The Heterogeneous Effects of Supply Shocks in Necessity Goods*

Tiago Bernardino†
IIES, Stockholm University

Pedro Brinca‡
Nova SBE

Saman Darougheh§
Danmarks Nationalbank

Márcia Silva-Pereira¶
Banco de Portugal and Nova SBE

October 2025

Abstract

We study the efficacy of price subsidies — implemented via lower value-added taxes — in addressing sectoral supply shocks in the presence of non-homothetic preferences and imperfect pass-through. In our heterogeneous-agent model, sectoral shocks generate recessions in both necessity and luxury goods, with necessity recessions imposing larger costs on the bottom 30% of the income distribution. A subsidy that lowers the price of necessity goods is welfare-reducing: high-income households respond disproportionately even though these goods account for a smaller share of their consumption. In contrast, cash transfers reduce the welfare losses associated with sectoral supply shocks. Central to our findings is that the model aligns closely with the empirically observed non-homothetic behavior of households and the imperfect pass-through, which we also document.

Keywords: Nonhomothetic Preferences; Incomplete Pass-through; Heterogeneous Agents; Sectoral Shocks; Price Subsidies

JEL Classification: E21; E32; E62

---

*T. Bernardino acknowledges financial support from the Jan Wallanders & Tom Hedelius Foundation (H21-0020). P. Brinca acknowledges funding from Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019, PTDC/EGE-ECO/7620/2020, and Social Sciences DataLab, LISBOA-01-0145-FEDER-022209), POR Lisboa (LISBOA-01-0145-FEDER-007722, LISBOA-01-0145-FEDER-022209), POR Norte (LISBOA-01-0145-FEDER-022209) and CEECIND/02747/2018. M. Silva-Pereira acknowledges financial support from Fundação para a Ciência e a Tecnologia (2021.04898.BD). The views expressed herein are those of the authors and should not be attributed to any of the affiliated institutions.

†Email: tiago.bernardino@iies.su.se.
‡Email: pedro.brinca@novasbe.pt.
§Email: sdma@nationalbanken.dk.
¶Email: mlspereira@bportugal.pt.

1 Introduction

Between 2021 and 2023, negative supply shocks affected countries around the globe, increasing the prices of food, electricity, and other essential goods (see Figure 1). Many European countries responded by subsidizing prices.[1] They argued that the types of goods affected were "necessary goods" — goods which, to some extent, must to be consumed, are difficult to fully substitute away from, and consequently are an important part of the consumption basket of low-income consumers. Therefore, subsidizing the price of necessity goods by means of lowering value-added taxes (VAT) should be an effective way of supporting low-income households.[2]

Figure 1: Relative price of food and energy vs. other goods

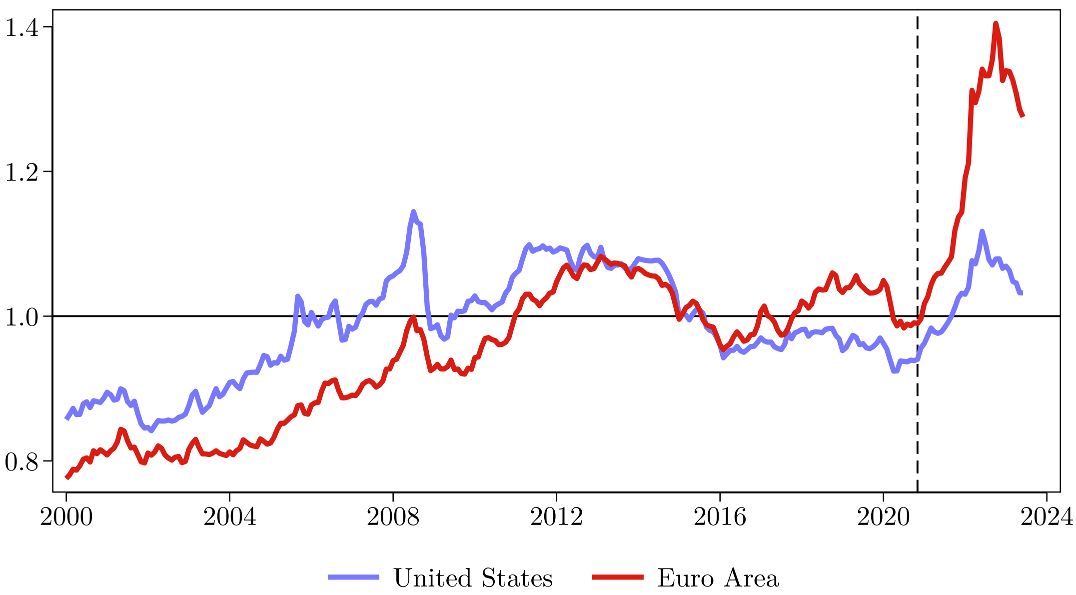

Notes: This figure shows the ratio of the price of food and energy relative to non-food and non-energy goods in the US (red line) and in the Euro Area (blue line). The numerator is computed as the weighted average of the CPI series of food and energy items. The denominator corresponds to the weighted average of non-food and non-energy CPI items. The weights correspond to CPI weights.

In addition to the conventionally recognized distortions arising from subsidizing prices (and financing this intervention), we highlight two additional caveats in this context. First, in practice, many goods turn out to be both necessities and luxuries: electricity may be necessary to run a household, but can at the same time be used to heat outdoor pools.[3] Lowering the price of necessity goods can, therefore, subsidize the consumption of high-income individuals — an unwanted and expensive side effect. Second, the

[1] 22 out of the 27 EU countries implemented price subsidies in the form of value-added tax cuts between 2022 and 2024, mostly on energy, gas, and food items, according to the OECD (2023).

[2] This is especially the case when some countries such as Germany are limited in their ability to employ targeted cash transfers.

[3] Indeed, the French government subsidized electricity but asked citizens to stop heating their pools in the spring of 2023.

pass-through of VAT on final prices is often thought to be imperfect, which adds to the potential inefficiencies stemming from the deadweight loss of taxation.

In this paper, we take these arguments seriously by studying empirically and with a quantitative model the effects of sectoral shocks that raise the relative price of necessity goods and the effects of a government intervention that subsidizes the price of necessities. We make three contributions. First, we empirically investigate the foundation underlying the two caveats mentioned above: we study the consumption patterns of necessity and luxury goods across the income distribution, and estimate the imperfect pass-through of VAT changes in European countries. We then build a heterogeneous-agents model with nonhomothetic preferences that is consistent with our empirical findings.Our second contribution is the comparison of recessions in the luxury and necessity sector, which play out differently in our model, with higher welfare costs associated with the latter, in particular for low-income workers. Third, we then study the role of VAT subsidies in alleviating the welfare costs in the context of recessions in the necessity sector. We now discuss these contributions in detail. Finally, we study fiscal policies that dampen the negative effects of the recession. Here, we compare the price subsidy to untargeted and targeted transfers. The advantage of a price subsidy to the necessity good might be that poorer households are more exposed to the necessity good price, and that thus subsidizing necessity goods is a good substitute for a targeted transfer in countries where such transfers are technologically infeasible. However, due to the elastic demand of high-income individuals for the price of necessities, even the unconditional transfers outperform the VAT subsidy.

First, we estimate nonhomothetic preferences in the United States by categorizing consumption categories as either necessities or luxuries depending on whether the income elasticity of expenditures was above or below one. We then demonstrate one of the drawbacks of price subsidies: the highly elastic demand curves of high-income individuals. We do so by estimating the response of expenditure composition for households across the income distribution after a relative price shock to necessities, instrumented with oil supply news shocks. We find that all households respond by increasing the budget share allocated to necessity goods, with the response increasing in income: the policy disproportionately increases the consumption of workers who were not intended to be its main beneficiaries. We also demonstrate the imperfect pass-through of value-added taxation. We combine data on value-added taxes and prices at the country-category-month level. We find that roughly one fourth of VAT changes are transmitted to consumer prices.

Second, we build a quantitative model to study the role of value-added taxation in the context of nonhomothetic preferences and imperfect pass-through. At the core of our model, households face idiosyncratic productivity shocks which translate to income volatility that they can insure against by saving in an asset with fixed positive supply (as in Huggett, 1993). Two goods exist in this economy, "necessities" and "luxuries".

We employ nonhomothetic preferences following Comin, Lashkari, and Mestieri (2021) which allow us to match the extent to which expenditure shares in necessities fall in income.[4]. Each good is produced by a sector comprised of a continuum of firms. To generate the imperfect pass-through of price subsidies to prices, we assume decreasing-returns-to-scale production function and monopolistic competition.

We introduce a negative sectoral productivity shock to either the necessities or the luxuries sector. Both productivity shocks lead to recessions that feature an increase in the relative price of the affected sector, a fall in consumption, and an increase in labor supplied. Our preliminary findings are that both recessions come with considerable welfare effects: workers in the first income decile would be willing to give up 14% (7%) of their quarterly income to prevent the necessity-sector (luxury-sector) recession: lower-income households suffer more in necessity-recessions than in luxury recessions (the opposite is true for higher-income households).

Finally, we study the extent to which fiscal policies can be an effective tool in the context of necessity-good recessions. We study the effects of a price subsidy (as implemented in many countries via a temporary lowering of value-added taxes) comparable in magnitude to that implemented by Portugal or Spain in 2023 and compare it to targeted and untargeted transfers. All interventions have the same budgetary cost and are financed with a proportional labor tax.

We find that price subsidies actually reduce aggregate welfare despite minimizing consumption losses. This welfare loss occurs because the price subsidy benefits all households, including those less affected by the shock, while being financed through distortionary taxes that generate a deadweight loss.

In contrast, both types of transfers are welfare-improving, which is mostly driven by the welfare gains of the bottom third of the income distribution. Targeted transfers emerge as the dominant policy, delivering average welfare gains of 2.3% of total income — substantially outperforming both lump-sum transfers (0.7%) and price subsidies (-0.2%).

Taking a step back, the extent to which price subsidies are actually welfare reducing will depend on the exact quantitative estimation. We do however tentatively conclude that for countries that are unable to target their cash transfers to low-income households, untargeted transfers are still a better fiscal tool than price subsidies.

Related literature. We contribute to the burgeoning literature on the role of nonho-

[4]Notably, the commonly used Stone-Geary preferences fall short: they imply a steep fall in the expenditure share of necessity goods among the first income deciles, followed by approximately constant expenditure shares for higher-income households.

mothetic preferences in the heterogeneous propagation of shocks. Engel (1857) first documented the negative relationship between income and the budget share allocated to necessity goods, an empirical regularity widely supported by the data. While consumption inequality has long been studied empirically (Aguiar and Bils, 2015; Attanasio and Pistaferri, 2016), only in recent years has income-dependent demand gained prominence as a key channel in macroeconomic research (Boppart, 2014; Comin, Lashkari, and Mestieri, 2021).

Hochmuth, Pettersson, and Weissert (2023) explore the consequences of consumption inequality on prices. They construct a nonhomothetic price index, providing a more precise measure of the cost of living, and find that inflation is 2.5 times more volatile for lower-income households than for higher-income households. Orchard (2024) documents that necessity prices rise more during recessions, relative to the prices of luxury goods, and that the share of necessity goods is countercyclical. Relatedly, Andreolli, Rickard, and Surico (2024) show that spending on non-essential goods is more sensitive to business cycles than on essential goods, which ultimately amplifies the effects of monetary policy. In the same spirit, Sonnervig (2023) shows that lower-skilled households — which likely correlates to lower-income households in our setting — experience smaller consumption changes over the business cycle. We contribute to this literature by understanding how nonhomotheticity shapes the aggregate and distributional responses to shocks, in our case focusing on a sectoral supply side shock.

We also contribute to the literature on optimal policies, particularly on the role of fiscal policy in mitigating the welfare costs of macroeconomic shocks. Correia et al. (2013) show that a combination of consumption and labor taxes can mimic the effect of monetary policy stimulus. More recently, Carroll et al. (2023) compare the effectiveness of different policy responses — unemployment insurance (UI), stimulus checks and a cut in labor taxes — after a negative income shock, concluding that UI is the more cost-effective policy. In our case, we look at the welfare effects of the fiscal stimulus policies most commonly implemented during the recent inflationary shock — lump-sum transfers, targeted transfers and a sectoral price subsidy to the sector directly affected by the shock.

Roadmap. The remainder of the paper is organized as follows. In Section 2 we describe the two empirical findings on nonhomothetic preferences and imperfect pass-through. Section 3 shows, in a simple environment, how nonhomothetic preferences on the household side and a decreasing-returns-to-scale production technology manages to yield model outcomes that are consistent with the empirical analysis. Section 4 presents the dynamic general equilibrium model with incomplete markets that we use to study the quantitative effects of a sectoral recession and the government policies effects. In Section 5 we discuss how we calibrate this model. Section 6 presents the effects of a sectoral recession both at the aggregate level and at the individual level. Section 7 compares the effects of three alternative fiscal policies. Section 8 concludes.

2 Empirical Analysis

Many policymakers have frequently cited the nonhomothetic nature of household preferences as a key motivation for fiscal policies during the inflationary episode of 2022. They argued that rising prices for necessities like food and energy disproportionately burden low-income households, since these goods constitute a larger share of their consumption baskets. However, using VAT reductions to subsidize essential goods raises a critical question: to what extent would such tax cuts actually be passed through to final consumers rather than captured by producers and retailers? In this section, we will address both claims.

First, we document the nonhomothetic nature of household preferences by examining how different income groups adjust their expenditure composition when facing a shock to the relative price of necessity goods. We find that high-income households are more responsive to price shocks than middle-income and lower-income groups — a pattern that is consistent with nonhomothetic preferences but difficult to replicate with Stone-Geary preferences (Matsuyama, 2023). This motivates the use of nonhomothetic CES preferences by Comin, Lashkari, and Mestieri (2021) when modeling household consumption choices — a more general case of nonhomothetic preferences with a constant elasticity of substitution not equal to one — in order to match these consumption patterns.

We begin by estimating semi-elasticities of category-level budget shares to income. This then allows us to categorize a given good as a necessity or a luxury good and motivates the need to model household preferences through the lens of a nonhomothetic structure, particularly to understand the effect of a sectoral supply shock. We then document the change in the composition of expenditure between necessity and luxury goods following an unexpected increase in the relative price of necessity goods. Since relative price changes are endogenous to changes in consumption and our goal is to estimate the short-term response after the price shock, specifically one that is supply-side driven, we instrument the change in the relative price of necessities' with an oil supply news shock. We show that this shock is instrument relevant, since it affects mostly the price of necessity goods, and argue that it satisfies the exclusion restriction.

Second, we examine the pass-through of price subsidies. We find that value-added tax changes are, on average, imperfectly transmitted to consumer prices. This aligns with standard producer theory, where pass-through rates can range from zero to 100% depending on the relative elasticities of demand and supply (Benzarti, 2024). Imperfect pass-through reduces the efficacy of sectoral price subsidies as a fiscal tool for mitigating welfare costs from sectoral shocks. We incorporate this empirical finding in our quantitative model.

2.1 Data

This section describes the data sources for our two empirical exercises. First, we use US expenditure-price data to document nonhomotheticity in household preferences. Second, we use EU price-VAT data to establish that VAT pass-through is, on average, imperfect.

2.1.1 US expenditure-price data

Expenditure data. We use US household-level quarterly consumption data for the period 2000-2023 from the Consumption Expenditure Survey (CEX) published by the Bureau of Labor Statistics (BLS). The CEX tracks spending in all product categories such as food, housing, utilities, transportation, health, or education, which are then used as weights for the Consumer Price Index (CPI). It has two different datasets: the interview survey and the diary survey. The interview survey is of quarterly frequency and consists of a rotating panel in which, in each quarter, between 5000 and 8000 households are interviewed, and each household stays in the sample, at most, for five consecutive quarters. In the initial interview, interviewers collect information on households' sociodemographic information and, in the subsequent four interviews, they collect expenditure data for the three months prior to the month of the interview. It is used more frequently in the literature since it captures a broader range of expenditures (Aguiar & Bils, 2015; Hobijn & Lagakos, 2005). The diary survey, on the other hand, only interviews each household once regarding their expenditures on a detailed list of goods over a period of two weeks. We use the interview survey since it contains a cross-sectional dimension, even if limited to at most four quarters, to estimate short-run consumption responses.

We aggregate consumption expenditure into 22 categories, closely following Hobijn and Lagakos (2005). We deviate from their aggregation only in that we exclude owner-occupied and rental housing to, first, ensure that our results are not driven by housing market fluctuations[5] and, second, since these are less likely to change contemporaneously in response to a relative price shock. We weigh different households according to their representativeness in the population using the sampling weights included in the CEX.

We exclude observations with zero or negative total expenditure. Following Aguiar and Bils (2015), we exclude households where the head of household is younger than 18 years old. As previously mentioned, in the interview survey, households are interviewed on

[5]Ex-ante, the concern would be that housing has big weights in consumption baskets and might thus be the main driver behind any empirical finding. We find, however, that the results are robust to the inclusion of housing.

their expenditures for up to four quarters. Since we will want to analyze how different households respond to changes in income or relative prices, we will be leveraging the cross-sectional dimension of the survey. Thus, we keep only households that stay in the sample for the maximum period, leading to a total of 76,484 households over 2000-2023.

Price data. We use Consumer Price Index (CPI) data produced by the BLS at a quarterly frequency for each of the selected categories from the CEX. The mapping between the CEX and CPI data is presented in Table XX. In the CEX's interview survey households are asked to report what they have consumed in a given category (e.g., food at home) over the three months before the interview. This means that, in order to have a correct mapping between expenditures and prices, we should match those expenditures with the average of the category-specific CPI during the three months before the interview as well. Thus, we construct 3-month moving averages of those series for the months preceding the reporting month, which are then matched to the expenditures reported in the same month. The price indices are based in the first quarter of 2000.

Oil shocks. We use identified oil supply and oil supply news shocks from Känzig (2021). Using high-frequency identification, Känzig (2021) identifies oil supply surprises by observing changes in the price of oil futures in a tight window around OPEC announcements about oil supply, thus isolating the impact of news about future oil supply. The author then proceeds to use the resulting series in a Vector Autoregression (VAR) model in order to identify a structural oil supply news shock.

In order to match this series to the previously constructed CEX-CPI dataset for 2000-2023, since the shocks are at the monthly level, we construct a 3-month rolling sum of the oil supply news shocks. The rationale for this is the same as previously explained for prices: since, in month t, we observe the sum of expenditures for months $t-1$, $t-2$ and $t-3$, we need, in month t, average prices for those three months and the sum of shocks occurring in the three months prior to month t as well. We then proceed to combine this 3-month rolling sum with the CEX-CPI dataset.

2.1.2 EU price-VAT data

Price data. We use monthly price data from Eurostat's Harmonized Index of Consumer Prices (HICP) at the COICOP5 level (*Classification of Individual Consumption by Purpose*), covering 2000-2019 to match the available VAT data period.

VAT rates. We merge the price data with the historical VAT rates database compiled by Benzarti et al. (2020) and extended by Benzarti and Tazhitdinova (2021). This database contains all value-added tax changes by product category for each European country over the same period.

2.2 Classification of consumption goods

Methodology. We use the expenditure data from the CEX described in Section 2.1.1 to estimate how households change the budget shares allocated to different consumption categories after a change in income. If household choices are consistent with homothetic preferences, budget shares should be independent from income, meaning that a change in income would not impact the expenditure share of any given good in households' total expenditure.

In order to categorize goods as a necessity or a luxury good, we estimate semi-elasticities of expenditure shares for each category to income using the following specification, estimated using an OLS estimator:

$$\omega^i_{h,t} = \alpha^i + \beta^i \log(y_{h,t}) + \delta_t + \mathbf{X}_h + \epsilon^i_{h,t}, \quad i \in [1,22] \cap \mathbb{Z}, \tag{1}$$

where $\omega^i_{h,t}$ is the budget share of category i for household h at time t, $y_{h,t}$ is the quarterly pre-tax income of household h at time t, δ_t are time fixed effects and $\mathbf{X}_h$ are a set of household-specific controls comprising marital status, gender of the head of household, age of the head of household, education level of the head of household, region, family size, family type, a dummy if there is a person above 64 years old living in the household, the race of the head of household and the number of earners in the household. We include time fixed effects since we use quarterly data for the period of 2000-2023, which was subject to time-specific shocks that we want to condition on. Unfortunately, we cannot include household fixed effects since each household only reports income once, even if interviewed for four quarters. We weighted different observations according to their sampling weights and standard errors were clustered two-way at the quarter and income quintile level.

Results. The resulting estimates for β^i are presented in Figure 2 for each of the 22 consumption categories considered. Most goods present semi-elasticities between -1 and 1 percentage points. The response of food at home contrasts with that of the remaining categories, since, after a 1% increase in household income, the budget share allocated to food consumed at home decreases by, approximately, on average, 3 percentage points.

These estimates clearly distinguish two groups of categories: goods and services for which households, on average, *decrease* spending after an increase in income, which we will classify as *necessities*, and goods and services for which they either *do not change* or *increase* spending after a rise in income, which we will classify as *luxuries*. We therefore classify 7 out of the 22 categories as necessities: food at home, utilities such as electricity, telephone, water and natural gas, gasoline and tobacco; and the remaining 15 categories as luxury goods: fuel, other vehicles, health drugs, used vehicles, education, personal care, public transportation, alcohol, apparel, health services, other lodging, household operations, food away, entertainment and new vehicles.

Figure 2: Semi-elasticities of expenditure share to household income

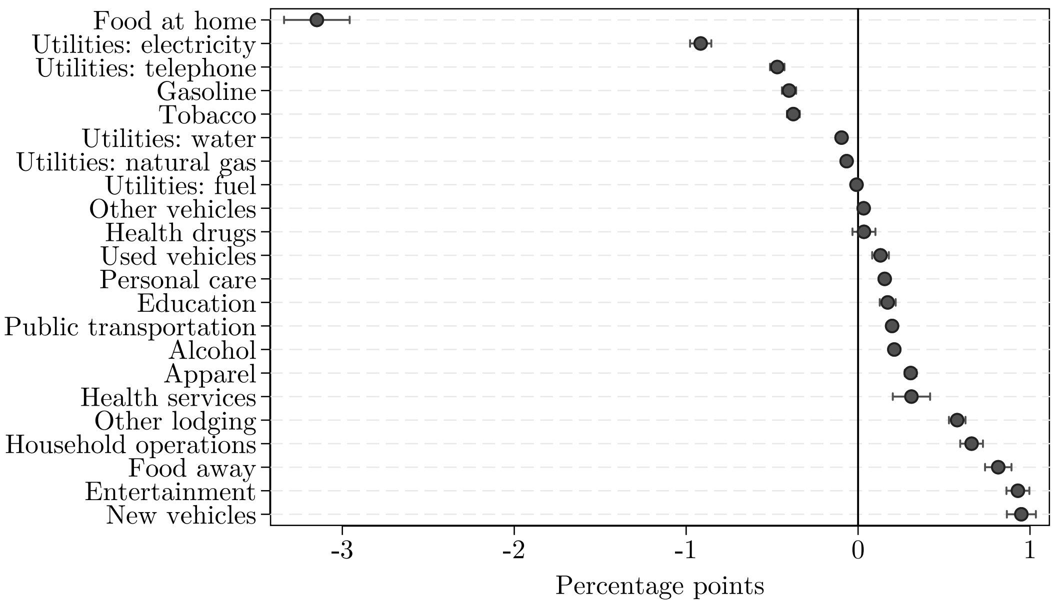

Notes: This figure shows the coefficient estimates $\widehat{\beta_1^i}$ from the model in Equation (1), estimated separately for each of the 22 consumption categories considered. The coefficient estimates can be interpreted as the percentage point change in the budget share associated with a 1% increase in pre-tax household income. Standard errors are clustered both at the level of the year-quarter and income quintile.

Alternatively, one could classify consumption goods using the slope of their Engel Curves. We show in Appendix A.1 that the two classifications yield similar results.

Using this classification, we then analyze how the budget share of necessities varies along the income distribution. Figure 3 presents the budget share of necessity goods along income quintiles. The expenditure share of necessities is a negative function of income, with the gap between the first and the last quintiles reaching approximately 19 percentage points. Note that this is not true by construction as we classified consumption goods using their expenditure elasticities, not their Engels curves. Averaging out across income groups, the average necessities' share corresponds to 49.1% and is represented by the dashed line. This average value is used as a calibration moment in the model, as described in Section 5.

Figure 3: Necessities' budget share along income

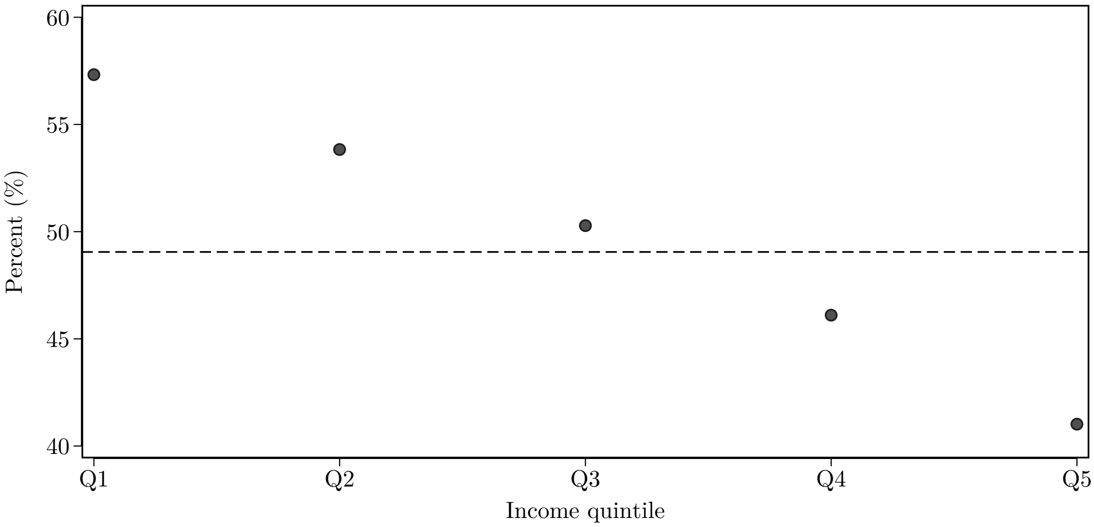

Notes: This figure presents the budget share allocated to necessity goods, in percent, along pre-tax household income quintiles. The dashed line represents the average share of necessities across this period and income, which is 49.1%. This is calculated using data between 2000 and 2023 and weighting observations differently according to sampling weights.

2.3 Response of budget composition to a relative price shock

Having established that choices on expenditure composition vary along income, we now document that household responses to a relative price shock are also income-dependent: higher-income households respond more that lower-income ones.

Methodology. We assess the impact on the composition of expenditure of a relative price shock, which increases the price of necessities relative to luxury goods, by estimating the following model:

$$\Delta\omega^n_{h,t} = \alpha + \beta\Delta\mathbf{P}_t + \epsilon_{h,t}, \tag{2}$$

where $\Delta\omega^n_{h,t}$ is change of the expenditure share on necessities (n) of household h in period t and $\Delta\mathbf{P}_t$ is the change in the relative price of necessities in period t, which here is standardized for interpretation purposes. We cluster standard errors at the household level and use the sample weights provided in the CEX to ensure representativeness of our estimates in the US population.

We conduct this analysis using the previously described dataset, where both the changes in budget shares and in relative prices are measured over consecutive three-month intervals. We use first-differences since we want to understand how the composition of expenditure changes after a change in the relative price, and not the level *per se*. Using the level of both variables would capture longer-term trends in these two variables, instead of how local changes affect households' decisions. Due to using first-differences, we are already differencing out any time-invariant information, thus we do not include household-specific information in this models such as household controls or fixed effects.

We construct the relative price index between necessity and luxury goods, **P**, using the same classification of necessity and luxury goods. First, we compute a price index for both groups of consumption goods separately. For that, we weight each category's CPI with the expenditure in that category over total expenditure in either necessities or luxuries. The relative price then corresponds to

$$\mathbf{P} = \frac{P_t^n}{P_t^\ell} = \frac{\sum_{i \in \mathcal{N}} p_t^i w_t^{i,n}}{\sum_{j \in \mathcal{L}} p_t^j w_t^{j,\ell}}, \tag{3}$$

where P_t^n is the price index for necessities in period t and P_t^ℓ is the price index for luxuries in period t; $\mathcal{N}$ is the set of necessity goods and $\mathcal{L}$ is the set of luxury goods; $w_t^{i,n}$ corresponds to the weight of category i within necessities in period t and $w_t^{j,\ell}$ to the weight of category j within luxuries in period t.

In order to causally estimate the effect price changes in necessity goods on expenditure composition, we instrument the relative price of necessities, **P**, with the oil supply news shock series estimated by Känzig (2021).[6] A valid instrument must satisfy two key conditions: exogeneity and relevance. The instrument is exogenous since, as described in Section 2.1.1, the shocks are estimated by exploiting the effects on oil future prices of OPEC announcements on oil supply on around a narrow window around the event. These are then fed into a VAR model to recover the structural oil supply news shock. The instrument is relevant because oil shocks are likely to have a bigger, immediate impact on the goods we classify as necessities, such as electricity and natural gas — a hypothesis which we test later on. Therefore, we expect that these shocks raise primarily the price of necessities over luxuries, creating a relative price shock for necessity goods.

Results. The relative price series constructed as in (3) is presented in Figure 4. The relative price of necessities increased steadily from 2000 until the beginning of the Great Financial Crisis (GCF) in 2008, which represented the biggest peak in the series so far during the period from 2000 until 2023. After a sharp fall, the relative price of necessities continued to increase, reaching another local maximum in the beginning of the sovereign debt crisis. The relative price decreased almost continuously from 2014 until the start of the pandemic, having reached 2000 values in late 2020. The start of 2021 saw a rise in the relative price of necessities relative to luxury goods, which translated into an inflationary

[6] As it stands, estimating the model in Equation (2) would simply capture any correlation between the budget share of necessities and the relative price of necessities, where the price response can and will likely be endogenous to any other shock that might change the composition of consumption. An example of such a shock would be an income shock which makes everyone poorer, thus leading to a higher demand for necessities, increasing their relative price.

episode, peaking in the second quarter of 2022.

Figure 4: Relative price between necessity and luxury goods

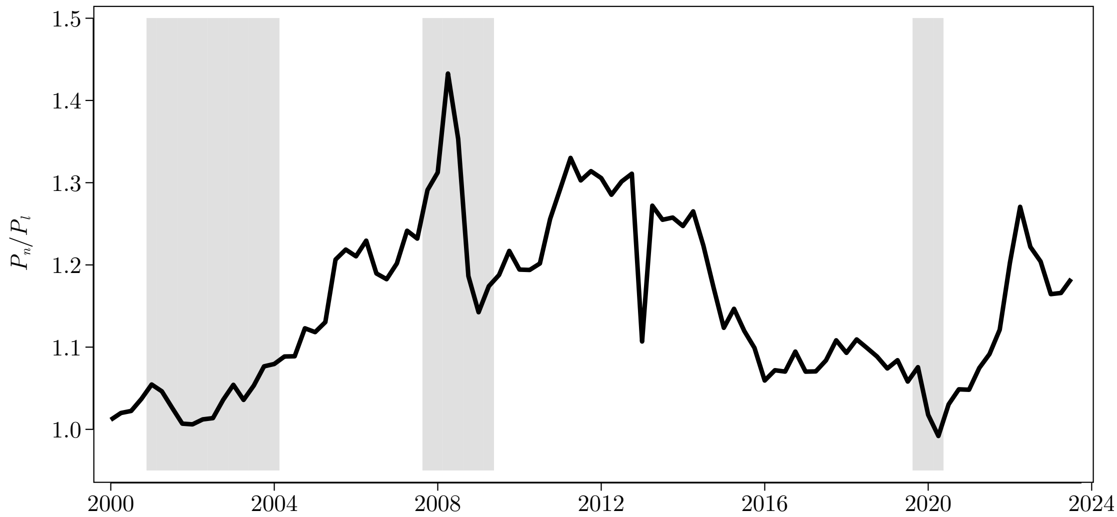

Notes: This figure presents the quarterly relative price series between necessity and luxury goods, computed according to Equation (3), using combined data from the CEX and CPI for 2000-2023. The relative price index is equal to 100 in the first quarter of 2000, since both price indices used to construct it are also based in that period. The gray areas correspond to US recessions according to the NBER's Business Cycle Dating Committee.

Before estimating the model in Equation (2) using an instrumental variable estimator, we formally test for the relevance of oil supply news shocks in driving the relative price between necessity and luxury goods. We assess for instrument relevance in two ways: by running a local projection of oil supply news shock on the relative price and, second, by formally testing this through the first-stage regression of the model in Equation (2). For the local projection, we regress the identified oil supply shock and the oil supply news shock on the relative price of necessity goods using the following local projection specification:

$$\Delta \mathbf{P}_{t+h,t-1} = \beta_0^h + \beta^h z_t + \sum_{l=1}^{4} \gamma_l^h \, \Delta \mathbf{P}_{t-l,t-l-1} + \zeta_t^h, \quad h \in [0, 12] \cap \mathbb{Z}, \tag{4}$$

where $\Delta \mathbf{P}_{t+h,t-1}$ is the change in the relative price of necessities between periods $t+h$ and $t-l$, z_t is the 3-month sum of the monthly oil supply news shock at time t and $\Delta \mathbf{P}_{t-l,t-l-1}$ are lagged values of the relative price. We consider a horizon of 12 quarters and 4 lags for the relative price. Since the shock is structurally identified, we do not consider any lags; however, we tested this considering 4 lags of the shock and the results are almost identical.

The resulting estimates are presented in Figure 5. The oil supply news shock has a statistically significant and immediate effect on the relative price of necessities across all relevant significance levels. Specifically, a one standard deviation shock increases the relative price by an average of 0.21 standard deviation units, which corresponds to a raw relative price change of approximately 0.02. In other words, the relative price responds by about one fifth the magnitude of the shock. The effect persists into the following period, albeit with a smaller magnitude: the lagged response is estimated at 0.17 standard deviation units. After that, the shock dissipates, which is expected due to the nature of the shock and given that we considering its effect on short-term changes in the relative price.

Figure 5: Impulse response of the relative price to an oil supply news shock

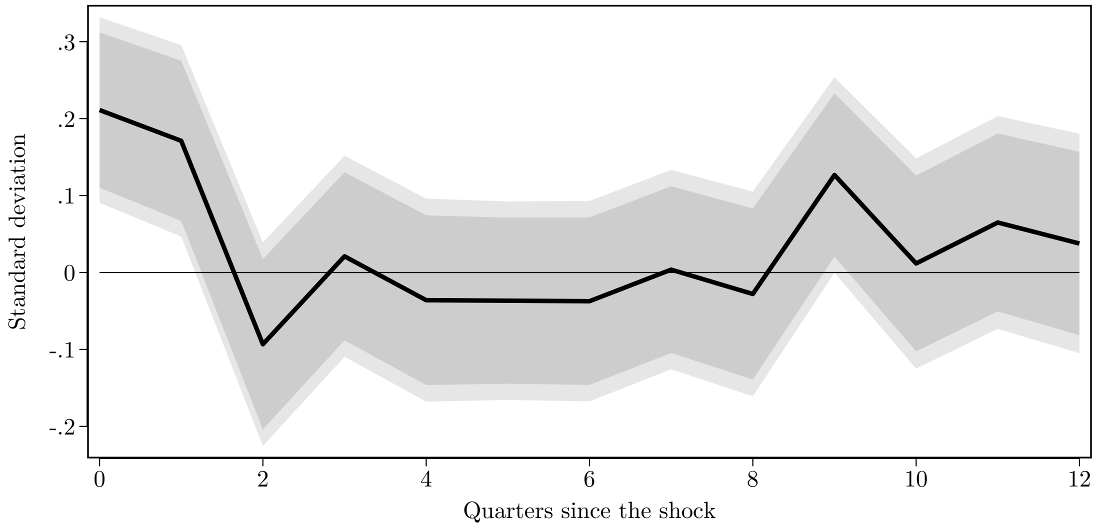

Notes: The figure presents the impulse response of the relative price to an oil supply news shock, estimated using the model in Equation (4), using combined data from the CEX, CPI and Känzig (2021)'s oil supply news shocks for 2000-2023. The dark grey and the light grey areas correspond to the 90% and the 95% confidence intervals, respectively.

We now formally test both the validity and strength of our instrument using an F-test in the first-stage regression. The F-statistic in Table 1 demonstrates the strength of our instrument.[7] As suggested by Andrews, Stock, and Sun (2019), we also compute the effective F-statistic of Montiel-Olea and Pflueger (2013), which is a more conservative procedure and has the advantage of adjusting for heteroskedasticity, serial correlation or clustering in the second stage, which is our case.[8] The effective F-statistic shows that the instrument surpasses the most demanding critical value for $\tau = 5\%$, therefore allowing us to reject that the instrument is weak.

We proceed to estimate the model in Equation (2), instrumenting the relative price

[7]To ensure that we are not in the presence of a weak instrument, when using an external instrument, the F-statistic should be above the threshold value of 10 (Andrews, Stock, & Sun, 2019). We only have a single endogenous regressor, thus the Kleibergen-Paap (KP) F-statistic is robust as it does not rely on the on the assumption that the errors are independent and identically distributed (i.i.d.). Table 1 reports that the KP F-statistic is 19.280, surpassing the standard threshold value.

[8]See Montiel-Olea and Pflueger (2013) for test-specific critical values.

Table 1: First stage regression results

	Oil supply news shock
Coefficient	0.168***
	(0.038)
Clusters	97,126
Observations	279,136
Kleibergen-Paap F-stat	19.280
Effective F-stat	80,514.253
T2LS/LIML CV τ=5%	37.418

Notes: *** corresponds to the 99% significance level. The table shows the results of the first stage regression of the change in the relative price of necessities, $\Delta\mathbf{P}$, on the oil supply news shock, z. The effective F-statistic corresponds to the Montiel-Olea and Pflueger (2013) robust weak instrument test. For the latter, we compute the critical values for the two-stage least squares (TSLS) and the limited information maximum likelihood (LIML) estimators, which, for $\tau = 5\%$, coincide. Standard error in parentheses.

by the oil supply news shock. We estimate the regression for the entire sample of households, for the bottom and top 50% of the distribution of income, and, finally, by income quartiles. We create these income groups according to the distribution in each year-quarter, taking into account sample weights.

The resulting estimates $\widehat{\beta}$ are presented in Table 2. The column "Total" shows the results for the entire sample of 97,126 households. The point estimate implies that a one-standard deviation increase in the relative price of necessities leads to an increase in the budget share of necessities by 1.49 percentage points. This means that, on average, households respond to a relative price increase by increasing their share of expenditures allocated to necessity goods. Columns Q1-Q4 report the results by income quartile and show that the response increases in income.

The main result emerges when comparing the bottom and top 50% of the sample. After a one-standard deviation increase in the relative price of necessities, the bottom 50% increase their necessities' expenditure share by around half the magnitude observed for the top 50%. We will discuss this finding in the context of the quantitative model developed in Section 5 as we use both the average estimate and the bottom and top 50% estimates for calibration purposes.

The same pattern appears when examining quartiles, though the larger standard errors prevent us from rejecting that neighboring quartiles are significantly different.[9] Quartile

[9] Limited data (maximum three observations per household after first-differencing) reduces statistical power for quartile comparisons, though a clear positive relationship with income remains. Larger groups like the bottom and top 50% are statistically different at the 90% confidence level.

1 shows no statistically significant response, consistent with consumption-constrained households already at subsistence levels who cannot adjust their spending composition. The remaining quartiles increase their necessity expenditure shares, with the response increasing with income. A one-standard deviation price increase raises necessities' share by 1.2 percentage points for quartile 2 and 2.6 percentage points for quartile 4 — roughly double the effect. This pattern can reflect two mechanisms: higher-income households can better absorb price increases while maintaining consumption levels, and strong complementarities between necessities and luxuries (e.g., larger homes requiring more electricity) compel wealthier households to maintain necessity consumption. Poorer individuals, on the other hand, may have fewer or weaker complementarities. These findings contradict Stone-Geary preferences, which predict weaker responses at higher incomes due to their asymptotically homothetic nature far enough from the subsistence level.

Table 2: Change in the necessities' share after a relative price change

	Total	B50	T50	Q1	Q2	Q3	Q4
$\Delta\mathcal{P}$	1.493***	0.970**	1.987***	0.753	1.207**	1.541**	2.558***
	(0.397)	(0.444)	(0.476)	(0.593)	(0.534)	(0.619)	(0.542)
Households	97,126	53,215	60,208	28,880	34,353	36,235	34,165
Observations	279,136	128,918	150,218	61,798	67,120	71,006	78,825

Notes: ***, ** and * correspond to 99%, 95% and 90% significance levels. The coefficient estimates presented correspond to the $\widehat{\beta}$ from Equation (2), with the relative price shock instrumented by the oil supply news shock, for different income groups: total, bottom/top 50%, and quartiles. The model was estimated using an IV estimator on combined data from the CEX, CPI and Känzig (2021)'s oil supply news shocks for 2000-2023. Standard errors in parentheses, clustered at the household level.

2.4 The pass-through of price subsidies

Value-added tax cuts emerged as a primary fiscal instrument to mitigate the welfare costs of inflationary shocks by directly reducing prices. Critics argued that such subsidies were both distortionary and potentially ineffective, since tax reductions might not translate into lower consumer prices. We provide evidence supporting the latter concern: the pass-through of price subsidies is, on average, imperfect.

Methodology. We estimate the pass-through semi-elasticity of prices to tax changes. We use a high-dimensional fixed effects estimator for the following specification:

$$\Delta \log(p_{ict}) = \alpha + \sum_{j=-6}^{6} \beta_j \, \Delta\tau_{ic,t+j} + \zeta_c + \eta_i + \gamma_t + \epsilon_{ict}, \tag{5}$$

where $\Delta \log(p_{ict})$ is the change in the log price of category i in country c in month t, $\Delta\tau_{ic,t+j}$ is the change in the tax rate of category i in country c in month $t+j$, with $j \in [-6, 6]$, ζ_c are country fixed effects, η_i are category fixed effects and γ_t are time fixed effects. We use two-way clustering for the standard errors at the country-category level.

Following Benzarti et al. (2020), we include six lags and leads of the change in taxes, thereby allowing for possible anticipatory or delayed effects after the reform.

Results. Figure 6 presents the results for the model in Equation (5) for $\sum_{j=-6}^{6} \beta_j$. Similarly to Benzarti et al. (2020), we find an average pass-through of 21.3% on impact. The following two months after implementation still display a statistically significant effect, whereas the fourth month after implementation sees a decrease in the transmission to consumer prices, thus cumulating to a pass-through of 24.2% four months after the reform. We find no evidence of pre-trends nor anticipatory behavior.

We incorporate this finding in our model since it is crucial for determining the efficacy of price subsidies. If retailers capture most of the benefit from the tax cut, the possible resulting increase in profits flows to higher-income households through dividends.

Figure 6: Pass-through of value-added tax changes

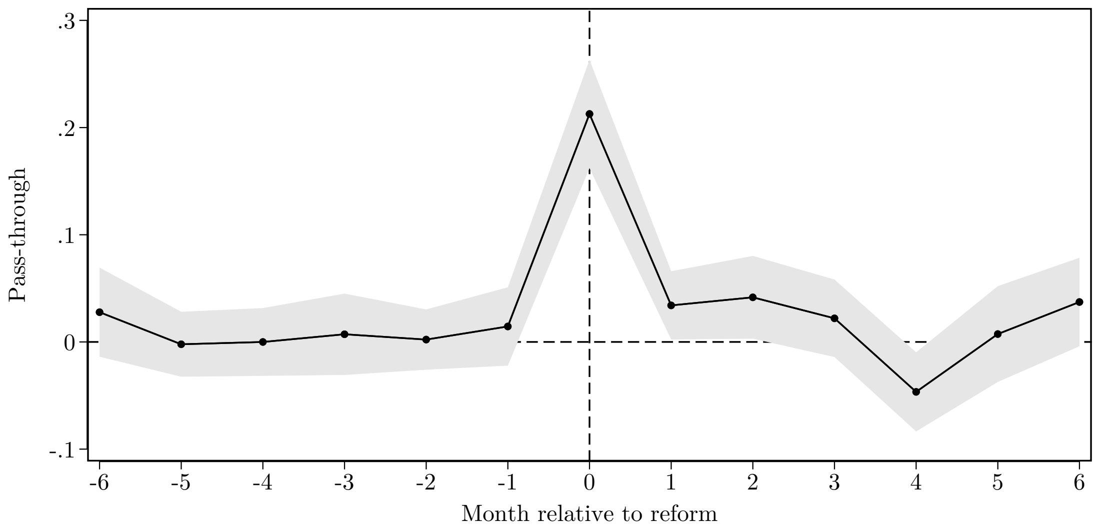

Notes: The figure shows the coefficient estimates $\sum_{j=-6}^{6} \hat{\beta}_j$ from the model in Equation (5) using combined data for COICOP5-level prices and VAT changes from Benzarti et al. (2020), later extended by Benzarti and Tazhitdinova (2021). The grey area corresponds to the 95% confidence interval.

3 Stylized Model

In this section, we consider a static partial-equilibrium model that captures the two empirical facts documented in the previous section. On the household side, we assume a nonhomothetic preference relation over consumption goods as in Comin, Lashkari, and Mestieri (2021), which can capture the differences in budget shares between income groups. Firms produce under monopolistic competition with a decreasing returns to scale technology, which allows us to match the incomplete pass-through of consumption taxes documented above. With this simple framework, we show the relevance of nonhomothetic preferences and decreasing returns to scale for consumption and welfare effects.

In the next section, we extend this model to a dynamic version with heterogeneous agents to quantify the effects of recessions and study alternative fiscal policies.

3.1 Households

Households have preferences over consumption over two goods: necessities (n) and luxuries (ℓ). These are implicitly defined through a nonhomothetic Constant Elasticity of Substitution aggregator as in Comin, Lashkari, and Mestieri (2021):

$$\left(\Omega_n c^{\epsilon_n}\right)^{\frac{1}{\sigma}} c_n^{\frac{\sigma-1}{\sigma}} + \left(\Omega_\ell c^{\epsilon_\ell}\right)^{\frac{1}{\sigma}} c_\ell^{\frac{\sigma-1}{\sigma}} = 1, \tag{6}$$

where c is the aggregate consumption index, c_s is the consumption of sector s for $s \in \{n, \ell\}$, $\sigma \in (0, 1)$ is the price elasticity of substitution, $\epsilon_s > 0$ is the income elasticity for good s, and $\Omega_s > 0$ is a sector-specific taste shifter. A uniform scaling of $(\epsilon_n, \epsilon_\ell)$ or (Ω_n, Ω_ℓ) implies the same observable choice behavior, as only the ratio between the income elasticities and the CES weights matters. Therefore, we set good ℓ as the base good, and thus ϵ_ℓ and Ω_ℓ are equal to 1.

Households receive an endowment e that they can spend on either c_n or c_ℓ given prices p_n and p_ℓ. Households choose c_n and c_ℓ by solving an expenditure minimization problem:

$$\begin{aligned} &\min_{c_n, c_\ell} E = p_n c_n + p_\ell c_\ell \\ &\text{s.to} \quad \left(\Omega_n c^{\epsilon_n}\right)^{\frac{1}{\sigma}} c_n^{\frac{\sigma-1}{\sigma}} + (c)^{\frac{1}{\sigma}} c_\ell^{\frac{\sigma-1}{\sigma}} = 1. \end{aligned} \tag{7}$$

The solution of the problem yields the sectoral Hicksian demands given by

$$c_n = \Omega_n \left(\frac{p_n}{E}\right)^{-\sigma} c^{\epsilon_n}, \tag{8}$$

$$c_\ell = \left(\frac{p_\ell}{E}\right)^{-\sigma} c, \tag{9}$$

where the output of sector s is a function of the sectoral price, p_s, expenditure, E, and the consumption index, c. As there is only one period, there will be no savings, and hence expenditure is equal to the total endowment.

Note that, for $\epsilon_n = 1$, the nonhomothetic CES in Equation (6) collapses to the regular

CES aggregator, recovering the homothetic case. For $\epsilon_n < 1$, the expenditure share that a household chooses in the necessity sector is a decreasing function of income, and the smaller the income elasticity, the more convex the Engel curve is.[10]

3.2 Firms

There are two final goods sectors that produce necessities and luxuries. Each of them aggregates a continuum of intermediate inputs indexed by $j \in [0, 1]$ according to

$$Q_s = \left(\int_0^1 q_s(j)^{\frac{\theta_s - 1}{\theta_s}} dj \right)^{\frac{\theta_s}{\theta_s - 1}}, \tag{10}$$

where θ_s is the elasticity of substitution across intermediate inputs of sector s. Given a level of sectoral demand Q_s, the final producer's cost minimization problem implies the demand for intermediate inputs given by

$$q_s(j) = \left(\frac{p_s(j)}{P_s} \right)^{-\theta_s} Q_s, \tag{11}$$

where $p_s(j)$ is the price of intermediate input j and P_s is the equilibrium price of the final good, expressed as

$$P_s^{1-\theta_s} = \int_0^1 p_s(j)^{1-\theta_s} dj. \tag{12}$$

Intermediate firms of sector s produce under monopolistic competition using labor as their only input in a decreasing-returns-to-scale production technology:

$$q_s(j) = Z_s n_s(j)^{\alpha_s}, \tag{13}$$

where $n_s(j)$ is the labor input and Z_s is the parameter that regulates the level of technology, common to all firms in the sector. These firms also face a fixed cost, F. They choose prices to maximize profits according to:

$$\max_{\{p_s(j)\}} \ \pi_s(j) = p_s(j) q_s(j) - w n_s(j) - F \tag{14}$$
$$\text{s.to Equations (11) and (13).}$$

[10]This is true under the assumption that $\sigma < 1$. For $\sigma > 1$, $\epsilon_n > 1$ would indicate that the good produced in sector n is a necessity good.

The optimal price-setting rule is given by

$$p_s(j) = \left[\frac{\theta_s}{\theta_s - 1}\frac{w}{\alpha_s Z_s^{1/\alpha_s}}\left(Q_s P_s^{\theta_s}\right)^{\frac{1-\alpha_s}{\alpha_s}}\right]^{\frac{\alpha_s}{\alpha_s+\theta_s-\alpha_s\theta_s}}. \tag{15}$$

3.3 Results

We use the stylized model to illustrate two key mechanisms of the model. First, we show that the decreasing-returns-to-scale technology allows us to capture the empirically observed incomplete pass-through. Second, we show that our preference structure allows us to match the nonhomothetic consumption behavior documented in the previous section.

3.3.1 Incomplete Pass-Through

When the production function features decreasing returns to scale, the pass-through is no longer perfect. To see this, note that the first-order condition associated with (14) implies

$$0 = p(j) + p(j)\frac{\partial q(j)}{\partial p(j)} - C'(q)\frac{\partial q(j)}{\partial p(j)}.$$

Under constant returns to scale (CRS), marginal costs are constant, and the first-order condition collapses to the well-known result of constant markups, where the price is equal to the markup times the marginal cost. This means that the price set by the firms does not depend on the consumer's demand. Under decreasing returns to scale (DRS), the marginal cost is not constant anymore and increases with the scale: a larger scale is less efficient.

The introduction of a price subsidy lowers the price consumers face and hence increases the demand for consumption goods. As DRS implies that marginal costs are not constant, this will change the price that firms set and, in turn, the final price that consumers face. Figure 7 plots the pass-through against the returns to scale parameter α. The stronger the returns to scale, the higher is the pass-through of the price subsidy. When returns to scale are constant, the pass-through is perfect. This illustrates how we can use the returns-to-scale parameter α to match the empirically observed pass-through of VAT.[11]

[11] We will employ sector-specific values of α, not because we necessarily believe that differences in returns to scale are the fundamental driver that drives cross-sector differences in returns to scale. Rather, differences in the returns to scale are an easily implementable approach that generates the observed differences in the pass-through rates.

3.3.2 Consumption behavior across the income distribution

Next, we illustrate how the household's consumption decision depends on the degree of nonhomotheticity, as captured by the preference parameter ϵ_n. We compare our nonhomothetic preferences ($\epsilon_n = 0.6$) to the homothetic case with $\epsilon_n = 1$. We also compare our results to those stemming from Stone-Geary preferences (Geary, 1950; Stone, 1954), which the majority of the literature has relied on so far.[12] Finally, we add the necessities' share estimated from the data as presented in Figure 3 for different income quintiles.

The left panel in Figure 8 displays necessity budget shares as a function of their income endowment. Empirically, we find that the budget share allocated to necessities corresponds to just below 60% for the lowest income quintile to 40% for the highest.[13] When comparing these to the predictions of the various preferences listed above, we find that, naturally, homothetic preferences exhibit constant budget shares and fail to match the pattern. Stone-Geary preferences are nonhomothetic preferences and, as such, generate a fall in budget shares; however, most of that fall happens between the first and second income quintile.[14] Finally, we display in orange the expenditure shares stemming from our nonhomothetic CES, which replicates best what is empirically observed.

These differences matter for our quantitative framework: the Stone-Geary preferences would not be able to match the differences in expenditure shares across the income distribution, which is a key determinant in the extent to which price subsidies lead to welfare gains and losses across the income distribution.

Next, we estimate the welfare losses — in the partial equilibrium sense — induced by the introduction of a consumption tax on necessities of 5%. To visualize these losses, the right panel of Figure 8 plots the income compensation required after a 40% increase in the relative price of necessities. Since the environment features multiple consumption goods, we refrain from employing the consumption equivalent variation. Instead, we compute for each household the income compensation they are willing to forego to return from P' to P, as a percentage of the income endowment.

While the compensation required by households along the income endowment distribution is constant in the case of homothetic CES, it declines in income with nonhomothetic CES. This follows from lower-income households being more affected by shocks to necessity goods than higher-income ones. Thus, the nonhomothetic structure of preferences has distributional welfare effects.

[12] Stone-Geary preferences are given by $U = \log(c_n - \bar{c}) + \log(c_\ell)$, and are not defined for $c < \bar{c}$. We use $\bar{c} = 0.1$ for the numerical illustration.

[13] Our empirical numbers stem from the results in Figure 3.

[14] This is intuitive since the subsistence point is quantitatively irrelevant once consumption is sufficiently large, with one of the major caveats of this class of preferences being that they quickly become asymptotically homothetic.

Note that the share of foregone income is strictly smaller than the household's exposure to the necessity good since they partially substitute non-essential consumption in response to the price shock. These elasticities – and the extent to which they vary along the income distribution – are key in estimating the welfare loss from price shocks.

Figure 7: Pass-through of a sectoral price subsidy

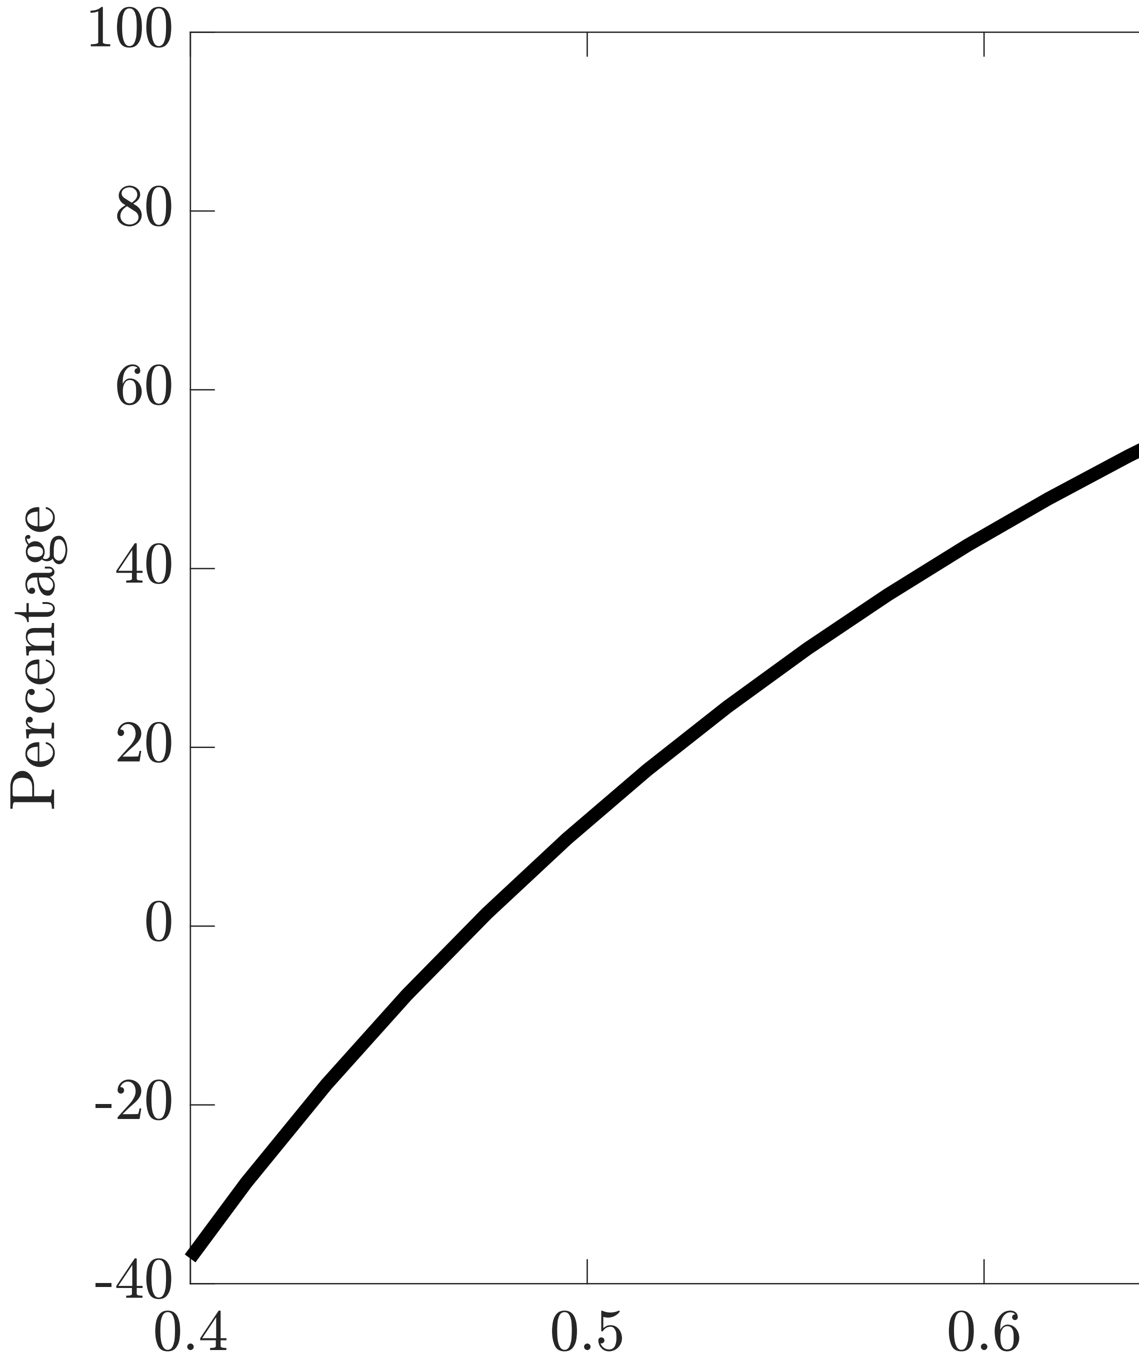

Notes: The figure shows the pass-through of a change in the consumption tax of necessities to the price of necessities for differe

Figure 8: Nonhomotheticity, consumption decisions and welfare

(a) Share of necessities **(b)** Welfare compensation

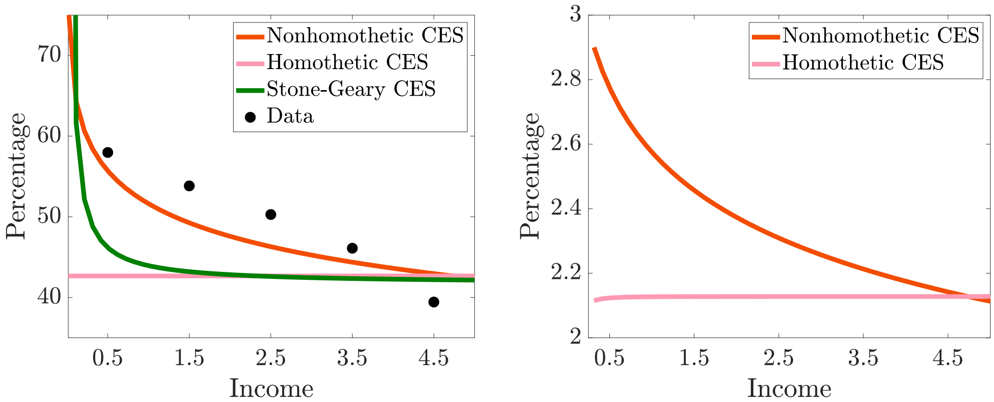

Notes: The left panel of this figure shows the budget share in the necessity sector as a function of the income endowment for nonhomethetic CES ($\epsilon_n = 0.6$), homothetic CES ($\epsilon_n = 1$), Stone-Geary preferences and the estimates from the data presented in Figure 3. The right panel of this figure corresponds to the welfare compensation required in income for households to be indifferent before and after a 5% increase in the consumption tax of necessities.

4 Quantitative Model

In this Section, we extend the model of the previous section to a dynamic setting with incomplete markets and heterogeneous agents as in Huggett (1993) or Aiyagari (1994). With this model, we will first study the aggregate and heterogeneous effects of a negative sectoral supply shock in either the necessity or the luxury goods sectors. We then study the welfare gains (and losses) associated with different fiscal policy interventions that attempt to mitigate the effects of said shocks.

The model is set up as follows. Workers receive wages in return for their production in firms in the necessity and luxury goods sector. Firms potentially make profits that are paid out as dividends to the workers. Workers also trade assets, which exist with a fixed positive supply. Finally, a government exists that is inactive in our baseline analysis. It becomes relevant during our policy analysis, in which it engages in fiscal policy, financed with taxes on labor.

We now discuss these three agents in detail.

4.1 Households and Firms

The economy consists of a continuum of agents normalized to measure 1 with CRRA preferences over consumption and additively separable preferences for leisure:

$$U = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, h_t), \tag{16}$$

where the per-period utility is

$$u(c,h) = \frac{c^{1-\gamma} - 1}{1-\gamma} - \chi \frac{h^{1+\psi}}{1+\psi}, \tag{17}$$

and $\beta \in (0,1)$ is the discount factor. The variable c indicates the consumption index, defined implicitly through the non-homothetic CES aggregator, as defined in the previous section in Equation (6).

Households are ex-ante heterogeneous in terms of their ability a. They are ex-post heterogeneous with respect to bond holdings, b, and idiosyncratic productivity, z. They rent their effective labor services, az_th_t, — productivity weighted hours — to firms for a real wage w_t. Households buy and sell bonds b_t at p_t^b.

Besides paying wages to households, firms also generate profits as they operate in monopolistic competition. Since households own those firms, they receive the generated profits in the form of dividend income, D_t. A household's dividend income depends on the household's ability.[15]

There is a government that collects taxes on consumption, τ_s^c, where $s \in \{n, \ell\}$, and labor income, τ^h, and gives subsidies, B, potentially only to some households. The household budget constraint reads as

$$\sum_{s\in\{n,\ell\}} (1+\tau_{s,t}^c) p_{s,t} c_{s,t} + p_t^b b_{t+1} = a z_t h_t w_t (1-\tau_t^h) + p_t^b b_t + D_t + T_t, \tag{18}$$

where $\log(z_{t+1}) = \rho_z \log(z_t) + \varepsilon_z$, with $\varepsilon_z \sim \mathcal{N}(0, \sigma_z^2)$. Households take prices as given. We can rewrite the household's problem recursively as follows:

$$\begin{aligned} V(a,b,z;\Xi) = & \max_{\{c_n,c_\ell,b',h\}} \quad u(c,h) + \beta \mathbb{E}\left[V(a',b',z;\Xi')\right] \\ \text{s.to} \quad & \sum_{s\in\{n,\ell\}} (1+\tau_s^c) p_s c_s + p^b b' \leqslant azhw(1-\tau^h) + p^b b + D + T \\ & 1 = (\Omega_n c^{\epsilon_n})^{\frac{1}{\sigma}} c_n^{\frac{\sigma-1}{\sigma}} + (c)^{\frac{1}{\sigma}} c_\ell^{\frac{\sigma-1}{\sigma}} \\ & \Xi' = \Psi(\Xi) \\ & b' \geqslant \bar{b}, \; c_n, c_\ell \geqslant 0, \; h \in (0,1), \end{aligned} \tag{19}$$

where $\Xi(a,b,z) \in \mathcal{M}$ is the distribution on the space of ability, asset holdings, and labor productivity across the population, which, together with the policy variables, determines equilibrium prices. Ψ is an equilibrium object that specifies the evolution of the distribution Ξ, and $\mathcal{M}$ is the set of probability measures.

The firm problem is the same as in the static model of Section 3.

4.2 Government

There is a government that obtains revenue from taxing labor and consumption, and uses those proceeds to finance transfers to households. The budget constraint faced by the government is given by

$$\int azhw\tau^h \, d\Xi + \sum_s \int \tau_s^c p_s c_s \, d\Xi = \int T \, d\Xi, \tag{20}$$

We examine three distinct fiscal policies to address the economic effects of sector-specific recessions. Our analysis compares a sectoral consumption tax cut in the affected sector, broad-based lump-sum household transfers, and means-tested stimulus payments restricted to households with incomes below a specified threshold.

[15] We thus ensure that high-ability workers (with on average higher income) receive a higher share of dividend income than workers with lower income.

4.3 Market Clearing and Competitive Equilibrium

Market clearing requires that labor demand by firms is equal to aggregate labor supplied by households, that the demand for bonds by households is equal to the existing outstanding bond, and that households consume the goods produced by firms in each sector:

$$\int \hat{h}_t(a,b,z)\, d\Xi_t = N_{n,t} + N_{l,t} \tag{21}$$

$$\int \hat{b}_t(a,b,z)\, d\Xi_t = \bar{B} \tag{22}$$

$$\int \hat{c}_{s,t}(a,b,z) d\Xi_t = Q_{s,t},\ s \in \{n,\ell\}, \tag{23}$$

where $\hat{h}_t(a,b,z), \hat{b}_t(a,b,z)$ and $\hat{c}_{s,t}(a,b,z)$ are the household policy functions of labor, savings, and sectoral consumption, and $\bar{B}$ is the net asset supply in the economy.

Definition: A competitive equilibrium is a sequence of tax rates τ^h_t, $\tau^c_{n,t}$ and $\tau^c_{\ell,t}$; government transfers, B_t; value functions V_t with policy functions $\hat{c}_{n,t}, \hat{c}_{\ell,t}, \hat{h}_t$ and $\hat{b}_t$; prices p^b_t, p^n_t, p^ℓ_t, and w_t; profits $\pi_{n,t}$ and $\pi_{\ell,t}$; and a law of motion Ψ, such that:

1. V_t satisfies the Bellman equation, Equation (19), with corresponding policy functions $\hat{c}_{n,t}, \hat{c}_{\ell,t}, \hat{h}_t$ and $\hat{b}_t$ given sequences of tax rates, government transfers, prices, and profits.
2. Firms maximize profits as in Equation (14).
3. The government runs a balanced budget as in Equation (20).
4. For all Ξ_t, the market clearing conditions (21) – (23) are satisfied.
5. Aggregate law of motion Ψ generated by the savings policy function.

5 Calibration

In this section, we discuss how to calibrate our quantitative model. The unit of time is one quarter.

Preferences. We set the discount factor, β, to 0.99, implying a natural annual interest rate of 4%. The risk-aversion parameter, γ, is set to 2.2 and the Frisch elasticity is set to 1, following micro-level estimates from Havránek (2015) and Chetty et al. (2011), respectively.

The intratemporal preferences given by Equation (6) depend on five parameters. As explained in Section 3.1, we can normalize either ϵ_n or ϵ_ℓ: only the relative difference between ϵ_n and ϵ_ℓ — and between Ω_n and Ω_ℓ — matter for the degree of nonhomotheticity and thus affect household choices. We set the luxury goods sector as the base sector, and set the associated parameters to 1. This leaves us with ϵ_n, Ω_n, and σ, together with the χ which regulates the disutility from working.

We discipline these parameters as follows. The preference parameters ϵ_n, Ω_n, and σ primarily jointly affect (i) the average expenditure share in necessity goods in the economy, (ii) the expenditure share in necessity goods of the bottom 20% of the income distribution, and (iii) the average price elasticity of necessity goods' demand. The disutility of working primarily affects our fourth moment, the number of hours worked. [16]

We obtain our consumption-related moments using the dataset described in Subsection 2.1.1. The average share of expenditure in necessity goods and the expenditure share in necessities of the bottom 20% are computed using the CEX dataset, described in Section 2. The average price elasticity of necessity goods demand corresponds to the one presented in Section 2. Finally, the average number of hours is the OECD average, normalized as a percentage of the total number of hours available in a year.

Education groups. We consider four education groups: (i) less than high school, (ii) high school, (iii) some higher education, (iv) BSc degree or more. These four groups differ in their permanent ability level, a, and the share of dividend income received out of the firm's profits.

We use the 2022 wave of the Survey of Consumer Finances (SCF) to compute the necessary moments by education group. First, we compute the share of dividend income accruing to each education group. Second, we compute the total income — the sum of labor and capital income — by education group.[17] We then set the parameters ASDF such that each education group receives dividend income in line with the empirical counterparts. Finally, we calibrate the permanent productivity of each education group a_i such that their total income matches the empirical counterpart. These empirical moments are summarized in Table 3.

Idiosyncratic Productivity Process. Households also face an idiosyncratic productivity shock, common to all ability types, that follows a log-AR(1) process as described in

[16] Since all parameters affect all four moments to some extent, we estimate them jointly in a Simulated Method of Moments (SMM) procedure.

[17] Capital income corresponds to the sum of the net annual income from interest, dividends, a sole proprietorship or a farm, mutual funds, the sale of stocks, bonds or real estate, non-taxable investments, and other businesses or investments, net rent, trusts, and royalties.

Table 3: Income and dividend distribution by education group

Education group	Average labor income	Average total income	Share of households	Share of dividend income
1	\$ 48,505	\$ 59,661	30.10%	3.75%
2	\$ 59,630	\$ 74,671	20.97%	4.94%
3	\$ 108,806	\$ 160,646	31.52%	37.56%
4	\$ 177,234	\$ 284,801	17.41%	53.75%

Notes: Computed using data from the 2022 SCF. We only include households with positive total income, i.e., the sum of labor income, unemployment benefits, and capital income. The calculation of the education groups and all other figures take into account the replicate survey weights. The education groups correspond to the following: (1) less than high school, (2) high school, (3) some higher education, (4) BSc degree or more.

Section 4. The values of the persistence and variance of the innovations are set according to the estimates by Krueger, Mitman, and Perri (2016).[18]

Firms. Production is characterized by two sectors, each producing a final good that is a combination of intermediate inputs that operate under monopolistic competition with an elasticity of substitution given by θ_n and θ_ℓ, respectively. Each intermediate input is produced under decreasing-returns-to-scale with curvature α and a sector-specific technology, Z_n and Z_ℓ.

We normalize Z_ℓ, the TFP parameter of the luxury goods sector, to 1. The other four parameters are calibrated using the SMM. We match four additional moments: (i) the price pass-through of a consumption tax change in the necessity goods sector, (ii) the price pass-through of a consumption tax change in the luxury goods sector[19], (iii) the relative price between the two sectors, and (iv) the earnings-to-profit ratio in the economy. In Panel III of Tables 4 and 5, we show the value of the parameters and the target moments fit associated with the production side of the economy. The average relative price between the two sectors is computed using the expenditure-price data from Section 2.

Summary. Table 4 lists the values of all the parameters. Table 5 shows the model calibration fit with the data and model moments used and found in the simulated method of moments.

[18]We then discretize the shock into a five-state Markov chain using the Tauchen (1986) method.

[19]We use preliminary estimates for the sector-specific pass-through based on the data described in Subsection 2.1.2.

Table 4: Model Parameters

Parameter	Name	Value	Observations
		I. Preferences	
β	Discount factor	0.99	Standard
ε_n	Necessity goods: homotheticity	0.57	Internal calibration
Ω_n	Necessity goods: taste-shifter	0.48	Internal calibration
σ	CES parameter	0.23	Internal calibration
γ	CRRA coefficient	2.20	Havránek (2015)
η	Frisch elasticity	1	Chetty et al. (2011)
χ	Disutility from working	214.75	Internal calibration
		II. Ability and Productivity	
ρ_z	Idiosync. productivity persis.	0.99	Krueger, Mitman, and Perri (2016)
σ_z	Idiosync. productivity std. dev.	0.1	Krueger, Mitman, and Perri (2016)
$[a_1, a_2, a_3, a_4]$	Ability grid	[0.7, 1, 3.5, 9.3]	Internal Calibration
		III. Production	
α_n	Curvature production function	0.22	Internal calibration
α_l	Curvature production function	0.29	Internal calibration
θ_n	EIS of necessity	6	Christiano, Eichenbaum, and Rebelo (2011)
θ_ℓ	EIS of luxury	6	Christiano, Eichenbaum, and Rebelo (2011)
Z_n	SS necessity sector TFP	0.57	Internal calibration
Z_l	SS luxury sector TFP	1	Normalized
F	Fixed cost	0.1	Internal calibration

Notes: The table displays the values of the parameters. Panel I contains the values of the preference parameters, Panel II the values of the ability factors and the idiosyncratic productivity process, and Panel III the values of the parameters of the production side of the economy. The right-most column has the source, in case the values are set directly.

Table 5: Calibration fit of the SMM

Moment	Data	Model
I. Preferences		
Average cons. share in necessities	49.1	49.1
Ratio necessities share bottom 20% to mean	1.18	1.15
Price elasticity of necessity goods	1.49	1.48
Average hours worked	0.19	0.19
II. Ability and Productivity		
Relative income of ability types	[0.81, 1, 1.82, 2.97]	[0.81, 1, 1.83, 2.97]
III. Production		
Pass-through of necessity goods sector	0.29	0.29
Pass-through of luxury goods sector	0.17	0.17
Relative price of necessity to luxury goods	1.15	1.15
Earnings to profit	3.73	3.72

Notes: The table displays the values of the calibration targets in the model and the data. Panel I contains the values of the preference parameters, Panel II the values of the ability factors and the idiosyncratic productivity process, and Panel III the values of the parameters of the production side of the economy.

6 Sectoral Supply Shocks

We now study the response of our economy to negative sectoral supply shocks in either the necessity good sector or the luxury good sector. We model the sectoral recessions as a one-time negative shock in the sectoral productivity parameter, Z_s, $s \in \{n, \ell\}$, of 10%. This shock is an unexpected, never-again-to-occur departure from the initial steady state value. After realization, sectoral productivity slowly returns to its steady-state value following an AR(1) process with persistence parameter of 0.8. Once realized, the entire forward path of productivity is known to the agents.

6.1 Aggregate and cross-sectional effects

In Figure 9 we plot the response of the economy to a negative supply shock in the necessity goods sector and the luxury goods sector. The first two panels (Figures 9a and 9b) plot the aggregate impulse responses, whereas the latter two panels (Figures 10c and 9d) plot the cross-sectional response in the first period following the shock.

In the aggregate, consumption falls more during luxury than during necessity recessions. This is driven by the consumption response of high-income workers whose consumption falls more in luxury recessions than in necessity recessions. Contrarily, low-income households adjust their consumption more in necessity than in luxury recessions. Their response is stronger of than that of high-income households, but they contribute less to aggregate consumption.

The next two subplots show the heterogeneity across consumption goods: during the luxury recession, consumption of either good falls proportionally across the income distribution. During the necessity recession, the fall in necessity consumption proportionately decreases in income: high-income households adjust less of either category, whereas low-income households adjust proportionately more of both. This is because of the constant elasticity of substitution across the two goods: low-income households cannot maintain their consumption standards and their consumption falls, disproportionately so for necessities. Their overall consumption share of necessities remains larger than that of luxuries, but they substitute away from necessities to (relatively cheaper) luxuries. This is also consistent with the fact that the expenditure share of luxuries falls for everyone during the necessity recession, even though households are substituting towards luxury goods. This pattern of expenditure shares aligns with our empirical findings in Table 2, despite being untargeted.

Aggregate expenditures increase more during necessity than luxury recessions. During the necessity recession, the increase in expenditure decreases in household income, whereas during the luxury recession, the increase in expenditure increases in income.

Figure 9: Impulse responses after negative sectoral productivity shocks

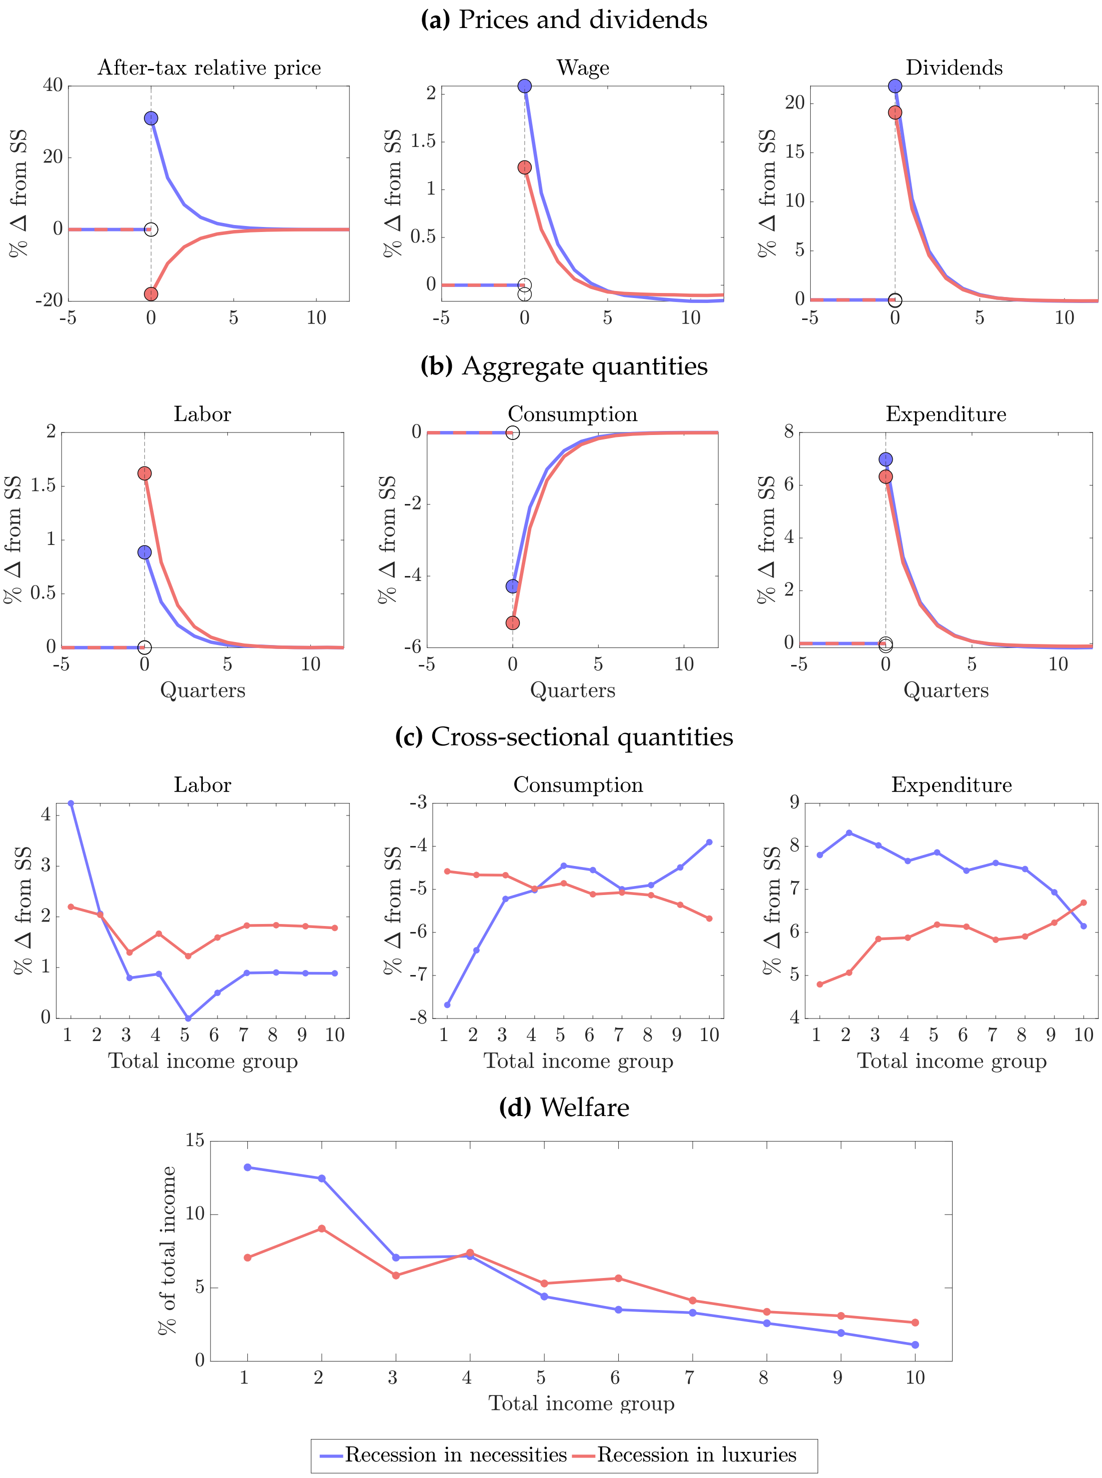

Note: The figure shows the aggregate and cross-sectional responses to a negative MIT shock on sectoral productivity of 10% for each sector. Panel (a) shows aggregate impulse responses at the time of the shock (dashed vertical line). Panel (b) shows the cross-sectional on-impact responses along total income deciles, where the response is computed as the percentage change one period after a negative sectoral productivity shock ($t = 6$) relative to the steady-state value. Total income is computed as the sum of labor income, dividend income and a household's asset position at each point in time. The last plot on welfare represents the compensation required in terms of total income to be as well off as in the absence of the recession - i.e., in the steady state. It is as a percentage of the steady state cash-on-hand value.

These findings are generally expected given our nonhomothetic preferences. The difference in responses across the income distribution stem from the fact that the expenditure shares of necessities fall in income. Moreover, since necessity goods are "necessary", households —especially low-consumption households— are less willing to tolerate a fall in necessity goods and thus willing to expend more to maintain their previous consumption levels.

This is also evident in the prices. In each recession, the price of the directly affected sector increases, while the price of the unaffected good falls. Yet, the price response in the necessities' recession is stronger than that during the luxuries' recession: after-tax relative prices increase by 30% during the necessity-sector recession, but only fall by 20% during the luxury-sector recession.

This pattern is less evident when looking at the labor supply response. Households increase their labor supply in response to both shocks. This is to finance higher expenditure, meaning that the income effect dominates over the substitution effect. However, aggregate labor increases more in the case of a recession in luxuries than in necessities. Cross-sectionally, we can see that low-income households increase their labor most during the necessity recession. The aggregate effects in labor however are mostly stemming from high-income households that increase labor supplied during the luxury recession.

Overall, aggregate consumption declines less during a necessity goods recession than during a luxury goods recession. This might suggest that necessity recessions are preferable to luxury recessions. However, our analysis indicates the opposite: while consumption of necessity goods falls by a smaller amount, this decline carries much greater welfare costs given the higher importance of necessities for households.

6.2 Welfare effects

Finally, we compute the welfare loss associated with each recession. For each worker, we compute the cash compensation required to render them indifferent utility-wise to the steady state. We do this at the onset of the recession. Our measure thus incorporates each worker's total loss of welfare, bot via the contemporaneous loss of utility, but also the lower expected future utility. Mathematically, our compensation is computed as κ:

$$V_{SS}(a, b, z; \Xi_{SS}) = V_R(a, b + \kappa, z; \Xi_R), \tag{24}$$

where the subindices SS and R denote the steady state, and the first quarter of the recession. For each worker, we scale κ by their total income.

Figure 9d presents the results for the welfare costs of sectoral recessions along total income deciles. We find that the average compensation required - i.e., the welfare loss

- in a necessities' recession corresponds to 5.7% of total income, whereas for luxuries' recessions this corresponds to 5.3% of total income.

Not only is the average compensation required higher for the necessity good recession, but there is also a bigger dispersion along income deciles: the absolute difference between the welfare loss experienced by the bottom and the top income deciles amounts to 12.1%, whereas, for luxuries, this corresponds to 6.4% — roughly half. Households in the bottom 30% of income are worse off in a recession in the necessities' sector, households in income decile 4 are roughly indifferent between the two, and households in the top 50% of income are worse off in a recession in the luxury good sector.

Our nonhomothetic preference structure creates cross-sectional differences in consumption patterns that translate into heterogeneous welfare costs of recessions, as measured by our percentage-loss-of-income measure. Necessity recessions lead to a larger loss of welfare, in particular for low-income households. This motivates fiscal policies that can alleviate these welfare losses, in particular for low-income households, which we will study next.

7 Fiscal Policy

Sectoral recessions impose considerable welfare costs, with recessions in the necessity sector disproportionately affecting lower-income households. This distributional asymmetry highlights the potential role for targeted government intervention through fiscal policy to mitigate recessionary impacts on vulnerable households or stimulate broader economic recovery.

We examine three fiscal policy options that gained prominence during the 2022-2023 inflationary episode: price subsidies on necessity goods, lump-sum transfers, and targeted transfers. Many countries adopted these measures primarily to reduce the effect of price increases on low-income households. Price subsidies, in particular, aimed both to relief low-income households and to directly stabilize prices by lowering necessity goods costs. In an attempt to target price subsidies to support low-income households, many countries focused these subsidies on necessity goods, to which low-income households are more exposed.

Real-world implementations varied considerably. Germany combined lump-sum energy allowances with targeted transfers for social benefit recipients and unemployed individuals, alongside price measures including temporary VAT cuts on natural gas, reduced fuel excises, and gas bill reimbursements. Portugal implemented targeted transfers to lower-income households and price subsidies through reduced taxes on oil products and temporary VAT cuts on certain food items.[20] Amores et al. (2023) provide a comprehensive overview of European fiscal policy responses to inflationary pressures.

We study whether a price subsidy in necessity goods actually alleviates the welfare losses associated with a recession in the necessity sector. In our model, we focus on effects in the immediate onset of the recession: the government subsidizes necessity good prices by 5% for one quarter.[21] As discussed previously, the government finances any fiscal policy intervention with proportional labor taxes.

We compare this price subsidy with two alternative fiscal policies. First, a targeted transfer, which mirrors the means-tested approach adopted by several countries, typically benefiting minimum-wage earners. These transfers reach approximately the bottom 20% of the population, which we replicate in our model. Second, for comparison, an untargeted transfer.

To maintain comparability, all three fiscal policies last one quarter, are budget-equivalent, and are financed through proportional labor taxes.

7.1 Aggregate and cross-sectional effects

All policies are only active in the first quarter after the shock. The economy responds to all policies similarly after the first quarter, and thus we focus here on studying the response of the economy during that first quarter.

Figure 10 plots in bar charts aggregate responses (such as prices in Panel 10a and quantities in Panel 10b), and in line-plots the cross-sectional response across the income distribution (Panel 10c).

For comparison, we plot in black the evolution of the economy absent any intervention. Nominal after-tax prices of the necessity good increase in all four scenarios, albeit to a lower extent during the VAT subsidy: the subsidy succeeds at lowering inflation, albeit at less than the 5% subsidy rate, since not all of the subsidy was passed along to the consumer. Wages increase somewhat during all interventions, but less so than after tax relative prices — real wages fall.[22]

Wages fall less in the price-subsidy scenario, since it stimulates production most — consequently, consumption falls less during the price subsidy than during the both transfers. Yet, relative to the baseline scenario, consumption falls during all interventions — even though labor increases —, indicative of the large distortions stemming from the

[20] See Bernardino et al. (2025) for an evaluation of this policy.

[21] Several countries implemented price subsidies to different necessity goods of similar magnitudes. For reference, Spain implemented a VAT cuts on selected food items of 4% and on electricity of 5%; Portugal implemented a VAT cut to selected food items of 6%.

[22] There is no single real wage in this economy, as each household has their own consumption bundle and price index. Yet, the strong increase in necessity prices implies that each household's real wage falls.

Figure 10: Impulse responses to a recession in the necessities' sector and fiscal policies

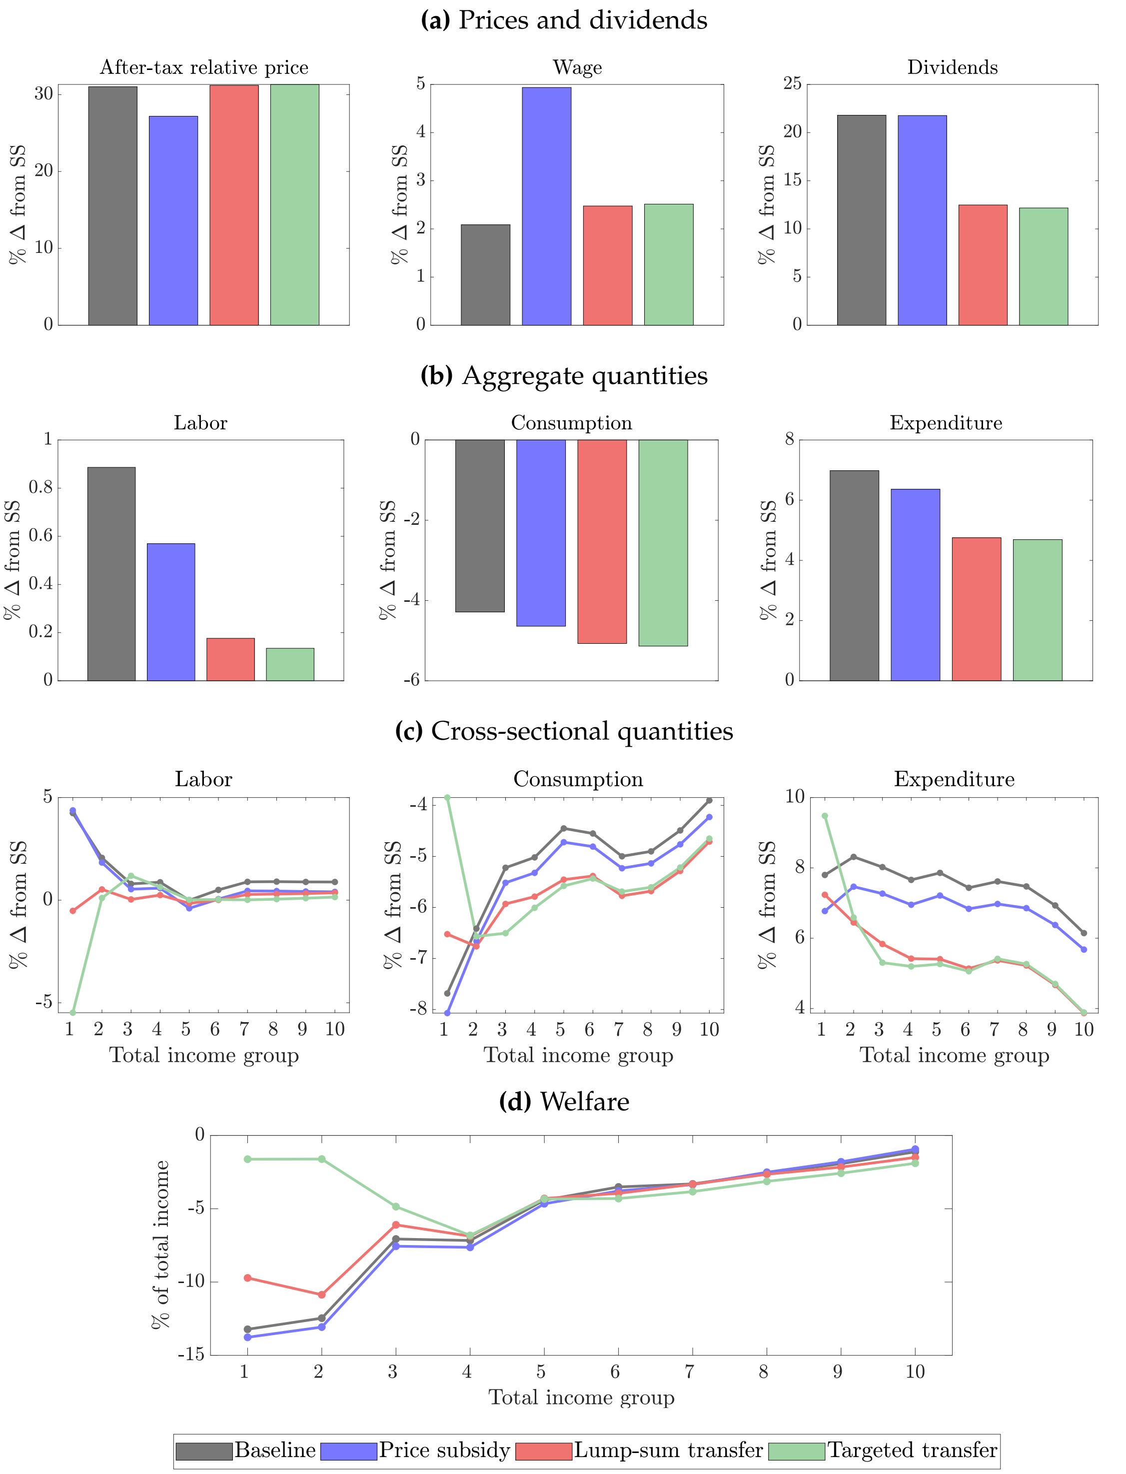

Notes: This figure shows the impulse responses to a 10% negative TFP shock in the necessities' sector combined with different fiscal policies. "Baseline" corresponds to the case with the shock and without any policy, whereas the other three correspond to cases where the same shock took place and the policy was implemented. The response is computed as the change on impact at the time of the shock ($t = 6$) relative to the steady-state value ($t = 1$). All policies are revenue-equivalent and last for one quarter. Panel (a) shows the aggregate response. Panel (b) shows the cross-sectional responses along total income deciles. The last plot on welfare corresponds to the amount of total income one would have to forego to go back to the steady-state as a percentage of the steady state cash-on-hand value.

increased labor tax.

The increase in the price subsidy leads to heightened dividends, disproportionately benefiting the higher-income demographic groups.

How does the fall in consumption and the increase in labor supply play out across the income distribution?

The changes in equilibrium labor supplied mostly stem from low-income households below the third income decile. In the price-subsidy scenario, they work as much as in the baseline scenario. In the transfer scenarios, they work less: their labor falls somewhat during the untargeted transfer. During the targeted transfer (i.e., when they receive a higher transfer), their labor supplied falls even below the steady state level. This is reasonable since lower productivity and an increased heightened labor tax should both push labor supplied below steady state values. For other income groups (and in other scenarios), labor supplied only remains above steady state to finance consumption.

Consumption is thus largely determined by heightened prices and lower labor income. The cross-sectoral consumption responses are consistent with that: consumption falls across the income distribution, but more so for low-income households that are more liquidity-constrained. The untargeted and targeted transfers put a dent in their liquidity constraint, thus increasing their consumption relative to the baseline scenario.

7.2 Welfare effects

Taking stock of these heterogeneous responses, we now ask what is the welfare gain of different policies relative to the baseline scenario without any policy. For that, we use as welfare measure our metric of compensating variation, as in the previous section: we compute the amount of cash-on-hand households are willing to pay to have the policy, which we then rescale by total income.

Averaging out across households, we estimate that the targeted transfer is the most welfare-improving policy, resulting in an average welfare gain of 2.3%, followed by the lump-sum transfer with a welfare gain of 0.7%. The price subsidy actually decreases average welfare by 0.2%.

Figure 10d plots the welfare gains of different policies across income deciles. We find both types of transfers do not significantly affect the welfare of workers with medium and high income: untargeted transfers improve the welfare of workers below the third decile — for the remainder of workers, the distortions introduced by the labor tax offset the transfers. Targeted transfers significantly improve the welfare of the bottom third deciles: the welfare effect is worth almost 10% of the income of workers in the lowest income decile.

Both transfer policies are on average welfare improving since they transfer resources to low-income workers with high marginal utility. The price subsidy, even though targeted on necessity goods, worsens welfare. The reason being that it subsidizes consumption throughout the entire income distribution, stimulating demand (of a good that is produced with lower productivity), which is then paid for by distortionary taxes. Untargeted transfers thus are better at alleviating the welfare losses done by the recession, even though they are to a large extent spent on medium and high-income workers.

8 Conclusion

We have shown that necessity good recessions impose substantially higher welfare costs than luxury good recessions, with the burden falling disproportionately on low-income households. While aggregate consumption falls more during luxury recessions, necessity recessions are most costly in terms of welfare, with a striking 12.1 percentage point difference in welfare losses between the bottom and top income deciles. This distributional asymmetry stems from low-income households' higher expenditure shares on necessities and their limited ability to substitute away from these essential goods when prices rise.

Perhaps more surprising is our finding that targeted fiscal interventions vary dramatically in their effectiveness, contradicting conventional policy intuitions. Targeted transfers to the bottom 20% of households generate the largest welfare gains at 2.3% of average income, with benefits reaching nearly 10% of income for the lowest decile. Untargeted lump-sum transfers provide moderate benefits at 0.7% of average income. Most counter-intuitively, price subsidies on necessity goods—despite directly addressing the source of the shock—actually reduce average welfare by 0.2%. This occurs because subsidies benefit all consumers regardless of income level, with much of the benefit flowing to higher-income households through increased firm dividends, while the distortionary labor taxes needed for financing impose costs across the income distribution.

These findings highlight a tension between theoretical economic efficiency and practical policy implementation constraints. While our analysis clearly demonstrates the superiority of targeted transfers, real-world examples like Germany's response to the 2022-2023 inflationary crisis reveal why countries often choose seemingly inferior instruments. Germany implemented extensive price subsidies—including VAT cuts on natural gas and food items—not because policymakers were unaware of targeting principles, but because they lacked the administrative infrastructure to rapidly deploy means-tested transfers to affected households. Price subsidies can be implemented immediately through existing tax and regulatory systems, requiring no new databases, eligibility verification processes, or distribution mechanisms. In contrast, targeted transfers demand sophisticated administrative capacity that may not exist or cannot be scaled quickly during crisis periods when immediate relief is paramount. Yet, we have shown in our analysis that even untargeted transfers outperform price subsidies.

References

Aguiar, M., & Bils, M. (2015). Has Consumption Inequality Mirrored Income Inequality? *American Economic Review*, *105*(9), 2725–2756.

Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, *109*(3), 659–684.

Amores, A. F., Basso, H. S., Flevotomou, M., Freier, M., Pidkuyko, M., Bischl, S., De Agostini, P., De Poli, S., Dicarlo, E., Maier, S., et al. (2023). Inflation, Fiscal Policy and Inequality. *European Central Bank Occasional Paper Series*, (330).

Andreolli, M., Rickard, N., & Surico, P. (2024). Non-essential business-cycles. *Mimeographed, London Business School*.

Andrews, I., Stock, J. H., & Sun, L. (2019). Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics*, *11*(1), 727–753.

Attanasio, O. P., & Pistaferri, L. (2016). Consumption inequality. *Journal of Economic Perspectives*, *30*(2), 3–28.

Benzarti, Y. (2024). Tax incidence anomalies. *Annual Review of Economics*, *17*.

Benzarti, Y., Carloni, D., Harju, J., & Kosonen, T. (2020). What Goes Up May Not Come Down: Asymmetric Incidence of Value-Added Taxes. *Journal of Political Economy*, *128*(12), 4438–4474.

Benzarti, Y., & Tazhitdinova, A. (2021). Do Value-Added Taxes Affect International Trade Flows? Evidence from 30 Years of Tax Reforms. *American Economic Journal: Economic Policy*, *13*(4), 469–489.

Bernardino, T., Gabriel, R. D., Quelhas, J., & Silva-Pereira, M. (2025). The full, persistent, and symmetric pass-through of a temporary vat cut. *Journal of Public Economics*, *248*, 105416.

Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, *82*(6), 2167–2196.

Carroll, C., Crawley, E., Frankovic, I., & Tretvoll, H. (2023). Welfare and spending effects of consumption stimulus policies.

Chetty, R., Guren, A., Manoli, D., & Weber, A. (2011). Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins. *American Economic Review*, *101*(3), 471–475.

Christiano, L., Eichenbaum, M., & Rebelo, S. (2011). When is the Government Spending Multiplier Large? *Journal of Political Economy*, *119*(1), 78–121.

Comin, D., Lashkari, D., & Mestieri, M. (2021). Structural Change with Long-Run Income and Price Effects. *Econometrica*, *89*(1), 311–374.

Correia, I., Farhi, E., Nicolini, J. P., & Teles, P. (2013). Unconventional fiscal policy at the zero bound. *American Economic Review*, *103*(4), 1172–1211.

Engel, E. (1857). Die Produktions- und Consumtionsverhältnisse des Königreichs Sachsen. *Zeitschrift des Statistischen Büreaus des Königlich Sächischen Ministeriums des Innern*, *8-9*(2), 1–54.

Geary, R. C. (1950). A Note on "A constant-Utility Index of the Cost of Living". *The Review of Economic Studies*, *18*(1), 65–66.

Havránek, T. (2015). Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting. *Journal of the European Economic Association*, *13*(6), 1180–1204.

Hobijn, B., & Lagakos, D. (2005). Inflation Inequality in the United States. *Review of Income and Wealth*, *51*(4), 581–606.

Hochmuth, P., Pettersson, M., & Weissert, C. J. (2023). A nonhomothetic price index and cost-of-living inequality. *Unpublished manuscript.(Cited on pp. 1–3, 5–9, 20.)*

Huggett, M. (1993). The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies. *Journal of Economic Dynamics and Control*, *17*(5-6), 953–969.

Krueger, D., Mitman, K., & Perri, F. Macroeconomics and Household Heterogeneity. In: In *Handbook of macroeconomics*. Vol. 2. Elsevier, 2016, pp. 843–921.

Känzig, D. R. (2021). The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements. *American Economic Review*, *111*(4), 1092–1125.

Matsuyama, K. (2023). Non-ces aggregators: A guided tour. *Annual Review of Economics*, *15*(1), 235–265.

Montiel-Olea, J. L., & Pflueger, C. (2013). A robust test for weak instruments. *Journal of Business & Economic Statistics*, *31*(3), 358–369.

OECD. (2023). Tax Policy Reforms 2023: OECD and Selected Partner Economies. *OECD Publishing, Paris*.

Orchard, J. (2024). Cyclical demand shifts and inflation inequality. *Working Paper*.

Sonnervig, M. (2023). Unequal business cycles. *Working Paper*.

Stepner, M. (2013). Binscatter: Stata module to generate binned scatterplots.

Stone, R. (1954). Linear Expenditure Systems and Demand Analysis: an Application to the Pattern of British Demand. *The Economic Journal*, *64*(255), 511–527.

Tauchen, G. (1986). Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions. *Economics Letters*, *20*(2), 177–181.

A Additional Empirical Figures and Tables

A.1 Engel curves

Using the estimated semi-elasticities presented in Figure 2 is consistent with classifying goods using the slope of Engel curves. For robustness to these semi-elasticities, we present the Engel curves for each category using a non-parametric binned scatter approach that groups income into equal-sized bins, computes the mean of both income and the budget share within each bin, and plots these points to non-parametrically obtain the conditional expectation of the budget share as a function of the logarithm of pre-tax income.

We consider the period of 2000-2021 and include time fixed effects to control for any time-specific shocks. The red line corresponds to a quadratic polynomial fit of the binned scatterpoints. See Stepner (2013) for a detailed description of Stata's `binscatter` command.

Figure 11: Budget shares as a function of income

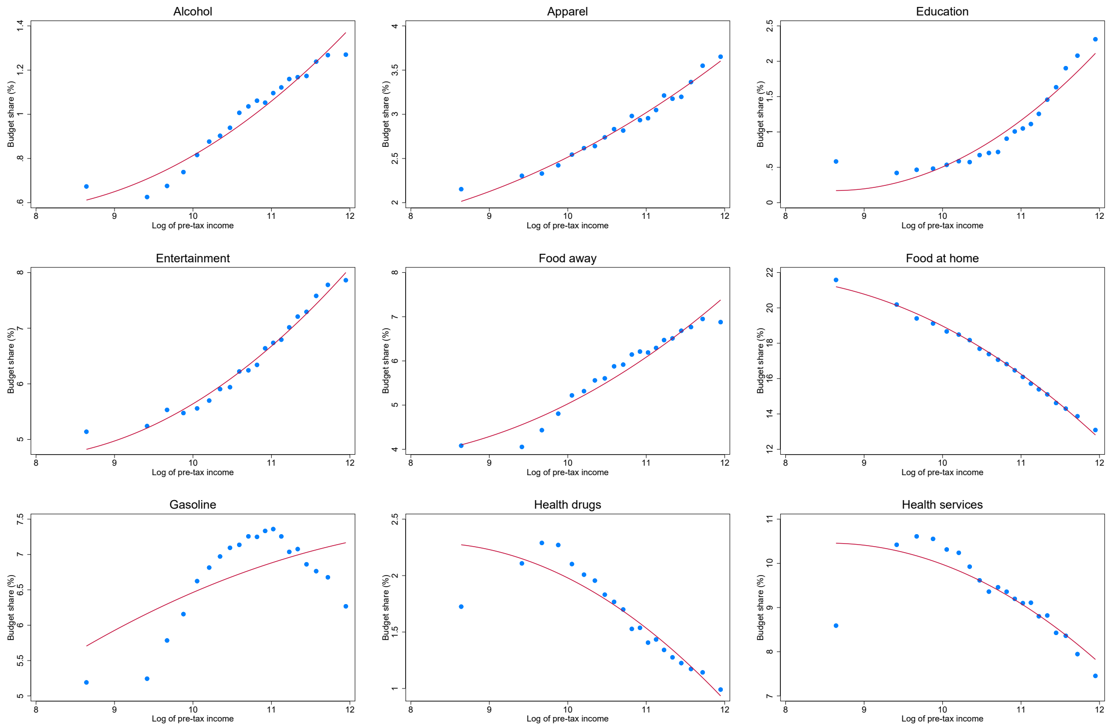

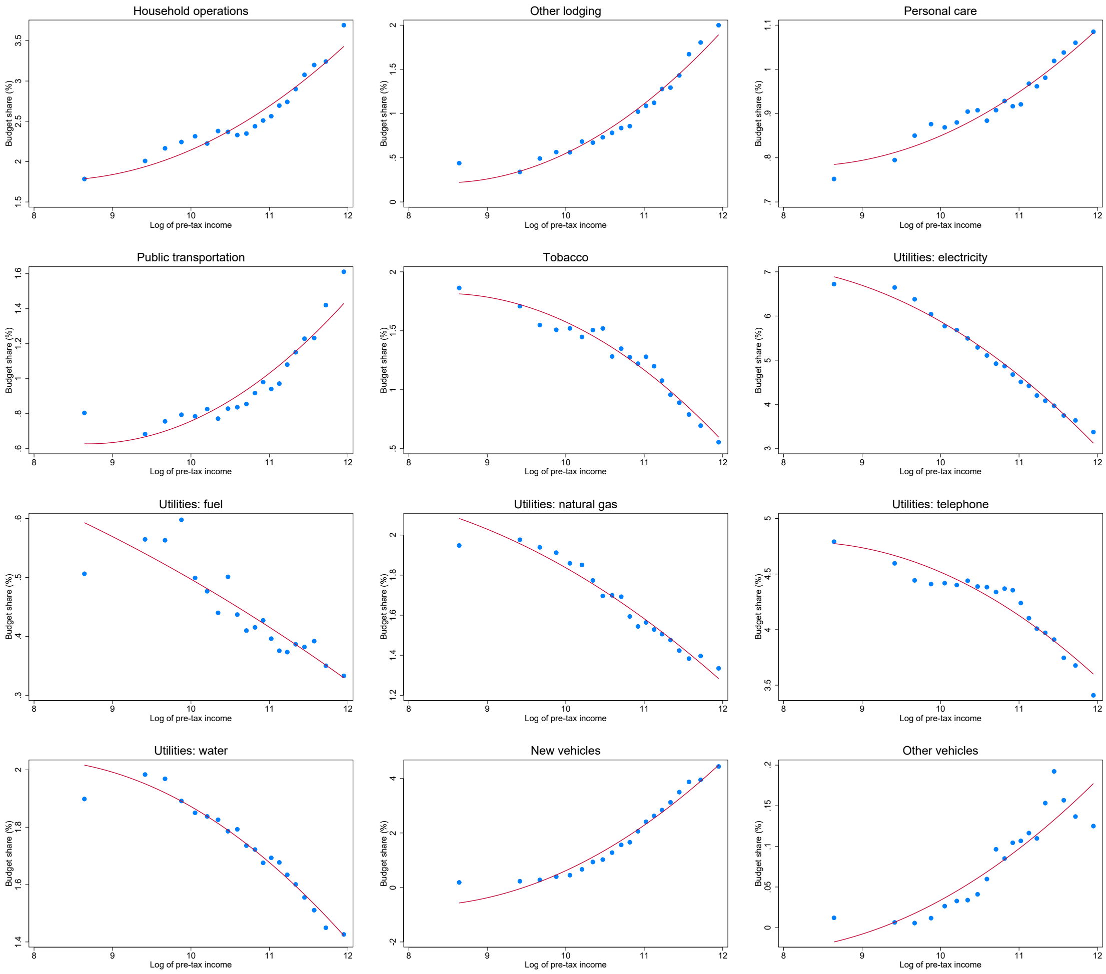

A.2: First stage results

Table 6: First stage

	$\Delta\mathcal{P}$
Coefficient	0.17
Standard error	0.04
Number of observations	279,136
Number of clusters	97,126
Kleibergen-Paap LM test	
Chi-squared	19.26
p-value	0.00
Kleibergen-Paap *F*-stat	19.28
Cragg-Donald *F*-stat	39.26
Stock-Yogo weak ID test CV	
10% maximal IV size	16.38
15% maximal IV size	8.96
20% maximal IV size	6.66
25% maximal IV size	5.53
Effective *F*-statistic	80,514.25
CV τ=5%	37.42
CV τ=10%	23.11
CV τ=20%	15.06
CV τ=30%	12.04

Notes: