B against the corresponding median constructed from our EPC-derived ensemble measure. The corresponding regression is presented in column 10 of Table 1 and Table 2 respectively.

Figure 6: Visualisation of the empirical link between house prices and the expected increase in average energy bills under Ofgem prices (market price proxy), the EPG, and the alternative two-tier tariff



Notes: Figures plot the relationship between median house prices at the MSOA-level against the expected increase in the energy bills to study the degree to which the specific measures are targeted in providing relief.

Table 1: Comparison of district-by-floor-area BEIS-reported average and median electricity and gas consumption vis-a-vis corresponding EPC-derived and rescaled proxy measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	<i>EPC</i>		<i>NEED</i>		<i>EPC + Local</i>		<i>EPC + NEED + Local</i>		<i>Average</i>	
<i>Panel A: No controls</i>										
Derived energy consumption proxy	0.859*** (0.011)	0.857*** (0.011)	0.924*** (0.015)	0.929*** (0.014)	0.977*** (0.007)	1.067*** (0.005)	1.035*** (0.010)	1.129*** (0.009)	0.919*** (0.010)	0.938*** (0.009)
R2	0.884	0.899	0.884	0.895	0.957	0.966	0.944	0.950	0.927	0.933
Observations	1650	1650	1650	1650	1650	1650	1650	1650	1650	1650
<i>Panel B: Floor Area Band FE</i>										
Derived energy consumption proxy	0.192*** (0.029)	0.298*** (0.031)	0.185*** (0.039)	0.298*** (0.041)	0.580*** (0.039)	0.806*** (0.037)	0.525*** (0.047)	0.714*** (0.044)	0.357*** (0.042)	0.484*** (0.042)
R2	0.953	0.947	0.952	0.946	0.971	0.973	0.965	0.964	0.958	0.955
Observations	1650	1650	1650	1650	1650	1650	1650	1650	1650	1650
<i>Panel C: District FE</i>										
Derived energy consumption proxy	0.901*** (0.009)	0.890*** (0.009)	0.960*** (0.012)	0.955*** (0.012)	0.986*** (0.007)	1.071*** (0.006)	1.046*** (0.009)	1.137*** (0.008)	0.939*** (0.008)	0.952*** (0.008)
R2	0.943	0.949	0.929	0.931	0.972	0.978	0.959	0.963	0.959	0.961
Observations	1650	1650	1650	1650	1650	1650	1650	1650	1650	1650
<i>Panel D: District FE and Floor Area Band FE</i>										
Derived energy consumption proxy	0.196*** (0.024)	0.286*** (0.026)	0.109*** (0.027)	0.185*** (0.029)	0.399*** (0.026)	0.606*** (0.028)	0.278*** (0.038)	0.448*** (0.042)	0.265*** (0.031)	0.367*** (0.034)
R2	0.986	0.982	0.984	0.979	0.989	0.988	0.986	0.983	0.986	0.983
Observations	1650	1650	1650	1650	1650	1650	1650	1650	1650	1650
Moment:	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median

Notes: Table presents regression results comparing the district-level average electricity and gas consumption average from BEIS micro data by floor-area type with the measures that we constructed as part of our proxy variables.

Table 2: Comparison of district-by-property-type BEIS-reported average and median electricity and gas consumption vis-a-vis corresponding EPC-derived and rescaled proxy measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	<i>EPC</i>		<i>NEED</i>		<i>EPC + Local</i>		<i>EPC + NEED + Local</i>		<i>Average</i>	
<i>Panel A: No controls</i>										
Derived energy consumption proxy	0.789*** (0.013)	0.859*** (0.011)	0.903*** (0.013)	0.975*** (0.010)	0.949*** (0.008)	1.071*** (0.011)	0.993*** (0.010)	1.121*** (0.012)	0.888*** (0.010)	0.975*** (0.009)
R2	0.756	0.850	0.823	0.896	0.887	0.942	0.881	0.927	0.845	0.918
Observations	2303	2303	2303	2303	2303	2303	2303	2303	2303	2303
<i>Panel B: Property Type FE</i>										
Derived energy consumption proxy	0.697*** (0.038)	0.782*** (0.037)	0.810*** (0.032)	0.892*** (0.029)	0.989*** (0.017)	1.111*** (0.028)	0.969*** (0.022)	1.102*** (0.032)	0.891*** (0.024)	0.968*** (0.024)
R2	0.822	0.888	0.846	0.906	0.899	0.947	0.882	0.928	0.871	0.929
Observations	2303	2303	2303	2303	2303	2303	2303	2303	2303	2303
<i>Panel C: District FE</i>										
Derived energy consumption proxy	0.807*** (0.009)	0.863*** (0.008)	0.930*** (0.011)	0.994*** (0.010)	0.937*** (0.007)	1.057*** (0.009)	1.004*** (0.010)	1.130*** (0.011)	0.888*** (0.008)	0.969*** (0.008)
R2	0.807	0.891	0.853	0.920	0.888	0.947	0.881	0.932	0.863	0.933
Observations	2303	2303	2303	2303	2303	2303	2303	2303	2303	2303
<i>Panel D: District FE and Property Type FE</i>										
Derived energy consumption proxy	0.707*** (0.033)	0.729*** (0.039)	0.855*** (0.043)	0.893*** (0.047)	0.951*** (0.024)	1.056*** (0.038)	1.020*** (0.040)	1.150*** (0.053)	0.893*** (0.027)	0.922*** (0.035)
R2	0.883	0.936	0.876	0.931	0.902	0.952	0.883	0.933	0.894	0.947
Observations	2303	2303	2303	2303	2303	2303	2303	2303	2303	2303
Moment:	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median

Notes: Table presents regression results comparing the district-level average electricity and gas consumption average from BEIS micro data by property type (detached, semi-detached, (mid/end) terraced, flat and/or bungalow) with the measures that we constructed as part of our proxy variables.

Table 3: Correlates of the exposure to the energy-price shock at the MSOA level - best-subset selection approach

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)		
median house price	2.686*** (0.337)	2.201*** (0.338)	1.685*** (0.337)	1.620*** (0.342)	1.446*** (0.346)	1.318*** (0.348)	1.187*** (0.335)	1.108*** (0.337)	1.079*** (0.334)	1.094*** (0.339)	1.069*** (0.340)	1.098*** (0.343)	1.091*** (0.342)	1.089*** (0.343)	1.078*** (0.348)	1.073*** (0.348)		
% population in social rented accommodation		-9.735*** (0.881)				-16.655*** (1.275)	-8.432*** (1.321)	-5.977*** (1.344)	-9.189*** (1.683)	-8.414*** (1.725)	-9.304*** (2.018)	-9.838*** (2.052)	-9.982*** (2.092)	-10.225*** (2.218)	-9.824*** (2.445)	-9.740*** (2.407)		
% population aged 65 plus			40.057*** (4.916)	57.770*** (6.409)	54.323*** (6.138)				30.258*** (3.906)	43.072*** (6.429)	35.037*** (6.883)	33.565*** (6.781)	30.095*** (8.526)	29.723*** (8.481)	31.090*** (7.950)	31.761*** (7.790)	31.624*** (7.741)	31.495*** (7.869)
% population disabled			-42.397*** (5.330)	-64.503*** (7.840)	-66.713*** (7.690)					-27.670*** (8.469)	-28.534*** (8.461)	-26.044*** (8.227)	-25.232*** (8.592)	-26.965*** (8.750)	-32.036*** (9.040)	-32.773*** (9.379)	-32.904*** (9.421)	-32.873*** (9.433)
% population in fuel poverty				28.528*** (5.586)	33.629*** (5.875)	54.710*** (6.898)	36.241*** (5.976)	38.844*** (6.482)	48.918*** (7.415)	49.522*** (7.347)	47.993*** (7.710)	47.439*** (7.654)	47.387*** (7.681)	46.945*** (8.262)	45.928*** (8.282)	45.810*** (8.356)		
population density					-3.968*** (0.815)	-4.567*** (0.801)	-4.778*** (0.756)	-4.609*** (0.741)	-4.295*** (0.735)	-4.023*** (0.695)	-4.004*** (0.699)	-4.077*** (0.703)	-4.065*** (0.705)	-4.079*** (0.700)	-4.031*** (0.709)	-4.041*** (0.702)		
% population with university degree						16.265*** (2.972)	16.503*** (2.849)	12.417*** (2.865)	13.467*** (2.911)	14.879*** (2.973)	15.236*** (3.017)	17.137*** (3.360)	17.217*** (3.380)	17.316*** (3.374)	17.632*** (3.526)	17.612*** (3.527)		
% population in private rented accommodation						-16.342*** (2.308)			-7.678*** (2.495)	-6.939*** (2.539)	-8.022*** (3.047)	-8.818*** (3.104)	-8.870*** (3.128)	-8.937*** (3.149)	-8.351** (3.301)	-8.311** (3.276)		
% households with more than 2 members							17.564*** (2.618)	15.411*** (2.686)	9.502*** (3.054)	9.435*** (3.121)	8.251*** (3.112)	8.762*** (3.115)	8.800*** (3.130)	8.740*** (3.335)	8.076** (3.335)	8.061** (3.345)		
% population commuting with public transport										-6.758* (3.465)	-6.719* (3.492)	-6.538* (3.474)	-6.676* (3.426)	-6.919** (3.428)	-6.867** (3.404)	-6.875** (3.404)		
% inactive											4.213 (3.779)	3.513 (3.869)	3.295 (3.728)	3.476 (3.876)	3.724 (3.857)	3.679 (3.829)		
average annual household income												-4.901 (3.679)	-4.979 (3.666)	-4.888 (3.654)	-5.139 (3.730)	-5.119 (3.740)		
% population in bad or very bad health													11.193 (17.024)	10.001 (16.086)	9.646 (16.120)	9.544 (16.054)		
% unemployed														6.269 (14.300)	6.986 (14.310)	6.813 (14.274)		
% population in whole house or bungalow															0.894 (1.547)	0.871 (1.553)		
% population in shared accommodation																-5.304 (14.218)		
Best Subset												X						
Observations	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789		
R2	.346	.359	.37	.381	.388	.394	.398	.4	.401	.402	.403	.403	.403	.403	.403	.403		

Notes: Table presents best-subset-selection regression results. The best-subset is indicated in the table footer. All regressions control for district fixed effects. The analysis presents correlational patterns that help characterize the incidence of the energy price shock in terms of the socio-economic characteristics of areas that are expected to be hit the most. Standard errors are clustered at the district level.

Table 4: Comparison of income-wealth gradient in the energy price shock incidence and the impact that the EPG or an alternative two-tier tariff has on the degree of progressivity

	(1)	(2)	(3)	(4)	(5)
Energy bill shock	<i>without intervention</i>	<i>with EPG</i>		<i>with two tier tariff</i>	
		Consumers	Govt	Consumers	Govt
<i>Panel A:</i>					
median house price	2.686*** (0.337)	1.492*** (0.186)	1.194*** (0.152)	2.852*** (0.321)	-0.166*** (0.058)
R2	0.315	0.324	0.301	0.394	0.188
Observations	6789	6789	6789	6789	6789
<i>Panel B:</i>					
average annual household income	34.175*** (2.176)	19.131*** (1.253)	15.044*** (0.933)	35.936*** (1.793)	-1.761** (0.749)
R2	0.276	0.284	0.266	0.326	0.186
Observations	6789	6789	6789	6789	6789
<i>Panel C:</i>					
Income Rank (IMD)	0.154*** (0.009)	0.085*** (0.005)	0.069*** (0.004)	0.156*** (0.007)	-0.003 (0.003)
R2	0.296	0.304	0.286	0.350	0.185
Observations	6789	6789	6789	6789	6789
District FE:	X	X	X	X	X

Notes: Table presents regression results showing how various energy-price shock measures on average bills vary with various socio-economic measures at the MSOA-level capturing income or wealth. Column 1 displays the overall shock to bills between 2021 and 2022. Columns 2 and 3 decompose this shock in the consumer- and the government-facing average bill increase under the EPG. Columns 4 and 5 present the same breakdown for the two-tier tariff intervention. All regressions include district fixed-effects. Standard errors are clustered at the district level with stars indicating *** p < 0.01, ** p < 0.05, and * p < 0.1.

Table 5: Correlates of the energy-savings potential at the MSOA level - best-subset selection approach

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
median house price	1.244*** (0.151)	0.900*** (0.146)	0.999*** (0.170)	0.856*** (0.160)	0.671*** (0.159)	0.612*** (0.154)	0.548*** (0.157)	0.505*** (0.159)	0.513*** (0.161)	0.543*** (0.164)	0.526*** (0.163)	0.511*** (0.162)	0.510*** (0.162)	0.499*** (0.163)	0.494*** (0.165)	0.495*** (0.166)
% population in social rented accommodation		-6.900*** (0.431)	-10.524*** (0.668)	-10.556*** (0.643)	-12.631*** (0.737)	-11.664*** (0.700)	-11.809*** (0.703)	-10.487*** (0.795)	-10.040*** (0.831)	-10.901*** (0.903)	-10.615*** (0.897)	-9.972*** (1.001)	-10.316*** (1.040)	-9.847*** (1.228)	-9.933*** (1.412)	-9.908*** (1.436)
% population in fuel poverty			19.510*** (2.891)	24.122*** (2.874)	35.739*** (3.959)	35.275*** (3.733)	36.745*** (3.854)	37.355*** (4.003)	37.669*** (3.972)	37.007*** (3.956)	36.381*** (4.072)	34.478*** (4.491)	33.909*** (4.785)	32.921*** (4.662)	32.667*** (5.132)	32.646*** (5.097)
% population aged 65 plus				13.496*** (1.695)		9.336*** (1.676)	8.262*** (1.599)	14.512*** (3.724)	13.707*** (3.596)	12.418*** (3.632)	12.101*** (3.718)	13.719*** (3.784)	15.099*** (3.564)	15.180*** (3.607)	14.783*** (4.540)	14.529*** (4.279)
% population with university degree					6.936*** (1.389)	6.667*** (1.313)	6.786*** (1.349)	5.208*** (1.306)	6.011*** (1.338)	8.141*** (1.557)	8.147*** (1.556)	8.866*** (1.580)	9.032*** (1.586)	9.377*** (1.652)	9.385*** (1.652)	9.376*** (1.661)
% population in private rented accommodation					-9.324*** (1.127)	-6.909*** (1.071)	-6.338*** (1.169)	-5.944*** (1.104)	-5.509*** (1.169)	-6.785*** (1.288)	-6.578*** (1.269)	-5.315*** (1.484)	-5.379*** (1.490)	-4.705*** (1.688)	-4.785*** (1.869)	-4.771*** (1.883)
population density							-1.364*** (0.368)	-1.306*** (0.356)	-1.154*** (0.329)	-1.224*** (0.330)	-1.256*** (0.327)	-1.320*** (0.318)	-1.340*** (0.314)	-1.290*** (0.311)	-1.282*** (0.318)	-1.285*** (0.320)
% population disabled								-11.828** (4.817)	-10.407** (4.611)	-12.390*** (4.563)	-12.464*** (4.574)	-11.186** (4.680)	-13.115*** (4.359)	-13.464*** (4.509)	-13.376*** (4.682)	-12.370** (4.920)
% population commuting with public transport									-3.805** (1.736)	-3.597** (1.739)	-3.640** (1.738)	-3.598** (1.733)	-3.973** (1.674)	-3.914** (1.649)	-3.923** (1.639)	-3.905** (1.620)
average annual household income										-5.514*** (1.882)	-5.455*** (1.878)	-5.846*** (1.831)	-5.793*** (1.811)	-6.141*** (1.871)	-5.991*** (1.950)	-5.974*** (1.943)
% population in shared accommodation											-16.695** (7.349)	-15.923** (7.240)	-15.520** (7.121)	-15.221** (7.197)	-15.071** (6.905)	-15.114** (6.836)
% households with more than 2 members												2.701 (1.671)	2.692 (1.665)	2.062 (1.842)	1.853 (1.828)	1.834 (1.843)
% unemployed													9.159 (7.977)	9.766 (7.909)	10.085 (8.621)	10.365 (8.252)
% population in whole house or bungalow														0.956 (0.921)	1.011 (0.855)	1.021 (0.863)
% inactive															0.583 (2.400)	0.638 (2.307)
% population in bad or very bad health																-2.351 (9.121)
Best Subset																X
Observations	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789	6789
R2	.314	.34	.36	.374	.379	.384	.387	.389	.39	.391	.391	.392	.392	.392	.392	.392

Notes: Table presents best-subset-selection regression results. The best-subset is indicated in the table footer. All regressions control for district fixed effects. The analysis presents correlational patterns that help characterize the variation in energy savings potential in terms of the socio-economic characteristics of areas. Standard errors are clustered at the district level.

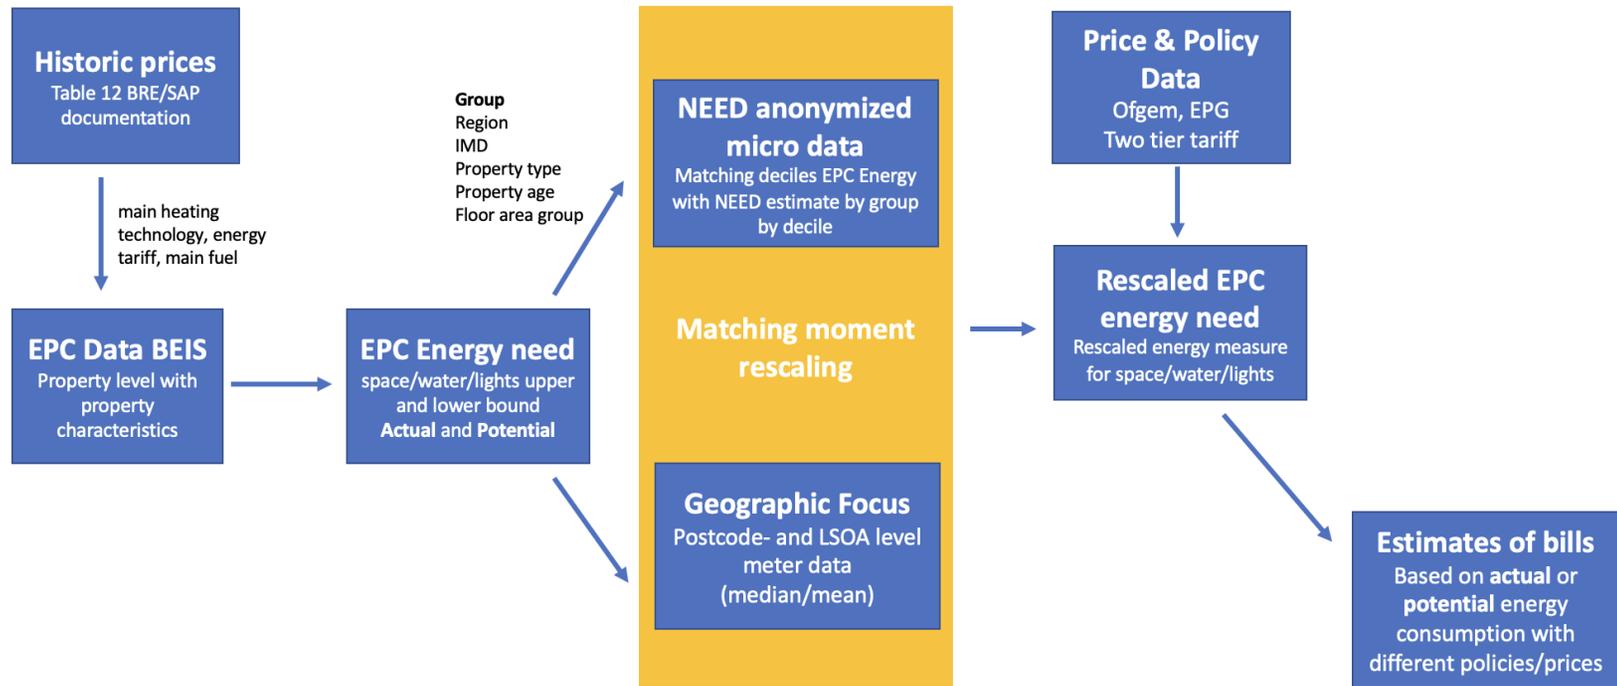
Table 6: Comparison of income-wealth gradient in the value of the energy-savings potential contrasting the full price shock, the EPG-mitigated, and the two-tier tariff mitigated shock

	(1)	(2)	(3)
Energy savings incentive	<i>without intervention</i>	<i>with EPG</i>	<i>with two tier tariff</i>
<i>Panel A:</i>			
median house price	1.244*** (0.151)	0.853*** (0.103)	1.526*** (0.178)
R2	0.281	0.288	0.302
Observations	6789	6789	6789
<i>Panel B:</i>			
average annual household income	15.029*** (0.968)	10.349*** (0.670)	18.436*** (1.110)
R2	0.245	0.250	0.258
Observations	6789	6789	6789
<i>Panel C:</i>			
income rank	0.070*** (0.004)	0.048*** (0.003)	0.085*** (0.005)
R2	0.264	0.270	0.279
Observations	6789	6789	6789
District FE:	X	X	X

Notes: Table presents regression results showing how the various energy-price shock measures affect the gradient of average energy savings potential with respect to various socio-economic measures at the MSOA-level capturing income or wealth. Column 1 displays the overall shock to bills between 2021 and 2022. Column 2 presents the energy savings potential under the EPG, while column 3 presents it under the two-tier tariff. All regressions include district fixed-effects. Standard errors are clustered at the district level with stars indicating *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

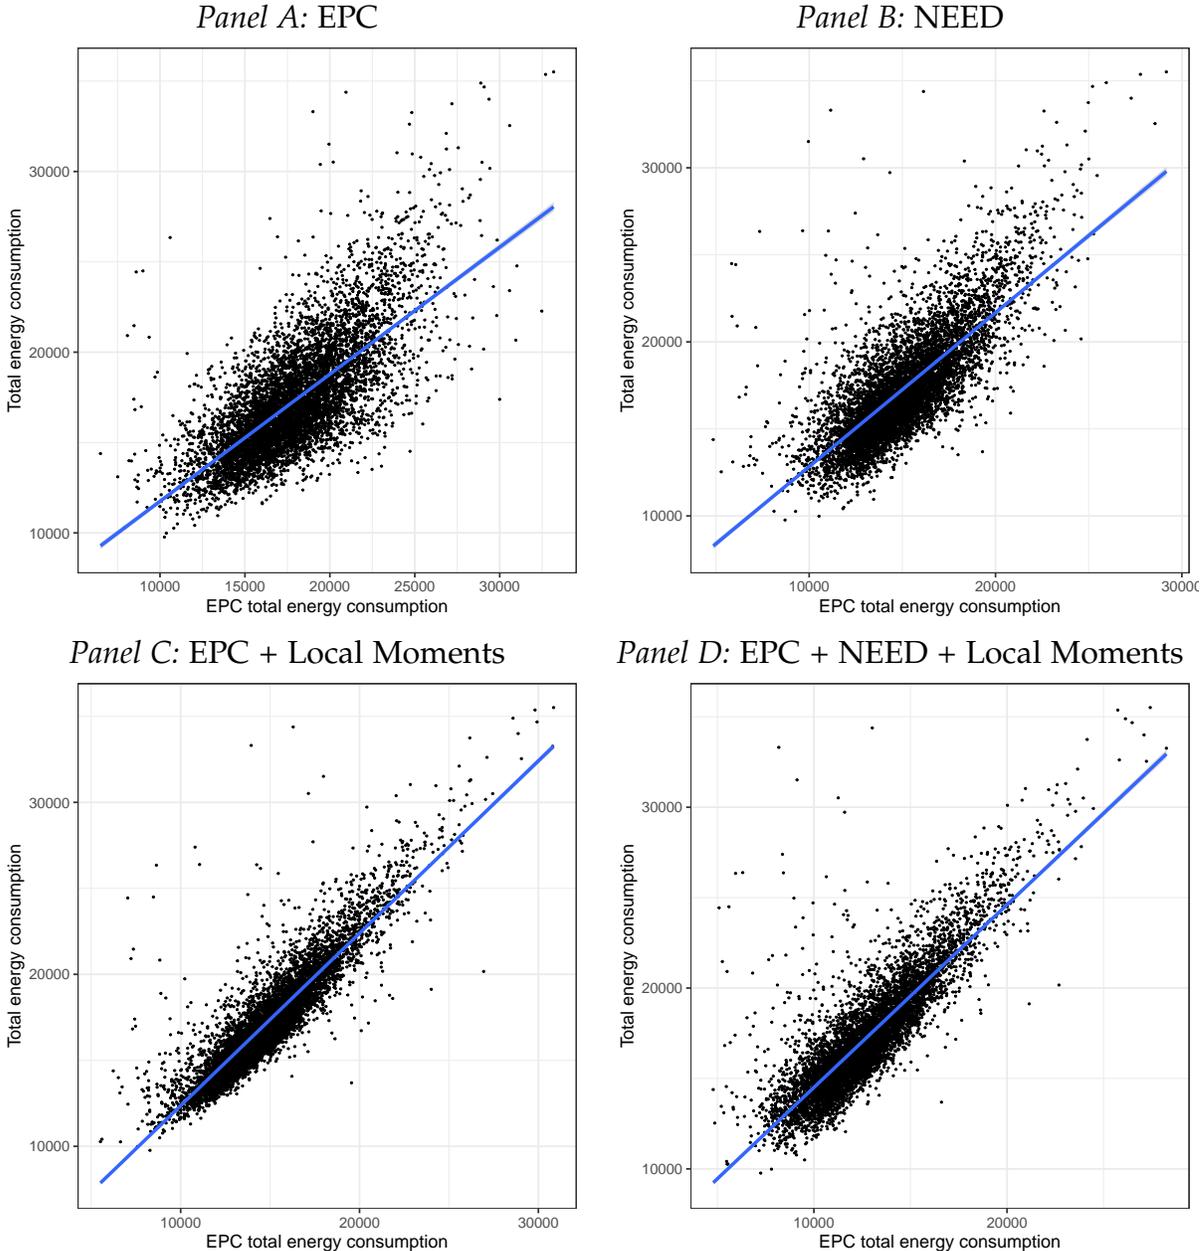
Figure A1: Schematic flowchart of the data processing pipeline to arrive at household-level energy price shock exposure measure

48



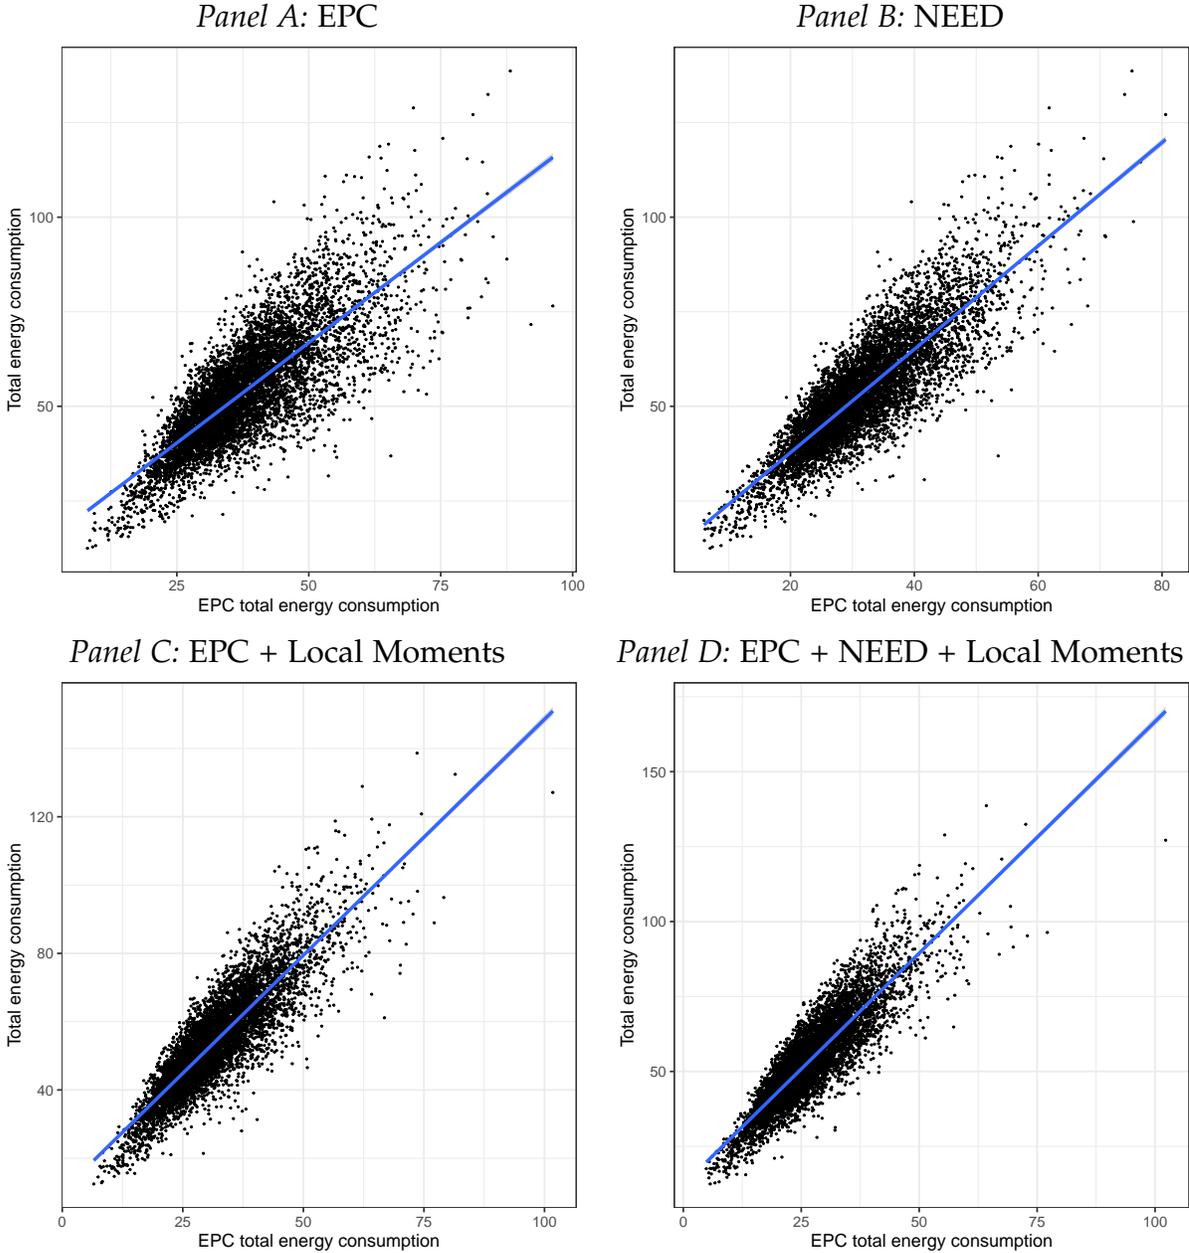
Notes: Figure provides a visual summary of the data construction process and the different steps and inputs that go into the derivation of the energy consumption and bill estimates.

Figure A2: Average property-level energy consumption measures at the MSOA-level compared with imputed energy consumption measures from EPC-NEED data



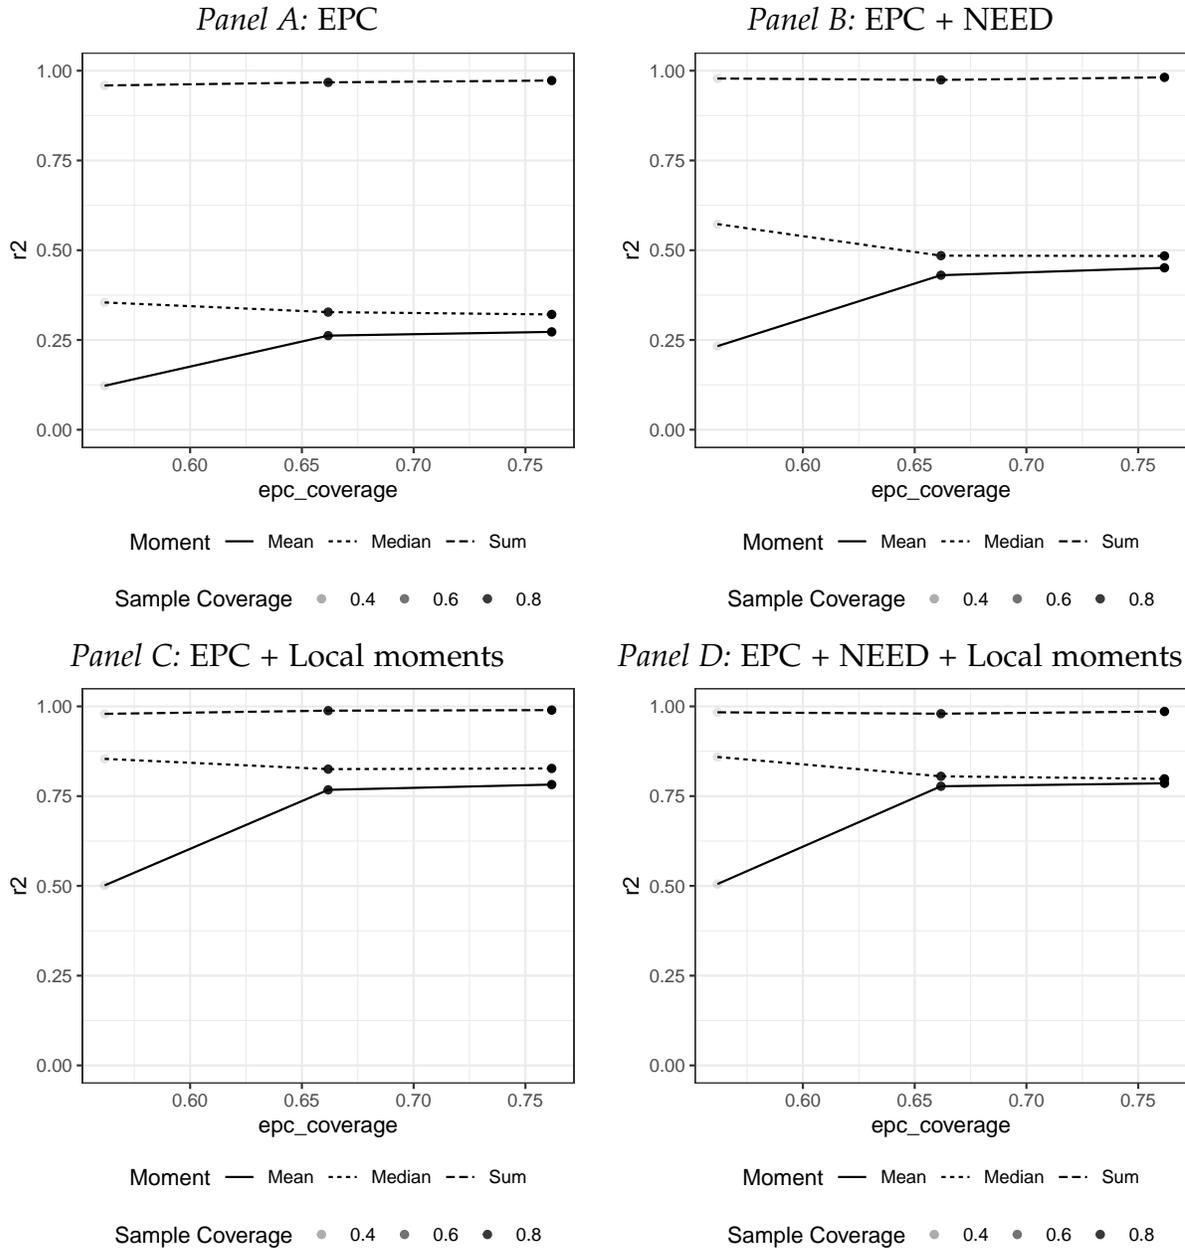
Notes: Figures provide a scatterplot of the mean energy consumption per meter estimates from published data at the MSOA level (for metered electricity and gas only) on the vertical axis and the average of various imputed energy consumption measures that leverage different data on the horizontal axis. Panel A provides the implied consumption estimates from the EPC data as is. Panel B augments the EPC data with a matching-of-moments approach based on anonymized individual level meter reading data collected under the NEED framework. Panel C uses the EPC raw energy consumption estimates and augments it with matched granular area-specific moments. Panel D is the final measure that combines the EPC raw data, the property-specific moment-matching and the local area specific moment matching.

Figure A3: Total property-level energy consumption measures at the MSOA-level compared with imputed energy consumption measures from EPC-NEED data



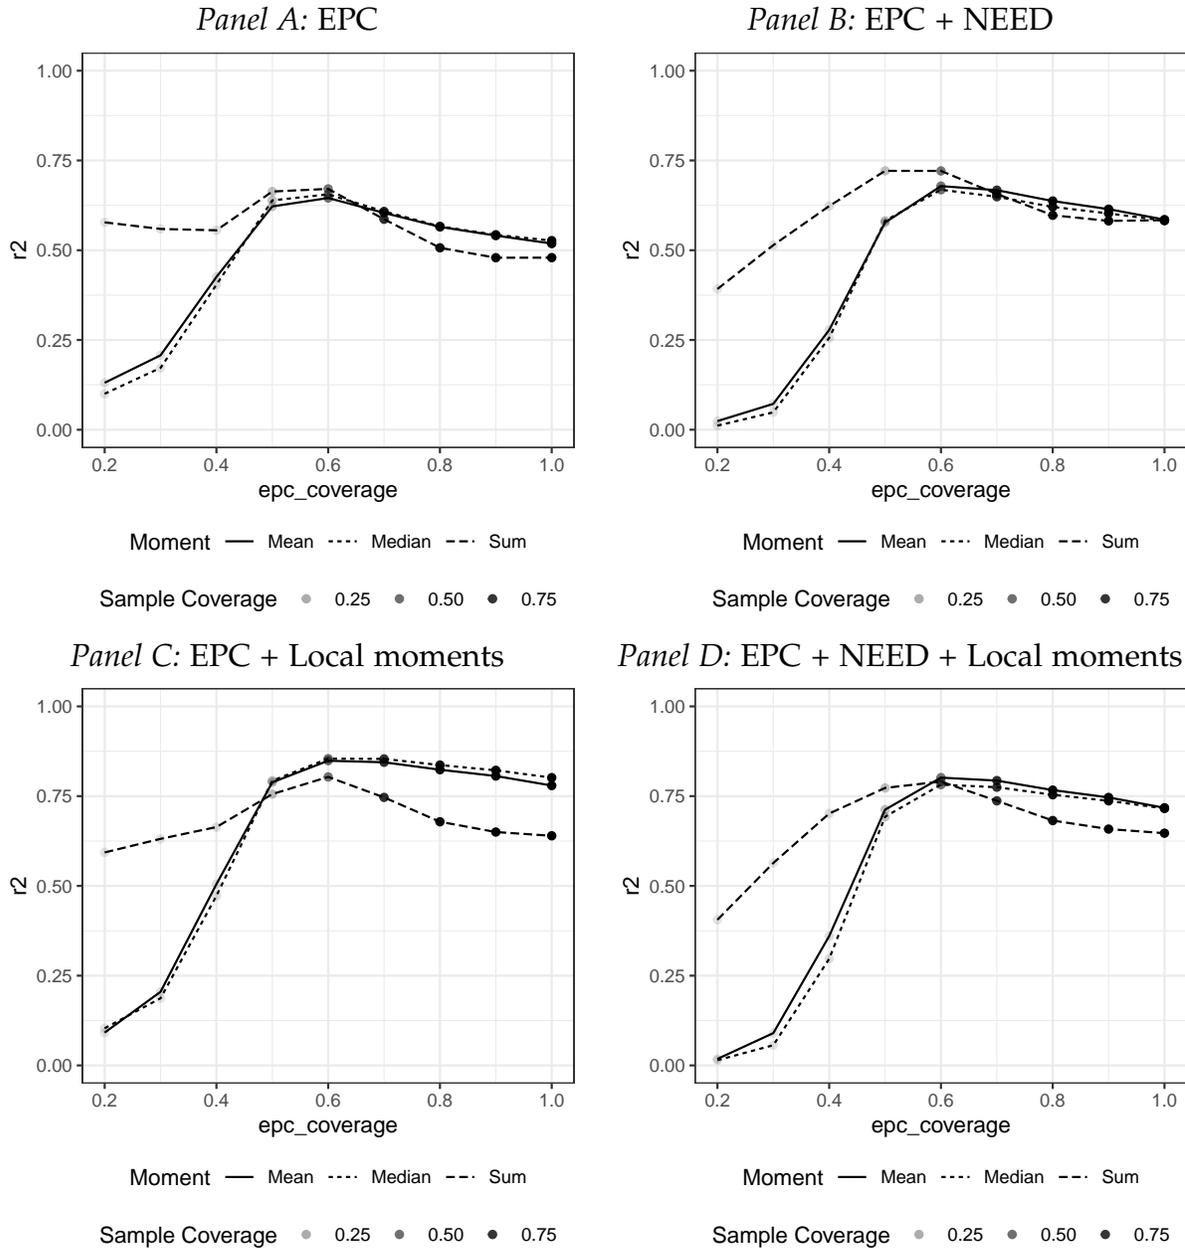
Notes: Figures provide a scatterplot of the total energy consumption per meter estimates from published data at the MSOA level (for metered electricity and gas only) on the vertical axis and the average of various imputed energy consumption measures that leverage different data on the horizontal axis. Panel A provides the implied consumption estimates from the EPC data as is. Panel B augments the EPC data with a matching-of-moments approach based on anonymized individual level meter reading data collected under the NEED framework. Panel C uses the EPC raw energy consumption estimates and augments it with matched granular area-specific moments. Panel D is the final measure that combines the EPC raw data, the property-specific moment-matching and the local area specific moment matching.

Figure A4: Illustration of goodness-of-fit of different derived consumption estimates at the Local Authority level



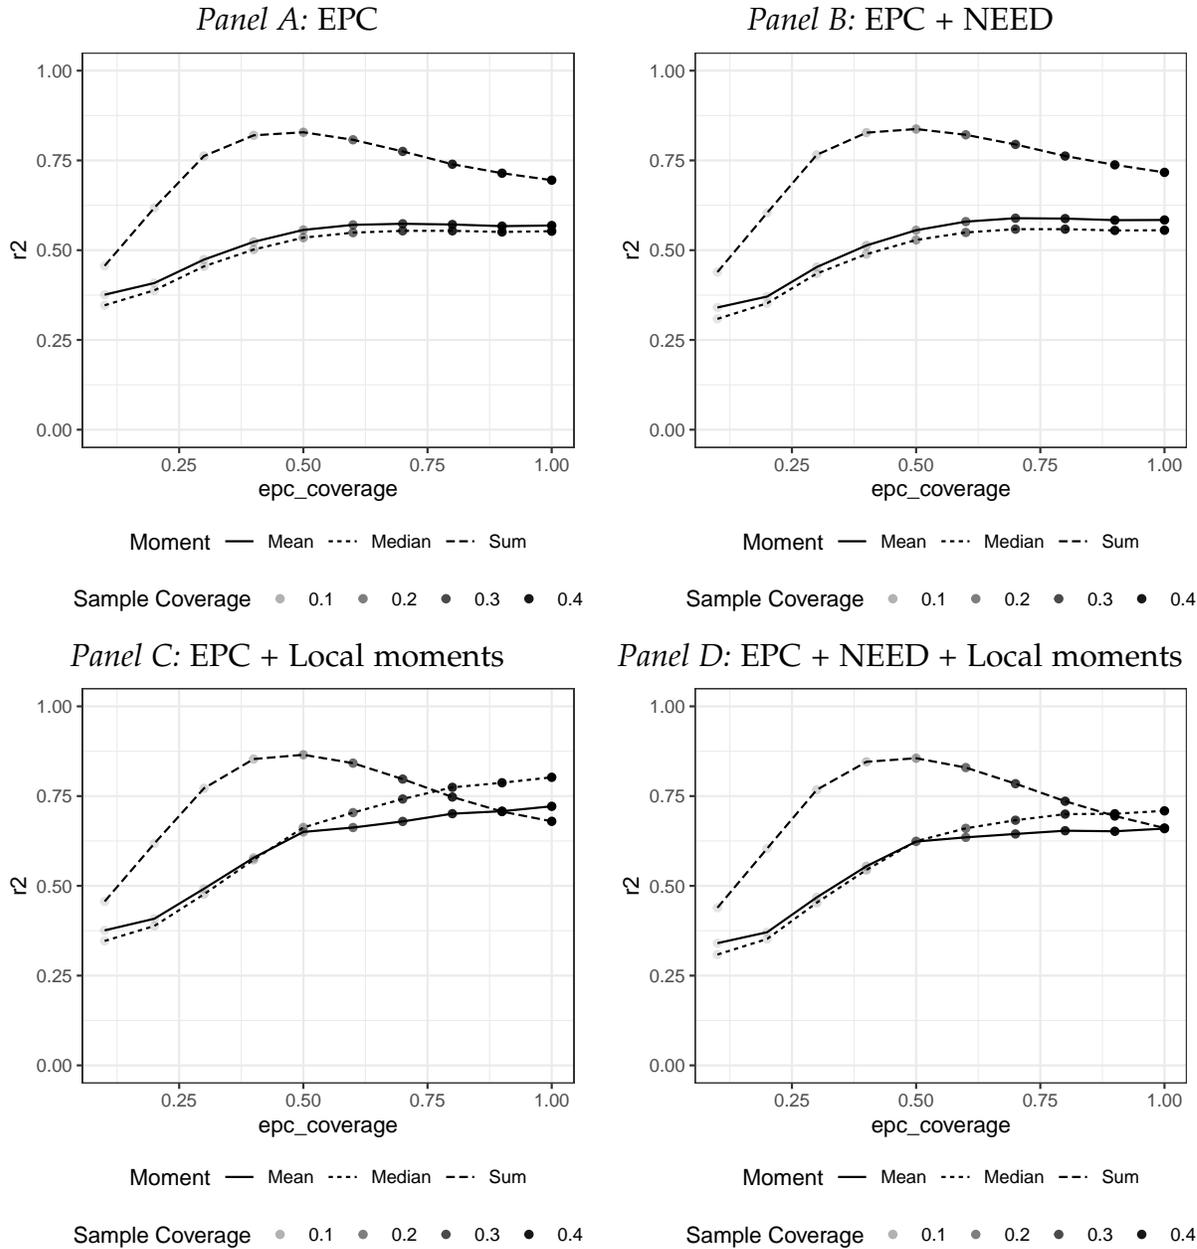
Notes: Figures plot the R^2 obtained from validating the derived implied consumption measures and three moments: the total consumption, the mean, and median consumption against actual consumption data that is published from gas and electricity meters across the country. For the four different derived measures, we compare the goodness-of-fit of the three moments against the corresponding moment from subnational statistics. The horizontal axis captures the ratio of the number of EPC properties against the population of properties in an area based on council tax data. A value of 0.4 on the axis implies that the estimating sample includes data from all districts that have at most 40% of their building stock captured in the EPC data. The EPC data covers at least 60% of the building stock in virtually all districts in England and Wales.

Figure A5: Illustration of goodness-of-fit of different derived consumption estimates at the LSOA level



Notes: Figures plot the R^2 obtained from validating the derived implied consumption measures and three moments: the total consumption, the mean, and median consumption against actual consumption data that is published from gas and electricity meters across the country. For the four different derived measures, we compare the goodness-of-fit of the three moments against the corresponding moment from subnational statistics. The horizontal axis captures the ratio of the number of EPC properties against the population of properties in an area based on council tax data. A value of 0.4 on the axis implies that the estimating sample includes data from all districts that have at most 40% of their building stock captured in the EPC data. The EPC data covers at least 60% of the building stock in virtually all districts in England and Wales.

Figure A6: Illustration of goodness-of-fit of different derived consumption estimates at the postcode level



Notes: Figures plot the R^2 obtained from validating the derived implied consumption measures and three moments: the total consumption, the mean, and median consumption against actual consumption data that is published from gas and electricity meters across the country. For the four different derived measures, we compare the goodness-of-fit of the three moments against the corresponding moment from subnational statistics. The horizontal axis captures the ratio of the number of EPC properties against the population of properties in an area based on council tax data. A value of 0.4 on the axis implies that the estimating sample includes data from all postcodes that have at most 40% of their building stock captured in the EPC data.

Appendix to “How large are the savings from energy efficiency upgrades in the UK?”

For Online Publication

A Step 1: Deriving a physical energy consumption measure for each property in the EPC data

An essential ingredient in our energy consumption calculations was the set of fuel prices faced by a given property for each type of energy consumption (space heating, water heating, and lighting). For example, while gas is the most common heating method across properties in the EPC data, many use either electricity or oil and therefore face different prices. Additional complexity follows from the range of possible tariffs used to price a household’s electricity use. The prices used in EPC calculations dating back to 2013 for all possible fuel types are published by BRE.¹ We had to infer which of these had been applied to each energy consumption type for each property in order to estimate expenditures.

To decide the assignment of fuel prices, we consulted four variables from the EPC database: main heating system (MAINHEAT_DESCRIPTION), water heating system (HOTWATER_DESCRIPTION), type of fuel used to power the central heating (MAIN_FUEL), and electricity tariff (ENERGY_TARIFF). For example, if the main heating system was recorded as “boiler and radiators, mains gas”, main fuel as “mains gas”, hot water system as “from main system” and energy tariff as “single”, a property was assigned the “mains gas” fuel price for space and water heating and the “standard tariff” price for lighting from the SAP price list. The raw data contain 9,796 unique combinations of these four variables, and so we restricted our attention to the 30 most common combinations, excluding those containing oil.² In total, these 30 combinations account for 85% of the sample. For the remaining 15%, we infer energy

¹Data are available here <https://bregroup.com/sap/standard-assessment-procedure-sap-2012/>

²We excluded properties using oil as there is no price cap for this fuel, which is our source of price data for gas and electricity (see the next section for details).

consumption using ENERGY_CONSUMPTION_CURRENT, a variable which estimates total energy consumption in kWh per metre squared of floor area. We scale this variable by codeFLOOR_AREA and multiply by the cost share of each energy use types to produce estimates of energy consumption for space and water heating and lighting.

We followed the SAP documentation to the best of our ability in the process of assigning fuel prices to energy consumption types, though in places the appropriate correspondence was not clear. Ambiguities also arose in interpreting how prices, which include a standing charge and price per kWh, had been applied to consumption to produce the spending estimates available in EPC data, complicating the reverse-engineering of this calculation. To account for this uncertainty, we have included a lower-bound estimate, which incorporates standing charges, as well as an upper-bound estimate, which excludes standing charges from consumption calculations. This inevitably introduces some measurement error, which we intended to tackle via spatial aggregation.

The consumption estimates we produce for water heating, space heating and lighting are *intention-to-treat* estimates, as the underlying physical SAP model used to produce the EPC data considers three factors:

1. The physical characteristics of a property, such as build-type, insulation technology, floor area, window area, number of rooms and light fixtures
2. Time-invariant climatic factors that affect fuel demand and are ultimately determined by property location
3. Fixed relationships between estimated use due to time-invariant estimates of likely occupation and use determined by the physical makeup of the property such as the number of bedrooms, the floor area, etc.

Consumption estimates stemming from EPC data will therefore not map one-to-one with consumption estimates which reflect actual patterns of energy consumption by residents, such as those produced from meter readings. Rather, EPC data is based on exogenous and typically fixed characteristics of the underlying buildings, a desirable feature for an econometrician. In this sense, our consumption estimates should be understood as *theoretical* as opposed to *real* consumption estimates.

B Step 2: Anchoring technically-required energy consumption measure with anonymized meter-level data

In step two, we produce a second consumption measure that incorporates anonymous micro data on energy and gas consumption from the National Energy Efficiency Data-Framework (NEED).³ We refer to these as *real* consumption estimates as, unlike the EPC-based estimates, they reflect patterns of energy consumption behaviour by households. The sample includes four million properties and is designed to be representative of domestic properties in England and Wales. Data are available annually for years 2005-19, of which our analysis uses 2017-19. The data include estimates of energy and gas consumption which are derived from meter readings, alongside a number of property and area-level characteristics.

We use this *real* consumption data to develop a refinement of the *theoretical* consumption measure derived in Step 1. We match moments of the NEED meter-level data with moments from the EPC-derived consumption measure, in effect rescaling our theoretical consumption measure. This is possible because the NEED data include a range of property characteristics which are also present in the EPC data:

- property type (six categories)
- property age band (four categories)
- an indicator for whether gas is the main heating fuel (two categories)
- floor area band (five categories)
- a measure of the relative deprivation of the area in which a property is located, measured in 2019 (five categories each for Wales and England)
- region (nine categories for England, one for Wales).

In theory, there are $6 \times 4 \times 2 \times 5 \times 5 \times 9 = 10,800$ unique combinations of these features in England and $6 \times 4 \times 2 \times 5 \times 5 = 1,200$ unique combinations in Wales.

For each unique combination, we calculate the deciles of combined gas and electricity consumption in the NEED data, excluding combinations which contain

³The data can be found here <https://www.gov.uk/government/statistics/national-energy-efficiency-data-framework-need-anonymised-data-2021>.

300 properties or fewer out of the total 12 million (4 million for each of the years 2017-19).

We then replicate this exercise using the EPC consumption estimates derived in Step 1. When calculating total consumption, we take weighted averages of the upper and lower bounds for our light, water, and space energy consumption estimates, before summing over these to derive aggregate energy consumption. The weight assigned to the upper bound of each consumption estimate is 5 minus the floor area band (1-5), meaning a higher weight is assigned to the lower bound for larger properties.

Next, we match the NEED energy consumption deciles for each unique combination of property attributes to the corresponding EPC energy consumption deciles. For example, a property that is in the top decile of *theoretical consumption* (derived from EPC data) among properties with the same combination of property attributes will be assigned the top decile of *real consumption* (derived from NEED data) for properties with these same attributes. The latter provides us with a potentially more accurate representation of real consumption behaviour at the property level.

We then update our property-level estimates of theoretical consumption by multiplying by the ratio of real to theoretical energy consumption (both actual and potential) for a property's attributes and consumption decile.⁴

We then perform a second rescaling using postcode-level gas and electricity consumption data, again for the years 2017-19.⁵ For each postcode, we compute the sum of median gas and electricity consumption across years. We then repeat this exercise for the theoretical consumption estimates developed in Step 1 as well as the estimates which were adjusted using NEED data. Next, we rescale the property-level theoretical and NEED-adjusted consumption estimates by multiplying by the ratio of the median postcode-level energy consumption from the postcode data to the corresponding value in the EPC data. We perform this rescaling of theoretical consumption estimates only for properties in postcodes with at least 25% coverage in the EPC data. Here, coverage is defined as the number of properties per postcode

⁴Note that energy consumption estimates for properties whose combinations of attributes included less than 300 properties were not rescaled.

⁵The electricity data can be found on <https://www.gov.uk/government/collections/sub-national-electricity-consumption-data>; the gas data is on <https://www.gov.uk/government/collections/sub-national-gas-consumption-data>.

in the EPC data relative to the number of energy meters used to form the energy consumption estimates in the postcode data.⁶ We then perform this same rescaling for NEED-adjusted consumption estimates. We exclude from both rescaling exercises properties in postcodes with five or fewer properties in the EPC data or five or fewer energy meters in the postcode consumption estimates. For properties in postcodes which fail these coverage requirements, we rescale consumption using the same methods but with LSOA-level data. Here, we impose a looser restriction of 50% coverage of EPC properties in a property's LSOA.

C Step 3: Converting consumption measures to time-varying spending estimates

In our third step, we convert the *time-invariant* consumption estimates from Steps 1 and 2, measured in kWh, into *time-varying* estimates of actual spending in GBP. In practice, this is not straightforward as the energy prices faced by households, which consist of a unit price and standing charge, depend on the particular energy supply contract which they have entered into.

We are interested in four types of price scenario:

1. Energy price cap.

The energy cap sets the maximum price that energy suppliers are allowed to charge customers, and is chosen by regulator Ofgem for gas and electricity prices to reflect the costs of supplying energy and to allow modest profits (Ofgem, 2022a). The cap has been updated every 6 months since its introduction in January 2019, but from October 2022 will be updated on a quarterly basis. The price cap was originally conceived to protect inattentive consumers from being charged unfair rates. In its early years, some energy contracts on the market were cheaper than the cap, but since the summer of 2021 the cap has been the cheapest rate available. This phenomenon is due to price increases between the time at which the price cap is set and the time at which it comes into effect (as of October 2022, this gap has been shortened from two

⁶The postcode-level data includes the number of meters used to form the estimates of median gas and electricity consumption respectively, and we use the highest of these two figures.

months to 25 working days) (Ofgem, 2022b). As such, the cap has been a more accurate reflection of the prices faced by households in recent months than in previous years. Our study incorporates price cap values from October 2021 and October 2022.

2. Energy Price Guarantee.

In September 2022, the UK government announced the Energy Price Guarantee programme as a response to the ongoing energy crisis. The EPG reduces the maximum per unit rate below the level of the October 2022 price cap in an attempt to limit the average household energy bill to around £2,500. As discussed in Fetzer (2022), the standing charge is maintained at the level of the October 2022 price cap.

3. Historical average energy prices.

The Department for Business, Energy and Industrial Strategy (BEIS) publishes data on average gas and electricity prices for 2010-2021.⁷ These data are particularly valuable for estimating energy bills pre-2019, when the energy price cap had not yet been introduced.

4. Two-tier tariff.

This is an alternative policy proposal to the energy-price guarantee that is discussed in more detail in Fetzer (2022). It consists of a two-tier tariff wherein the standing charge would be fixed at the level of the October 2021 price cap, as would unit prices for the first 9,500 kWh of natural gas consumption and the first 2,500 kWh of electricity consumption. As 50% of UK households consume less than 12,100 kWh of natural gas and 2,900 kWh of electricity, this would drastically limit energy price increases for the bulk of households.⁸ The second tier of the energy tariff would be set at steeper levels which could be aligned with the EPG. For example, a second tier unit price of 20 pence per kWh for natural gas and 60 pence per kWh for electricity, together with the first tier described above, would have a similar cost to the government as the

⁷Data are available here <https://www.gov.uk/government/statistical-data-sets/annual-domestic-energy-price-statistics>

⁸See <https://www.gov.uk/government/statistics/national-energy-efficiency-data-framework-need-consumption-data-tables-2021>.

EPG. This would offer much more targeted support without undermining the incentive to save energy created by higher unit prices.

Energy prices consist of a standing charge and unit rate which differ according to region, payment method (for example, direct debit versus pre-paid), fuel type (electricity versus gas), and electricity metering arrangement (whether the electricity tariff varies by time of use). We use information on these dimensions from the EPC data to assign the appropriate price to each property. In the absence of data on payment method, we assume direct debit for all households.⁹

We then estimate energy expenditure for a given property and energy use (space heating, water heating, or lighting) as follows:

$$spend_{ierfmt} = cons_{ierfm} \times price_{rfmt} + charge_{rfmt}$$

Here, $cons_{ierfm}$ is the consumption estimate for energy use type e by property i in region r with fuel f and metering arrangement m , as calculated in Steps 1 and 2. $price_{rfmt}$ and $charge_{rfmt}$ are the unit price and standing charge at which the cap has been set for their region, metering arrangement and fuel in period t (assuming payment by direct debit).

In essence, this spending calculation converts the intention-to-treat consumption estimate in physical energy units, which reflects the physical characteristics of a property, back into energy cost estimates that are, in turn, exogenous to household-specific choices with respect to their energy supply contract. This data structure is also ideal for merging in different price scenarios to forecast their likely impact on household bills across different groups and regions within the UK.

Most households are on one-year fixed term contracts at the energy supply contracts.

⁹Direct debit is the most popular payment method (Ofgem, 2019).

D Step 4: Energy efficiency upgrade recommendations and its costing

Lastly, we also examine the specific energy efficiency upgrade investments recommended in the EPCs. We were not able to confirm how the costing of these recommendations is done. We thereby convert the estimates of the costs for specific measures, which typically include an upper- and a lower-bound range to a further upper and lower bound based on an inflation rate estimate.

To do so, we construct a version of the cost estimate that is expressed in current GBP which effectively updates the upper- and lower-bound cost estimate by what we judged to be the most appropriate inflation rate from the lodgement date (the date the EPC was drawn up) to the current date.

E Data used for correlational analysis

The following covariates were sourced from the 2011 UK census:

Category	Covariates
Demographics	Highest qualification obtained, ethnicity, county of birth, age, household size, self-reported health, disability
Housing	Tenure, population density, dwelling type, method of commute
Economic activity	Economic activity, industry of employment (2007 SIC)

These data were supplemented by the following variables:

Household income. Model-based income estimates at the MSOA-level are produced by the Office of National Statistics (ONS).¹⁰ Our analysis used estimates of average total annual income for the year 2018.

Fuel poverty. Annual statistics on the number of individuals in fuel poverty at the LSOA-level are produced by the Department for Business, Energy & Industrial

¹⁰Data are available here <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/datasets/smallareaincomeestimatesformiddlelayersuperoutputareasenglandandwales>

Strategy (BEIS).¹¹ These adopt the Low Income Low Energy Efficiency (LILEE) metric of fuel poverty, which considers a household fuel poor if it lives in an energy inefficient property and has disposable income below the poverty line. Our analysis uses figures for the year 2020. These were aggregated from Lower Layer Super Output Area (LSOA) level to the MSOA level using population estimates.

Median house prices. Data on median house prices by MSOA are published by the ONS, calculated using open data from the HM Land Registry.¹² Our analysis uses figures for the year ending March 2022.

Index of Multiple Deprivation (IMD). English Indices of Deprivation (IoD) are published by the Department for Levelling Up, Housing & Communities.¹³ These are relative measures of deprivation which incorporate the seven following domains: income; employment; health deprivation and disability (in our analysis, we refer to this as health); education, skills and training (education); crime; barriers to housing and services (housing and services) and living environment. Our analysis uses rankings along these dimensions for each LSOA for the year 2019. These were aggregated to the MSOA-level using population estimates.

¹¹Data are available here <https://www.gov.uk/government/collections/fuel-poverty-sub-regional-statistics>

¹²Data are available here <https://www.ons.gov.uk/peoplepopulationandcommunity/housing/datasets/hpssadataset2medianhousepricebymsoaquarterlyrollingyear>

¹³Data are available here <https://www.gov.uk/government/statistics/english-indices-of-deprivation-2019>