Optimal Monetary Policy during a Cost-of-Living Crisis*

Alan Olivi,† Vincent Sterk,‡ and Dajana Xhani§

March 6, 2024

Abstract

How should monetary policy react to sectoral shocks in a world where consumption baskets and demand elasticities vary across households? We present a multi-sector New-Keynesian model with generalized, non-homothetic preferences and inequality. The output gap is governed by a *Marginal* Consumer Price Index (MCPI), rather than the regular CPI. Policy trade-offs are shaped by two novel wedges in the New-Keynesian Phillips Curve (NKPC). Analytical results and quantitative simulations show that, following a negative shock to necessity sectors, the NKPC is shifted upward, increasing CPI inflation but decreasing the output gap. We find that the optimal policy response is relatively accommodative.

JEL classification: E21; E25; E31; E52

Keywords: Sectoral shocks, Non-homothetic Preferences, Inequality, Optimal Policy

*For helpful comments, we are grateful to Jordi Galí, Greg Kaplan, Kurt Mitman, Benjamin Moll, Andreas Schaab, Kathrin Schlaffmann, Elisa Rubbo, Ludwig Straub, Gianluca Violante, and seminar/conference participants at the BSE Summer forum 2023, CEBRA 2023, Bank of Spain, CEMFI, De Nederlandsche Bank, ECB, Einaudi Institute for Economics and Finance, National Bank of Belgium annual conference 2022, NBER Summer Institute 2023, NHH Bergen, Queen Mary University of London, SED 2023, Sciences Po, St. Louis Fed, the Toulouse School of Economics, University College London, Universitat Pompeu Fabra (CREi), University of Bonn, University of Copenhagen, University of Exeter, University of Surrey, and the University of Tilburg.

†University College London. E-mail: a.olivi@ucl.ac.uk.

‡University College London, NHH Bergen, and CEPR. Email: v.sterk@ucl.ac.uk.

§Tilburg University. E-mail: D.Xhani@tilburguniversity.edu.

1 Introduction

Since 2020, many economies have been confronted with large supply disruptions, resulting from the COVID-19 pandemic, the war in Ukraine, and other shocks. A significant surge in inflation followed, particularly in sectors producing necessities like food and energy. Low-income households have often been disproportionately affected, as they tend to allocate a larger proportion of their expenditures towards such goods.[1] Indeed, the strong squeeze in real incomes, in particular among the poorest households, has led many commentators to declare the situation a "cost-of-living crisis".

To central banks, these events underscored important yet unresolved outstanding questions: How to conduct monetary policy in a world with diverse consumption baskets, and thus heterogeneity in inflation rates across households? Do supply shocks to specific sectors, producing either necessity or luxury goods, call for a specific policy response? How do the distributional implications of such shocks affect monetary policy trade-offs? Is the Consumer Price Index (CPI) still a suitable target for monetary policy?

To answer these questions in a comprehensive way, the standard New Keynesian model –a standard framework for monetary analysis– is arguably not well suited, even when extended with sectoral heterogeneity and inequality in household income and wealth. A key limitation is that preferences are typically assumed to be of a homothetic-CES (Constant Elasticity of Substitution) form, which implies that the composition of consumption baskets is equal across households. As a result, all households share the same price index and are equally affected by sector-specific price increases, unlike in reality.

This paper presents a novel New Keynesian model which incorporates (i) multiple sectors, (ii) permanent income and wealth heterogeneity, and (iii) generalized, non-homothetic preferences, represented through "sufficient statistics" rather than a specific functional form. In this setting, each household has an individual consumption basket, creating heterogeneity in individual inflation rates, real wages and real interest rates. The generalization of preferences may also give rise to heterogeneity in price elasticities of demand across consumers. For example, rich households may not only have high overall levels of expenditures, but may also react less strongly to changes in prices of individual goods. In this generalized environment, we examine both the positive and normative implications of aggregate and sector-level shocks.

Towards this end, we derive an analytical characterization of the model and show that two novel wedges emerge in the New Keynesian Phillips Curve (NKPC): a *non-homotheticity wedge* and an *endogenous markup wedge*. Importantly, the joint movement in these wedges shifts the NKPC in a direction that depends on the sectoral source of the shock. Specifically, a negative

[1] According to the Office for National Statistics, in October 2022, UK households in the lowest income decile faced on average a nearly 3 percentage points higher rate of inflation than those in the highest income decile, see ONS (2022).

productivity shocks to necessity sectors initially leads to an upward shift of the NKPC, increasing inflation but reducing the output gap. By contrast, shocks to aggregate productivity, or productivity in luxury sectors, tend to move the output gap and inflation in the same direction, as is usually the case in the New Keynesian model. A cost-of-living crisis thus poses a specific challenge to a central bank seeking to stabilize inflation and the output gap, even setting aside any distributional concerns pertaining to such a situation.

In order to draw normative lessons, we study the welfare-optimal monetary policy response to aggregate and sectoral productivity shocks, and compare it to the prescription of a standard interest rate rule targeting CPI inflation. In a simplified version of the model, we show analytically that the optimal policy stance following a negative necessity shock is initially relatively loose, because of the upward shift in the NKPC mentioned above. A swift and strong increase in interest rate could bring down inflation, but only at the expense of a strongly negative output gap, which is not optimal. However, later on the optimal policy tightens, which is qualitatively in line with the delayed tightening by several central banks in response to the recent shocks.

An important implication of non-homothetic preferences is that households devote a relatively large fraction of *marginal* spending to luxuries. Indeed, even a household which spends most of its budget on necessities may still allocate most of any additional spending to luxuries. Accordingly, the real wage which guides marginal saving and labor supply decisions is one which deflates the nominal wage with a *Marginal CPI* (MCPI), weighing sectors by marginal rather than regular budget shares and thus down-weighting necessities compared to the regular CPI. We show that output gap dynamics are associated with the MCPI rather than the regular CPI. Therefore, the MCPI complements the CPI as a natural metric to guide monetary policy.

To better understand the policy trade-offs, we study the two novel NKPC wedges in detail. The first is the *non-homotheticity wedge*. This wedge captures a labor market distortion which arises due to the gap in marginal and regular budget shares on different sectors, in the presence of price rigidities. To understand this wedge intuitively, consider a shock which simultaneously decreases productivity in necessity sectors but increases productivity in luxury sectors. Following this shock, luxury goods become cheaper relative to necessity goods. This increases the real wage in units of households' marginal consumption bundles, since at they margin they spend relatively more on luxuries. In turn, the increase in the marginal real wage induces households to optimally increase labor supply. However, when prices are sticky this increase is diminished, because relative prices react less strongly. As a result, labor supply is distorted downwards and the output gap becomes negative for a given inflation rate or –equivalently– inflation increases for a given output gap. A decrease in the relative productivity of necessity sectors thus shifts up NKPC, while a decrease in the relative productivity of luxury sectors has the opposite effect.[2]

The second wedge in the NKPC is the *endogenous markup wedge*, which arises from the fact that price elasticities of demand for goods vary across households and over time, once we move beyond homothetic-CES preferences over varieties within sectors. Realistically, poorer households are likely to be more price sensitive and demand elasticities may increase during recessions, as consumption falls. For firms, demand elasticities are in turn a key consideration when setting markups. Fluctuations in the level and distribution of consumption thus create fluctuations in demand elasticities and hence distortions in markups. Specifically, the wedge tends to shift the NKPC *downward* after negative productivity shocks. Compared to the non-homotheticity wedge, the movements in the endogenous markup wedge tend to be smaller but more persistent. The combined effect of the two wedges is that, following a negative shock to necessity sectors, the NKPC is initially shifted upward, but downward later on, calling for a specific dynamic policy response which depends on the sectoral origin of the shock.

In addition to the analytical results derived in the simplified model, we conduct a quantitative exploration in a full-blown version of the model, calibrated to the United Kingdom. The model features realistic heterogeneity in income, wealth, expenditure baskets, and marginal propensities to consume, disciplined by data from the Living Costs and Food (LCF) survey. We also allow for heterogeneity in price rigidities across sectors and input-output linkages. Despite its richness, the model is computationally tractable, up to a first-order approximation, as we can characterize the dynamic equilibrium with as a system of sector-level NKPCs and Euler equations, alongside two sector-level equations tracking the relevant aspects of the wealth distribution.

Model simulations reveal that the channels and policy trade-offs highlighted analytically are also important quantitatively. We observe that, under a standard interest rate rule targeting CPI inflation, negative shocks to necessity sectors, such as *Food* or *Electricity & Gas*, lead to an increase in CPI inflation but an initial decline in the output gap, followed by a subsequent upswing. By contrast, after a negative shock to productivity in all sectors, or only in luxury sectors, CPI inflation and the output gap both increase persistently. Regarding the distributional impact of aggregate and sectoral shocks, we also find strong heterogeneity in the consumption responses of individual households, depending not only on their income and wealth but also on their expenditure baskets.[3]

Next, we solve for the welfare-optimal interest rate path in response to sectoral shocks.

[2]While this channel also arises in a representative-agent version of the model with non-homothetic preferences, its strength depends on the degree of long-run inequality. And importantly, empirical discipline on the channel is critically obtained from cross-sectional evidence on the relation between income and expenditures on different goods, which is at odds with a representative-agent assumption.

[3]In addition to non-homothetic preferences, the model includes idiosyncratic preference shifters for goods from different sectors, allowing us to match exactly the heterogeneous consumption baskets observed in micro data.

We do so analytically in a simplified version of the model, as well as quantitatively in the full-blown model. We find that, compared to a standard interest rate rule, the optimal policy response to a negative necessity shock is initially significantly more accommodative, i.e. the interest rate is held relatively low. For shocks to luxury sectors, we find the opposite. Later on, the optimal policy stance tightens, in particular following necessity shocks. Moreover, we find that potential distributional considerations lead to an overall looser monetary policy reaction to negative productivity shocks, as lower interest rates redistribute wealth towards poorer households, who tend to be more heavily affected by such shocks.[4]

Relation to the literature. A main contribution of this paper is to embed a generalized, non-homothetic preference structure in a multi-sector New Keynesian model, allowing for household inequality. Empirical evidence supporting the relevance of non-homothetic preferences has a long history in the literature. A particularly famous and robust finding is that expenditure shares on food are negatively related to income (Engel, 1857; Houthakker, 1957). It is also understood that these patterns have important implications for the aggregate price indices and the measurement of inequality, see e.g. Hamilton (2001); Kaplan and Schulhofer-Wohl (2017); Jaravel (2019); Argente and Lee (2021). While in this paper we focus on monetary policy and business cycles, others have studied the implications for non-homothetic preferences for growth and structural transformation (e.g. Herrendorf et al. (2014); Boppart (2014); Comin et al. (2021)). Non-homothetic preferences are also recognized to have important implications for tax policy, see Jaravel and Olivi (2021). We further relate to literature which deviates from CES preferences over goods varieties, e.g. Kimball (1995); Amiti et al. (2019); Xhani (2021) and which studies how demand elasticities and markups vary across the income distribution, see e.g. Mongey and Waugh (2023); Nord (2023); Sangani (2023).

The New-Keynesian literature typically sticks to the simplifying assumption of homothetic-CES preferences.[5] Such models therefore abstract from heterogeneity in consumption baskets even when they feature household heterogeneity.[6] Indeed, the mechanisms that we highlight complement the channels highlighted in the literature on monetary policy transmission with heterogeneous agents, see e.g. McKay et al. (2016); Kaplan et al. (2017); Auclert (2019) and many others. This literature often emphasizes the role of heterogeneity in Marginal Propensities to Consume (MPCs), a micro-level non-linearity which makes distributions matter for macroe-

[4]This is the case even though our assumed social welfare function is such that monetary policy has no motive to affect steady-state inequality.

[5]Some authors in this literature have deviated from CES utility by assuming a Kimball demand function, see e.g. Smets and Wouters (2007). However, such preference preserve homotheticity and do not create endogenous markup fluctuations. Cavallari and Etro (2020) consider a representative-agent model with extended CES preferences which delivers a time-varying price elasticities of demand.

[6]One exception is Blanco and Diz (2021) who study a representative-agent household NK model with two consumption goods, one of which is subject to a subsistence point. Another one is Melcangi and Sterk (2019), who develop a heterogeneous-agents New Keynesian model with an infrequently consumed luxury good.

conomic outcomes. While we connect to this literature, our analysis highlights heterogeneity in consumption behaviour generated by non-homothetic, non-CES preferences. This form of household heterogeneity matters not only for the demand block of the model (characterised by Euler equations and household constraints) but also for the supply block of the model, as characterised by the NKPC. Indeed, we show that household heterogeneity affects both the slope of the NKPC and the time-varying wedges that emerge under generalized preferences.

The normative analysis in this paper connects to the literature on optimal policy in the NK model, see e.g. Galí (2015) and references therein, and on how inequality and redistribution affect optimal monetary policy trade-offs, including redistributive effects, see Challe (2020); Bhandari et al. (2021); Nuno and Thomas (2022); Dávilla and Schaab (2022); Acharya et al. (2023); McKay and Wolf (2023). Our model abstracts from idiosyncratic risk. Instead, we show how non-homothetic, non-CES preferences, combined with permanent inequality, gives rise to novel policy trade-offs. Finally, the multi-sector structure of our model connects our contribution to several recent papers on intersectoral transmission of shocks in neoclassical models and (HA)NK models, including Baqaee and Farhi (2019); Pasten et al. (2020); Rubbo (2023); LaO and Tahbaz-Salehi (2019); Baqaee et al. (2021); Moll et al. (2023); Schaab and Tan (2023); Auclert et al. (2023).

The remainder of this paper is organized as follows. Section 2 lays out the primitive model environment and provides an analytical characterization of the model. In Section 3 we inspect the mechanisms in a relatively simple version of the model, focusing on the role of the two new wedges in the NKPC. Results for the full quantitative model are presented in Section 4. Optimal policy is discussed in Section 5. Section 6 concludes. The Online Appendix provides supplemental, technical material.[7]

2 The model

2.1 Environment

Households. There is a continuum of heterogeneous households, of unit mass and indexed by i. In every period t, a household dies with a probability $\delta \in (0,1)$. Households consume goods from different sectors, indexed by $k = 1, 2.., K$. Within each sector, there is a unit mass continuum of differentiated varieties, indexed by j. The expected utility of household i at time t is given by:

$$\mathbb{E}_t \sum_{s=0}^{\infty} (\beta(1-\delta))^{t+s} \left(u_i(\mathbf{c}_{t+s}(i)) - \chi \left(\frac{n_{t+s}(i)}{\vartheta(i)} \right) \right), \tag{1}$$

[7] We apologize for the lengthy derivations underlying the model and in particular the optimal policy results. These derivations involve mostly straightforward algebra and we present them for full transparency.

where $n_{t+s}(i)$ is effective labor supply, $\vartheta(i)$ is labor productivity, $\beta \in (0,1)$ is the subjective discount factor, and $\mathbb{E}_t$ is the conditional expectations operator. Moreover, the utility from consumption depends on a vector $\mathbf{c}_t(i) = \{\mathbf{c}_{1,t}(i), .., \mathbf{c}_{K,t}(i)\}$, where $\mathbf{c}_{k,t}(i)$ is in turn a vector consisting of the consumption of each variety j in sector k. Specifically, the flow utility from consumption is given by:

$$u_i(\mathbf{c}_t(i)) = U_i(\mathcal{U}(\mathbf{c}_{1,t}(i)), ..., \mathcal{U}(\mathbf{c}_{K,t}(i))),$$

where $U_i(\cdot)$ is an outer utility function, defined over sectoral bundles, which may be household specific. We assume that $U_i(\cdot)$ is differentiable and weakly separable across sectors. The sectoral bundles are in turn given by $\mathcal{U}(\mathbf{c}_{k,t}(i))$. We further assume that the inner utility function $\mathcal{U}(\cdot)$ is a concave, C^3-function which is symmetric over varieties. Moreover, $\chi(\cdot)$ is an increasing, twice differentiable function capturing disutility from labor supply.

Households can save in one-period nominal bonds, denoted by $b_t(i)$ and they are born with different initial levels of nominal wealth. Households also differ in terms of their labor productivity, $\vartheta(i)$, which is constant over time. We thus abstract from idiosyncratic risk, aside from mortality risk. We do allow for the possibility that some households are Hand-to-Mouth (HtM) consumers, which we treat as a permanent characteristic.[8] HtM households cannot adjust their bond holdings, and thus consume their current incomes. Households who are not HtM can choose bond holdings freely, facing only a natural borrowing limit. Households further differ in their ownership of firms. The budget constraint of household i in period t is given by:

$$e_t(i) + \frac{b_{t+1}(i)}{R_t} = b_t(i) + n_t(i)W_t + \sum_k \varsigma_k(i) Div_{k,t}. \tag{2}$$

Here, $e_t(i) = \sum_{k=1}^{K} e_{k,t}(i) = \sum_{k=1}^{K} \int_0^1 p_{k,t}(j) c_{k,t}(i,j) dj$ denotes the household's total consumption expenditures, R_t is the gross nominal interest rate on bonds, which is set by a central bank, W_t is the nominal wage per effective unit of labor, $Div_{k,t}$ are total dividends from sector k and $\varsigma_k(i)$ is the equity share of household i in firms in sector k. We assume that equity portfolios are perfectly diversified within and across sectors.

In any period t, household i chooses consumption of each goods variety, $c_{k,t}(i,j)$, bond holdings, $b_t(i)$, and effective labor supply, $n_t(i)$, to maximize utility objective (1), subject to the budget constraint (2) and the laws of motion of equilibrium objects exogenous to households. HtM households in addition face the constraint $b_t(i) = b_{t-1}(i)$.

[8] Even without HtM households, distributional dynamics will generally matter for aggregates, due to the non-linearities embedded in the generalized, non-homothetic and non-CES preferences.

Table 1. Steady-state statistics

	Individual	Aggregate
Marginal Propensity to Consume:	$MPC(i) = \frac{\partial e_t(i)}{\partial b_t(i)}$	
Budget share:	$s_k(i) = \frac{e_k(i)}{e(i)}$	$\bar{s}_k = \frac{E_k}{E}$
Marginal budget share:	$\partial_e e_k(i) = \frac{\partial e_k(i)}{\partial e(i)}$	$\overline{\partial_e e_k} = \int \frac{e(i)}{E} \partial_e e_k(i) di$
Cross-price elasticity:	$\rho_{k,l}(i) = \frac{\partial c_k(i)}{\partial P_l} \frac{P_l}{c_k(i)}$	$\bar{\rho}_{k,l} = \frac{\partial C_k}{\partial P_l} \frac{P_l}{C_k}$
Demand elasticity:	$\epsilon_k(i) = -\frac{\partial c_k(i,j)}{\partial p_k(j)} \frac{p_k(j)}{c_k(i,j)}$	$\bar{\epsilon}_k = \int \frac{e_k(i)}{E_k} \epsilon_k(i) di$
Super-elasticity:	$\epsilon_k^s(i) = \frac{\partial \epsilon_k(i)}{\partial p_k(j)} \frac{p_k(j)}{\epsilon_k(i)}$	$\bar{\epsilon}_k^s = \frac{\partial \bar{\epsilon}_k}{\partial p_k(j)} \frac{p_k(j)}{\bar{\epsilon}_k}$
Markup sensitivity w.r.t. expenditures:	$\gamma_{e,k}(i) = \frac{\partial \mu_k}{\partial e_k(i)} \frac{E_k}{\mu_k}$	
Markup sensitivity w.r.t. wealth:	$\gamma_{b,k}(i) = \frac{\partial \mu_{k,t}}{\partial b_t(i)} \frac{E}{\mu_k}$	

Note: all statistics are evaluated in the deterministic steady state with zero inflation. $E_k = \int e_k(i) di$ are aggregate expenditures on sector k and $E = \sum_k E_k$ are total expenditures across all sectors. Moreover, $C_k = E_k / P_k$ is aggregate sectoral consumption. Finally, $\rho_{k,l}(i)$ is a compensated elasticity.

Some key statistics. In the absence of a parametric form for preferences, let us introduce some key concepts regarding household behavior. As discussed in Appendix A, we can express the demand of household i for a certain goods variety as a function of its price, $p_{k,t}(j)$, a vector of all other prices in the sector, denoted $\mathbf{p}_{k,t}$, and the total expenditures of the household on sector-k goods, $e_{k,t}(i)$. We denote this demand function by $c_{k,t}(i,j) = d_k\left(p_{k,t}(j), \mathbf{p}_{k,t}, e_{k,t}(i)\right)$.

We can now define a number of household-level statistics, evaluated at the deterministic steady state of the model, which we indicate by omitting the time subscript. We consider a steady state with zero inflation and therefore equal prices within sectors, i.e. $p_k(j) = P_k$ for any variety j in sector k where P_k is the sectoral price level. Note that in such a steady state it holds that $c_k(i,j) = c_k(i)$. Table 1 defines the statistics, which may all vary across households. The table also presents a number of aggregate counterparts that will play a role in the model.

The first statistic is the Marginal Propensity to Consume, often emphasized in the heterogeneous-agents literature. In our setting, we can derive $MPC(i) = \frac{R-1}{R} / \left(1 + \frac{Wn(i)\psi}{e(i)\sigma}\right)$ for non-HtM households and $MPC(i) = 1 / \left(1 + \frac{Wn(i)\psi}{e(i)\sigma}\right)$ for HtM households. Within both groups of households, there is MPC heterogeneity resulting from differences in the wealth effect on labor supply, which in turn is due to differences in the composition of financial versus human wealth.

The next three statistics in the table derive from the outer utility function $U_i(\cdot)$ and thus pertain to the allocation of household expenditures over sectors. First, $s_k(i)$, is the regular budget share, i.e. the fraction of expenditures that household i devotes to sector k. Its aggregate counterpart, $\bar{s}_k$, is used to construct the Consumer Price Index, which is defined as $P_{cpi} = \sum_k \bar{s}_k P_k$. Second, $\partial_e e_k(i)$, is the household's *marginal* budget share on sector k. It measures the fraction of each marginal unit of expenditures that the household devotes to goods in sector k. This statistic is not much emphasized in the literature on macroeconomic fluctuations. Indeed, under homothetic preference we obtain $\partial_e e_k(i) = s_k(i)$. However, in our model preferences are non-homothetic and the gap between the two statistics will play an important role. The aggregate (expenditure-weighted) counterpart of the marginal budget share is $\overline{\partial_e e_k}$. At the margin, households tend to spend less on necessity goods than they do on average, whereas the opposite is true for luxuries. Accordingly, we label k a necessity sector if $\overline{\partial_e e_k} < \bar{s}_k$, and a luxury sector if $\overline{\partial_e e_k} > \bar{s}_k$.

For later use, we define the *Marginal CPI* (MCPI) as $P_{mcpi} = \sum_k \overline{\partial_e e_k} P_k$. This price index weighs sectors by their marginal rather than their regular budget shares. Relative to the CPI, the MCPI thus overweights luxury sectors and underweights necessity sectors.[9] Note that under homothetic preferences over sectors, marginal and regular budget shares coincide, so that the CPI and MCPI become equal. The final statistic relating to the outer utility function is $\rho_{k,l}(i)$, the compensated elasticity of consumption by household i of sector-k goods with respect to a change in P_l, the price of sector-l goods.[10] Moreover, $\bar{\rho}_{k,l}$ is the aggregate counterpart.

The remaining statistics pertain to the inner utility $\mathcal{U}$, which defines utility over varieties within a sector. These statistics will be key determinants of markups in the model. The first, $\epsilon_k(i)$, is the elasticity of demand for a variety with respect to its price $p_k(j)$. Note that this elasticity varies not only across sectors, but also across households. When setting the markup, firms consider the aggregate demand elasticity for their good, $\bar{\epsilon}_k$, which weighs individual elasticities by expenditure shares. The steady-state markup is given by $\mu_k = \frac{\bar{\epsilon}_k}{\bar{\epsilon}_k - 1}$. While $\epsilon_k(i)$ denotes the demand elasticity at the steady state, the distribution of demand elasticities moves around over time: as households change their levels of expenditures, their demand elasticities change. The response of the individual demand elasticity to a change in the price is given by the price super-elasticity of demand, denoted by $\epsilon^s_k(i)$, as defined in the table.[11] Under CES preferences, demand elasticities are constant and hence $\epsilon^s_k(i) = 0$, but once moving beyond CES this is no longer the case. The super-elasticity of *aggregate* demand for sector-k varieties can be expressed as $\bar{\epsilon}^s_k = \left(\int \epsilon^s_k(i)\epsilon_k(i)\frac{e_k(i)}{E_k}di - \int \left(\epsilon_k(i) - \bar{\epsilon}_k\right)^2 \frac{e_k(i)}{E_k}di\right)/\bar{\epsilon}_k$. This object takes into account that a change in prices not only affects $\bar{\epsilon}_k$ via changes in individual demand (the first term) elasticities, but also through changes in the composition of demand (the second term).

When moving beyond CES preferences, different households thus contribute differently to markups, depending on their price elasticities of demand, their super-elasticities, and their share in aggregate expenditures. We define two additional statistics which capture the combined effects of this. First, $\gamma_{e,k}(i)$ measures the sensitivity of the markup with respect to in-

[9] One may think of "Core CPI" –a popular index in practice– as an extreme sibling of the MCPI, in the sense that it completely disregards prices in two of the most important necessity sectors: Food and Energy.

[10] More formally, we define in the derivation appendix $\rho_{k,l}(i) = \left(P_l \partial_{P_l} e_k(i) + \partial_e e_k(i) e_l(i)\right)/e_k(i)$ for $l \neq k$ and $\rho_{k,l}(i) = \left(P_k \partial_{P_k} e_k(i) + (\partial_e e_k(i) - 1)e_k(i)\right)/e_k(i)$. Given that subvariety prices are initially equal within sectors and preferences for subvariety are symmetric, $\rho_{k,l}(i)$ captures the change in spending on product k, through substitution effects, for any first order change in subvariety prices in sector l.

[11] Note that, due to symmetry and anticipating that in the steady state firms are identical within sectors, $\epsilon_k(i)$ and $\epsilon^s_k(i)$ do not depend on j, i.e. at the steady state these elasticities are the same for all varieties within a sector.

dividual i's expenditures on sector-k goods: $\gamma_{e,k}(i) = \left(1 - \frac{\epsilon_k(i)}{\bar{\epsilon}_k}\left(1 + \frac{\partial \epsilon_k(i)}{\partial e_k(i)} \frac{e_k(i)}{\epsilon_k(i)}\right)\right) \frac{1}{\bar{\epsilon}_k - 1}$. Intuitively, if there is a relative increase in expenditures among households who have relatively low demand elasticities, the aggregate demand elasticity decreases, pushing up markups. A similar effect takes place if there is a shift in expenditures towards households whose price elasticity of demand is relatively insensitive to the level of expenditures. The second statistic, $\gamma_{b,k}(i)$, captures the markup sensitivity with respect to individual wealth, which we can express as $\gamma_{b,k}(i) = MPC(i)\gamma_{e,k}(i)\partial_e e_l(i)/\bar{s}_k$. Note that under CES preferences we obtain $\gamma_{e,k}(i) = \gamma_{b,k}(i) = 0$.

Finally, we assume that the Elasticity of Intertemporal Substitution (EIS) and the Frisch elasticity of labor supply are homogeneous across households, and denote them by σ and ψ respectively. It is possible to allow for heterogeneity in these objects as well, at the expense of somewhat more complicated algebraic expressions.[12]

Firms. Firms are monopolistically competitive, each producing a single goods variety j in a certain sector k. Within each sector, firms are ex-ante identical but subject to a Calvo-style pricing rigidity: they are able to adjust their price only with a probability $1 - \theta_k$ in every period. This probability may vary across sectors. Firms in sector k operate the following technology:

$$y_{k,t}(j) = A_{k,t} F_k(n_{k,t}(j), \tilde{Y}_{1,k,t}(j), \tilde{Y}_{2,k,t}(j), ..., \tilde{Y}_{K,k,t}(j)), \tag{3}$$

where $y_{k,t}(j)$ is output, $F_k(\cdot)$ is a sector-specific production function with constant returns to scale and $A_{k,t}$ is an exogenous, sector-specific productivity variable. In the production function, $n_{k,t}(j)$ are effective units of labor hired by the firm, while $\tilde{Y}_{l,k,t}(j)$ is the quantity of intermediate inputs from sector $l = 1, 2, ..., K$ used in production by firm j in sector k. Intermediate goods are produced by competitive firms who bundle varieties and sell on the these bundles. The technology of these firms is given by $\tilde{Y}_{k,t} = \tilde{F}_k(\tilde{\mathbf{y}}_{k,t})$ where $\tilde{\mathbf{y}}_{k,t}$ is a vector of varieties used in production and where we assume that $\tilde{F}_k$ is twice differentiable, symmetric across varieties and has constant return to scale. We can express the demand of the intermediate goods producers for an individual variety j as $\tilde{y}_{k,t}(j) = \tilde{d}_k\left(p_{k,t}(j), \mathbf{p}_{k,t}\right)\tilde{Y}_{k,t}$.

Firms take as given the aggregate of household demand functions, as well as demand by intermediate goods producers. The total demand for a variety is given by:

$$y_{k,t}(j) = \int_0^1 d_k\left(p_{k,t}(j), \mathbf{p}_{k,t}, e_{k,t}(i)\right) di + \tilde{d}_k\left(p_{k,t}(j), \mathbf{p}_{k,t},\right)\tilde{Y}_{k,t}. \tag{4}$$

where the first term corresponds to household demand and the second to demand from intermediate goods producers. Under CES preferences, household demand for a variety can

[12] It is always possible to renormalize the utility function to obtain a common and arbitrary EIS and Frisch elasticity. Straub (2017) presents a model with EIS heterogeneity.

be expressed as a simple function of its relative price and total demand. In our more general setting, however, the composition of demand matters as well, as demand elasticities and super-elasticities vary across households.

Firms which are allowed to adjust their price do so to maximize the expected present value of profits. The decision problem of those firms is given by:

$$\max_{\substack{p_{k,t}(j),\{n_{k,t+s}(j),\\ y_{k,t+s}(j),\tilde{Y}_{l,k,t+s(j)}\}_{s=0}^{\infty}}} \mathbb{E}_t \sum_{s=0}^{\infty} \Lambda_{t,t+s}\theta_k^s \left(p_{k,t}(j)y_{k,t+s}(j) - (1-\tau_k)(W_{t+s}n_{k,t+s}(j) + \sum_l P_{l,t+s}\tilde{Y}_{l,k,t+s}(j)) - T_{k,t+s} \right), \tag{5}$$

subject to Equations (3) and (4), where $\Lambda_{t,t+s}$ is the firm's stochastic discount factor.[13] In the above equation, τ_k is a time-invariant, sector-specific subsidy which may be used by the government to correct markup distortions in the steady state, and $T_{k,t}$ a lump-sum tax to finance the subsidy, which can be arbitrarily differentiated across sectors, as long as the government budget constraints is satisfied.

Government Policy. We assume that the fiscal authority runs a balanced budget, which implies:

$$\sum_{k=1}^{K} \tau_k \int_0^1 (W_t n_{k,t}(j)) + \sum_{l=1}^{K} P_{l,t}\tilde{Y}_{l,k,t}(j))dj - \sum_k T_{k,t} = 0. \tag{6}$$

The nominal interest rate R_t is set by a central bank, taking fiscal policy as given. We will consider two versions of the model. In the first, the central bank follows a simple interest rate rule. In the second version, the interest rate is set optimally.

Demographics and Market Clearing. In any period, a fraction δ of all households dies. We assume that each deceased household is replaced by a new household of the same type. A household's type is pinned down by its labor productivity, $\vartheta(i)$, firm ownership, $\varsigma_k(i)$, initial bond holdings, $b_0(i)$, preferences, U_i, and HtM status. Bond market clearing implies that the average wealth of households is zero, and hence the same is true for deceased and newborn households, due to i.i.d. death probabilities. Therefore, the wealth given to newborn households can always be financed and the net inheritance from all deceased households is zero. From now on, we will assume that firm ownership is proportional to labor productivity. Clear-

[13]We assume that in the steady state $\Lambda_{t,t+s} = (1-\delta)^s\beta^s$. We do not need to make further assumptions on $\Lambda_{t,t+s}$ since we will linearize the model around a steady state with zero inflation.

ing in the labor market and the bond market requires, respectively:

$$\begin{aligned}\int_0^1 n_t(i)di &= \sum_k \int_0^1 n_{k,t}(j)dj, \\ \int_0^1 b_t(i)di &= 0.\end{aligned} \tag{7}$$

Goods market clearing requires, for any goods variety:

$$\int_0^1 c_{k,t}(i,j)di + \tilde{y}_{k,t}(j) = y_{k,t}(j). \tag{8}$$

and in every sector:

$$\tilde{Y}_{k,t} = \sum_l \int \tilde{Y}_{l,k,t}(j)dj. \tag{9}$$

An equilibrium is a law of motion for prices and allocations such that households, firms and the government behave as specified above, and markets clear. It is worth noting that in the deterministic steady state of the model, households keep their bond holdings constant over time.[14] The model is thus consistent with any arbitrary steady-state distribution of wealth, which in the calibration we will take from the data.

2.2 Dynamic Equilibrium

In order to study dynamics, we linearize the model around a deterministic steady state. We assume that the central bank targets long-run price stability, so steady-state prices are identical within sectors. We further assume that the government eliminates steady-state markup distortions using the subsidy τ_k.

We now present the system of equations that jointly characterize the dynamic equilibrium of the model, to a first-order approximation. Appendix A provides the underlying derivations, and Appendix B summarizes the equations. To ease the exposition, we present in the main text a simplified model version without HtM households and without Input-Output linkages. In the quantitative applications, we do include these features. Moreover, in Section 3 we will consider a version of the model that is further simplified and derive a number of analytical results which help to sharpen intuition.

New Keynesian Phillips Curve. The central equation in our analysis is the New Keynesian Phillips Curve (NKPC). Let $\hat{P}_{k,t} = \int \hat{p}_{k,t}(j)dj$ be the price of the sector-k goods, where hatted variables denote log deviations from the steady state and where we used that in the steady

[14] It can be shown that, in the absence of idiosyncratic income risk and aggregate shocks, the target level of wealth equals current wealth.

state prices are identical within sectors. We will denote steady-state variables by omitting the time subscript t. The steady-state interest rate equals $R = \frac{1}{\beta(1-\delta)}$. The net rate of inflation in sector k is given by:

$$\pi_{k,t} = \hat{P}_{k,t} - \hat{P}_{k,t-1}. \tag{10}$$

Moreover, individual consumption of sector-k goods is given by $\hat{c}_{k,t}(i) = \hat{e}_{k,t}(i) - \hat{P}_{k,t}$. The NKPC for sector k can be now expressed as:

$$\pi_{k,t} = \kappa_k \tilde{\mathcal{Y}}_t + \lambda_k \left(\mathcal{NH}_t + \mathcal{M}_{k,t} - \mathcal{P}_{k,t}\right) + \beta(1-\delta)\mathbb{E}_t \pi_{k,t+1}, \tag{11}$$

with:

$$\tilde{\mathcal{Y}}_t = \hat{\mathcal{Y}}_t - \hat{\mathcal{Y}}_t^*, \qquad \text{(Output gap)}$$

$$\mathcal{NH}_t = \sum_{l=1}^{K} (\overline{\partial_e e_l} - \bar{s}_l)(\hat{P}_{l,t} - \hat{P}_{l,t}^*), \qquad \text{(Non-homotheticity wedge)}$$

$$\mathcal{M}_{k,t} = \int \gamma_{e,k}(i) \frac{e_k(i)}{E_k} \hat{c}_{k,t}(i) di - \Gamma_k \tilde{\mathcal{Y}}_t, \qquad \text{(Endogenous markup wedge)}$$

$$\mathcal{P}_{k,t} = (\hat{P}_{k,t} - \hat{P}_{cpi,t}) - (\hat{P}_{k,t}^* - \hat{P}_{cpi,t}^*), \qquad \text{(Relative price wedge)}$$

and the following slope coefficients:

$$\kappa_k = \lambda_k \left(\frac{1}{\sigma} + \frac{1}{\psi}\right)\left(1 + \frac{\sigma\psi}{\sigma+\psi}\Gamma_k\right),$$

$$\lambda_k = \frac{(1-\theta_k)(1-\theta_k/R)}{\theta_k} \frac{\bar{\epsilon}_k - 1}{\bar{\epsilon}_k - 1 + \bar{\epsilon}_k^s},$$

$$\Gamma_k = \frac{R}{R-1} \frac{\sigma+\psi}{\sigma} \int \gamma_{b,k}(i) \frac{Wn(i)}{WN} di.$$

Before explaining our generalized NKPC in detail, let us note that it is a generalization of the 'standard' NKPC. As usual, the equation relates current sectoral rate of inflation, $\pi_{k,t}$, to the discounted expected rate of inflation, $\beta(1-\delta)\mathbb{E}_t\pi_{k,t+1}$, and an 'output gap', $\tilde{\mathcal{Y}}_t$.

In addition, a number of wedges emerge in the NKPC, which affect the joint dynamics of the output gap and inflation. The first of these, $\mathcal{NH}_t$, arises due to non-homothetic preferences over sectors, which makes the composition of consumption baskets vary across households and over time. The second, $\mathcal{M}_{k,t}$, arises due to changes in markups due to fluctuations in the price elasticities of demand faced by firms, which are no longer constant once one deviates from CES preferences over varieties within sectors. We label this wedge the *endogenous markup wedge*. The two new wedges will affect the trade-offs between the output gap and inflation, faced by the central bank. Finally, there is a relative price wedge $\mathcal{P}_{k,t}$ which generally arises in

New Keynesian models with sectoral asymmetries.

Slope of the NKPC. Let us now discuss the equation in more detail, starting with κ_k, the slope coefficient with respect to the output gap. The first term within this coefficient, λ_k, captures the micro-level pass-through of marginal costs to prices and in turn consists of two components. The first component within λ_k, i.e. $\frac{(1-\theta_k)(1-\theta_k/R)}{\theta_k}$, is due to sticky prices and is standard in the NK model. The second component, $\frac{\bar{\epsilon}_k-1}{\bar{\epsilon}_k-1+\bar{\epsilon}_k^s}$, arises because of the endogeneity of demand elasticities. Intuitively, a firm realises that if it raises its price, demand will fall and, as a result, consumers may become more price sensitive. This component does not appear under CES preferences ($\bar{\epsilon}_k^s = 0$), but it does appear under for instance Kimball (1995) preferences. In a typical calibration it holds that $\bar{\epsilon}_k^s > 0$, which implies that the pass-through from marginal costs to prices is less than one-for-one, even when prices are fully flexible.

The second term in the definition of κ_k, i.e. $\left(\frac{1}{\sigma}+\frac{1}{\psi}\right)$, is standard in the NK literature. The third term, $\left(1+\frac{\sigma\psi}{\sigma+\psi}\Gamma_k\right)$, is again due to non-CES preferences. However, this time it captures an aggregate spending effect: when household change their consumption levels, demand elasticities react, which induces firms to change markups. When markups tend to be increasing in wealth ($\gamma_{b,k}(i) > 0$) then an increase in aggregate income makes consumers less price sensitive, therefore pushing up markups. Again, the term vanishes under CES preferences.[15]

Note further that in the general setting, κ_k depends on the entire steady-state distribution of expenditures, through Γ_k and $\bar{\epsilon}_k^s$. Thus, long-run changes in inequality affect the slope of the NKPC. As such, our environment differs from typical HANK settings, in the sense that inequality affects not only the demand block of the model, as represented by consumption Euler equations and budget constraints, but also the supply block, as formed by the NKPCs.

Output gap. The first term on the right-hand side of the NKPC, is the output gap. Here, $\hat{\mathcal{Y}}_t$ is an aggregate demand index, and $\hat{\mathcal{Y}}_t^*$ is its 'natural' counterpart, indicated by a star and defined as its level in a parallel economy without markup distortions. As in the standard NK model, the output gap captures distortions in the labor market due to time-varying markups. To see this concretely, one can express the output gap alternatively as a (household) wage gap: $\tilde{\mathcal{Y}}_t = \frac{\psi}{1+\frac{\psi}{\sigma}}\left(\hat{w}_{h,t}-\hat{w}_{h,t}^*\right)$, where $\hat{w}_{h,t} = \hat{W}_t - \sum_{l=1}^K \overline{\partial_e e_l}\hat{P}_{l,t}$ is the real wage, computed using the Marginal CPI (MCPI) as the deflator, which is the relevant wage for marginal labor supply decisions. Moreover, $\hat{w}_{h,t}^* = \sum_{l=1}^K \overline{\partial_e e_l}\hat{A}_{l,t}$ is the natural counterpart of the real wage. This expression for the output gap also obtains in the standard NK model, in which the CPI and MCPI coincide.

In an economy with heterogeneous agents and multiple sectors there are in principle many

[15]It also vanishes under Kimball (1995) preferences, since such preferences are homothetic, in the sense that they are scaled to be invariant to total demand.

ways in which one could measure aggregate labor market distortions. However, in Section 6 and Appendix E.3 we show that above formulation of the output gap gives precisely the distortion that enters into a planner's social welfare objective (see Result 7 for a simple case).

Dynamically, the output gap index evolves according to the following Euler equation:

$$\tilde{\mathcal{Y}}_t = \mathbb{E}_t \tilde{\mathcal{Y}}_{t+1} - \sigma \mathbb{E}_t \left(\hat{R}_t - \pi_{mcpi,t+1} - \hat{r}_t^* \right). \tag{12}$$

This Euler equation has the standard form, except that the real interest rate is computed with $\pi_{mcpi,t} = \sum_{l=1}^{K} \overline{\partial_e e_l} \pi_{l,t}$, i.e. MCPI rate of inflation, rather than the regular CPI. Intuitively, when households decide on consumption today versus consumption tomorrow, they consider on which sectors they spend at the margin. In the Euler equation, $\hat{r}_t^*$ is the natural real interest rate associated with the demand index, i.e. the real interest rate that satisfies the Euler Equation for the natural level of aggregate demand. We can express this rate as:

$$\hat{r}_t^* = \frac{1}{\sigma + \psi} \sum_{l=1}^{K} \left(\psi \overline{\partial_e e_l} + \bar{s}_l \right) (\hat{A}_{l,t+1} - \hat{A}_{l,t}), \tag{13}$$

Moreover, we can express as the natural level of demand and the natural sectoral price as $\hat{\mathcal{Y}}_t^* = \sum_{l=1}^{K} \frac{\psi \overline{\partial_e e_l} + \bar{s}_l}{1 + \psi / \sigma} \hat{A}_{l,t}$ and $\hat{P}_{k,t}^* = -\hat{A}_{k,t}$, respectively.

Note that in the equation for the natural rate, both regular budget shares ($\bar{s}_l$) and the marginal budget shares ($\overline{\partial_e e_l}$) enter. Indeed, in this economy, both the regular CPI and the MCPI matter for aggregate demand. To clarify this point further, let us express the natural level of demand as $\hat{\mathcal{Y}}_t^* = -\frac{1}{1+\psi/\sigma} \hat{P}_{cpi,t}^* - \frac{\psi}{1+\psi/\sigma} \hat{P}_{mcpi,t}^*$, i.e. as a weighted sum of the natural CPI and MCPI. Intuitively, sectoral productivity shocks directly affect aggregate income by shifting the productive capacity of the economy. For this effect, the regular budget shares (i.e. CPI sectoral weights) are the relevant sectoral weights. Secondly, sectoral shocks have an indirect equilibrium effect on households' marginal saving and labor supply decisions. For these decisions, the marginal budget shares are the relevant sectoral weights.

Non-homotheticity wedge. We now discuss the two novel NKPC wedges. The first of these, $\mathcal{NH}_t = \sum_{l=1}^{K} (\overline{\partial_e e_l} - \bar{s}_l)(\hat{P}_{l,t} - \hat{P}_{l,t}^*)$, is a wedge which arises due to non-homothetic preferences combined with distortions in relative sectoral prices. This wedge increases when prices are distorted downward ($\hat{P}_{l,t} < \hat{P}_{l,t}^*$) in necessity sectors ($\overline{\partial_e e_l} < \bar{s}_l$), but falls when prices are distorted downward in luxury sectors. Indeed, the movements in this wedge will depend critically on the sectoral nature of shocks. Note that $\mathcal{NH}_t$ is the same for all sectors, since it derives from a distortion in the aggregate labor market. Note further that under homothetic preferences, marginal and regular budget shares coincide and hence $\mathcal{NH}_t = 0$. But under non-homothetic preferences, the wedge moves over time, shifting the NKPC. The direction and magnitude of this shift depends on the gap $\overline{\partial_e e_l} - \bar{s}_l$, which in turn depends on the extent of steady-state inequality.[16]

To understand the wedge, it is useful to consider an alternative formulation, given by $\mathcal{NH}_t = (\hat{w}_{f,t} - \hat{w}^*_{f,t}) - (\hat{w}_{h,t} - \hat{w}^*_{h,t})$. Here, $\hat{w}_{f,t} = \hat{W}_t - \hat{P}_{cpi,t}$ is the real wage according to the CPI, which is relevant to the marginal cost of the firm (weighted by sales), and $\hat{w}^*_{f,t} = \sum_{l=1}^{K} \bar{s}_l \hat{A}_{l,t}$ is its natural counterpart. Recall that $\hat{w}_{h,t} = \hat{W}_t - \sum_{l=1}^{K} \overline{\partial_e e_l} \hat{P}_{l,t}$ is the real wage according according to the MCPI deflator, which is relevant to households' marginal labor supply decisions, and $\hat{w}^*_{h,t} = \sum_{l=1}^{K} \overline{\partial_e e_l} \hat{A}_{l,t}$ is its natural counterpart. We now observe that $\mathcal{NH}_t$ can be interpreted as a term capturing the extent to which real wage distortions differ between households and firms. As such, $\mathcal{NH}_t$ can be interpreted as a labor wedge, akin to a labor income tax distortion.

To obtain further intuition, note that in an economy with non-homothetic preferences, labor supply optimally responds to changes in relative sectoral productivities, even if aggregate productivity (i.e. sales-weighted sectoral productivity) does not change. Intuitively, when the relative productivity of luxury sectors increases, and relative prices in these sectors fall, households optimally increase labor supply since at the margin they spend relatively more on luxuries. To see this concretely, note that when CPI weighted aggregate productivity does not move, then $\hat{\mathcal{Y}}^*_t = -\frac{\psi}{1+\psi/\sigma} \hat{P}^*_{mcpi,t}$. Given this, any increase in the relative productivity of luxury sectors means that the natural MCPI declines, which leads to an increase in labor supply, increasing the natural level of output. However, when prices are sticky, the relative price movements are muted, and as a result $\mathcal{Y}_t$ increases by less than its natural counterpart, i.e. the output gap becomes negative.

Endogenous markup wedge. The second novel wedge, $\mathcal{M}_{k,t}$, captures the evolution of the distribution of price elasticities of demand for individual goods varieties, which affects the markups set by firms. The distribution of demand elasticities in turn fluctuates with the distribution of expenditures. The distributional origins of the wedge become clear by observing the first term in its definition, $\int \gamma_{e,k}(i) \frac{e_k(i)}{E_k} \hat{c}_{k,t}(i) di$, which integrates over individual households. Here $\hat{c}_{k,t}(i)$ is the consumption change of household i, $\frac{e_k(i)}{E_k}$ is the household's share in total sectoral consumption, and $\gamma_{e,k}(i)$ the sensitivity of markups with respect to individual expenditures. The second term, $-\Gamma_k \tilde{\mathcal{Y}}_t$, subtracts the endogenous markup response due to fluctuations in the output gap, as this effect has been subsumed in the NKPC slope κ_k.

The endogenous markup wedge arises due to non-CES utility over varieties within sectors.[17] To see this, note that under CES preference we obtain $\gamma_{b,k}(i) = \Gamma_k = 0$, as demand elasticities are constant, which in turn implies that $\mathcal{M}_{k,t} = 0$. Moving beyond CES, the wedge takes the same form as exogenous markup shocks often considered in New Keynesian models. However, in our setting it is a rich endogenous object, which is shaped by the distribution of

[16] Under non-homothetic preferences, budget shares are non-linear functions total expenditures, hence a long-run change in inequality will generally change the gap between marginal and regular budget shares.

[17] Note that preferences may be homothetic but non-CES and vice versa.

expenditures across households, and it therefore moves over time, along with the distribution of wealth. Nonetheless, it turns out that the evolution of the endogenous markup wedge can be represented in a tractable way. Specifically, it can be decomposed as:

$$\mathcal{M}_{k,t} = \Gamma_k \hat{\mathcal{Y}}_t^* + \mathcal{M}_{k,t}^P + \mathcal{M}_{k,t}^D. \tag{14}$$

The first component, $\Gamma_k \hat{\mathcal{Y}}_t^*$, arises due to changes in demand elasticities in response to changes in the natural level of aggregate demand. Intuitively, when $\Gamma_k > 0$ then during an economic downturn households cut expenditures and become more price-sensitive, which induces firms to reduce price markups.

The second component captures how substitutions in response to changes in prices in other sectors affect demand elasticities:

$$\mathcal{M}_{k,t}^P = \sum_{l=1}^{K} \mathcal{S}_{k,l} \cdot \left(\hat{P}_{l,t} - \hat{P}_{k,t}\right), \tag{15}$$

where $\mathcal{S}_{k,l} = \int_i \frac{e_k(i)}{E_k} \gamma_{e,k}(i) \rho_{k,l}(i) di$ captures the effect of cross-price substitution on demand elasticities, and hence markups.

The third component, $\mathcal{M}_{k,t}^D$, summarizes the effects of changes in the distribution of household-level real expenditures on markups. For instance, a redistribution from poor to rich households may give rise to an increase in markups, if rich people are less price sensitive. The evolution of $\mathcal{M}_{k,t}^D$ can be characterized by the following equation:

$$\mathcal{M}_{k,t}^D = \mathbb{E}_t \mathcal{M}_{k,t+1}^D - \sum_{l=1}^{K} \sigma_{k,l}^{\mathcal{M}} \left(\hat{R}_t - \mathbb{E}_t \pi_{l,t+1}\right) - \frac{\delta}{1-\delta} \mathbb{E}_t \mathcal{M}_{k,t+1}^0, \tag{16}$$

for any sector k, where $\sigma_{k,l}^{\mathcal{M}} = \sigma \int \gamma_{e,k}(i) \frac{e(i)}{E_k} \partial_e e_k(i) \partial_e e_l(i) di - \sigma \overline{\partial_e e_l} \Gamma_k$. In Equation (16), $\mathcal{M}_{k,t+1}^0$ captures the dynamics of the wealth distribution, insofar relevant for the markup wedge. It is pinned down by the following equation:

$$\begin{aligned} \mathcal{M}_{k,t}^0 = & \frac{1}{(1-\delta) R} \mathbb{E}_t \mathcal{M}_{k,t+1}^0 + \int \gamma_{b,k}(i) \frac{b(i)}{RE} di \left(\hat{R}_t - \mathbb{E}_t \pi_{cpi,t+1}\right) \\ & - \sum_{l=1}^{K} \int \gamma_{b,k}(i) \left(\frac{e(i)}{E} \left(s_l(i) - \bar{s}_l\right) + \frac{\psi W n(i)}{WN} \left(\partial_e e_l(i) - \overline{\partial_e e_l}\right) \right) di \hat{P}_{l,t} - \frac{R-1}{R} \mathcal{M}_{k,t}^D. \end{aligned} \tag{17}$$

Here, the second term on the right hand side captures markup effects due to changes in real interest rates. Intuitively, an increase in interest rates may redistribute wealth towards richer households, who have lower price elasticities of demand. This shifts the composition of demand towards households who are less price sensitive. In response, firms increase markups. Similarly, the third term captures the markup effects of redistributions due to changes in relative prices.

Relative price wedge. The final wedge in the NKPC, $\mathcal{P}_{k,t} = (\hat{P}_{k,t} - \hat{P}_{cpi,t}) - (\hat{P}^*_{k,t} - \hat{P}^*_{cpi,t})$ arises due to distortions in relative sectoral prices. Specifically, $\hat{P}_{k,t} - \hat{P}_{cpi,t}$ is the sectoral price, relative to the CPI and $\hat{P}^*_{k,t} - \hat{P}^*_{cpi,t}$ is its natural equivalent. The wedge $\mathcal{P}_{k,t}$ is generally present in multi-sector extensions of the standard NK model, if sectors are asymmetric in some way, e.g. if they differ in the degree of price rigidity or if there are sectoral shocks.

Monetary policy. In the positive part of our analysis, we will consider a simple interest rate rule of the following form:

$$\hat{R}_t = \sum_k \phi_k \pi_{k,t}, \tag{18}$$

where setting $\phi_k = \phi \bar{s}_k$ delivers a rule which responds to the CPI inflation rate. In Section 5, we will move beyond the simple rule and instead consider the fully optimal Ramsey policy.

Dynamic Equilibrium. Equations (10)-(18) constitute a system of $5K + 3$ equations in $5K + 3$ endogenous variables, given by $\{\hat{P}_t, \pi_{k,t}, \mathcal{M}^D_{k,t}, \mathcal{M}^P_{k,t}, \mathcal{M}^0_{k,t}\}_{k=1}^K, \tilde{\mathcal{Y}}_t, \hat{R}_t, r^*_t$. We can thus characterize the model with a core block of equations, despite the fact that fluctuations in the distribution of income and wealth matter for the aggregate equilibrium outcomes. The equations for $\mathcal{M}^D_{k,t}$ and $\mathcal{M}^0_{k,t}$ keep track of the relevant distributional moments in a tractable way.

Distributional dynamics. While we do not need to keep track of the full distributional dynamics in order to solve for the aggregate equilibrium, it is straightforward to solve for such dynamics. Here, we focus on the distribution of consumption. Let us define the response of real consumption expenditures of household i as $\hat{c}_t(i) = \hat{e}_t(i) - \sum_{l=1}^K s_l(i)\widehat{P}_{l,t}$. Moreover, let $\boldsymbol{\omega}$ be a vector defining a weight $\omega(i)$ on each household i, with $\int \omega(i)di = 1$. We can thus use $\boldsymbol{\omega}$ to select and weight any arbitrary subset of households.

Now consider some moment of the consumption distribution, $\hat{C}_t(\boldsymbol{\omega}) = \int \omega(i)\hat{c}_t(i)di$. For instance, if we set $\omega(i) = e(i)/E$, then this moment corresponds to the aggregate response of real expenditures. We could also set $\omega(i) = 1$ for only one specific household i and zero for all others. In that case, $\hat{C}_t(\boldsymbol{\omega})$ corresponds to the individual consumption response of a particular household. Alternatively, one can choose $\boldsymbol{\omega}$ to compute the average response among households with certain characteristics. We can characterize $\hat{C}_t(\boldsymbol{\omega})$ with the following Euler equation:

$$\mathbb{E}_t\hat{C}_{t+1}(\boldsymbol{\omega}) - \hat{C}_t(\boldsymbol{\omega}) = \sigma\left(\int \omega(i)di\hat{R}_t - \textstyle\sum_k \int \omega(i)\partial_e e_k(i)di\, \mathbb{E}_t\pi_{k,t+1}\right) + \frac{\delta}{1-\delta}\hat{C}^0_t(\boldsymbol{\omega}), \tag{19}$$

where wealth dynamics are captured by:

$$\hat{C}_t^0(\boldsymbol{\omega}) - \frac{1}{(1-\delta)\,R}\mathbb{E}_t\hat{C}_{t+1}^0(\boldsymbol{\omega}) = \int \omega^0(i)\frac{b(i)}{RE}di\left(\hat{R}_t - \mathbb{E}_t\Sigma_k\bar{s}_k\pi_{k,t+1}\right) + \left(1+\frac{\psi}{\sigma}\right)\int \omega^0(i)\frac{Wn(i)}{WN}di\hat{\mathcal{Y}}_t$$
$$-\Sigma_k\int \omega^0(i)\left(\frac{e(i)}{E}\left(s_k(i)-\bar{s}_k\right)+\frac{Wn(i)}{WN}\psi\left(\partial_e e_k(i)-\overline{\partial_e e_k}\right)\right)di\hat{P}_{k,t} - \frac{R-1}{R}\hat{C}_t(\boldsymbol{\omega}), \tag{20}$$

where we defined $\omega^0(i) = \frac{R-1}{R}\frac{\omega(i)}{e(i)/E+Wn(i)/WN\frac{\psi}{\sigma}}$.

3 Understanding the NKPC Wedges

Before studying the model quantitatively and deriving the optimal policy, we present a number of analytical results which help to understand how the wedges respond to aggregate and sectoral shocks, how they affect aggregate dynamics, and to what extent it is possible for policy to neutralize the distortions they create. In order to derive these results, we consider two simplifying assumptions:

Assumptions:
(**A.1**) The slope of the NKPC with respect to the output gap is homogeneous across sectors, i.e. $\kappa_k = \kappa > 0$ for any sector k.
(**A.2**) There is no steady-state wealth heterogeneity, i.e. $b(i) = 0$ for any household i.[18]

We impose these assumptions throughout this section in order to understand the key novel mechanisms, but we will dispose of them the quantitative analysis. Indeed the literature has shown the quantitative importance of sectoral heterogeneity in NKPC slopes (see e.g. Pasten et al. (2020); Rubbo (2023)) and of wealth heterogeneity (see e.g. Kaplan et al. (2017)). Under the assumptions, we can derive a number of results. The proofs of these are provided in Appendix C.

Result 1 (policy invariance of the sectoral wedges): *Under (**A.1**)-(**A.2**), $\mathcal{NH}_t$, $\mathcal{M}_{k,t}$ and $\mathcal{P}_{k,t}$ evolve independently of monetary policy.*

The key insight behind our first analytical result is that all three wedges can be expressed as functions of relative sectoral prices and relative nominal wealth positions only. When the slope of the NKPC with respect to the output gap is homogeneous across sectors and there is no initial

[18] As above, we also abstract from Input-Output linkages and Hand-to-Mouth agents. Note that we do allow for income heterogeneity and for endogenous wealth heterogeneity in response to shocks. Moreover, there is still wealth inequality out of steady state. Finally, all the results go trough under a generalized assumption $\int \gamma_{b,k}(i)b(i)di = 0$, i.e. what matters is that wealth positions are orthogonal to markup contributions.

nominal wealth heterogeneity, the central bank has no levers to move these relative outcomes, and hence the wedges become invariant to monetary policy. The wedges then become similar to exogenous markup shocks often introduced to NK models, but with potentially richer dynamics propagated through movements in the wealth distribution.

In the full model, assumptions (**A.1**)-(**A.2**) do not apply. It then becomes possible for policy to affect the wedges, but only via two specific channels: relative sectoral prices and nominal redistributions. Thus, even if a central bank's mandate refers only to aggregate inflation and the output gap, wealth heterogeneity and movements in sectoral prices become intermediate targets for policy.

3.1 The role of the $\mathcal{NH}$ wedge

Let us now explore the wedges in more detail, starting with the non-homotheticity wedge, $\mathcal{NH}$. In order to focus exclusively on this wedge, let us assume, in addition to (**A.1**)-(**A.2**), that $\mathcal{M}_t = 0$, i.e. preferences are are homothetic over sectors. We do preserve the other wedges, i.e. $\mathcal{NH}_t \neq 0$ and $\mathcal{P}_{k,t} \neq 0$. We can now derive our second analytical result, highlighting the relevance of the MCPI index:

Result 2 (divine coincidence without endogenous markup wedge): *If (**A.1**)-(**A.2**) hold and $\mathcal{M}_t = 0$, then fluctuations in the output gap can be eliminated by stabilising Marginal CPI inflation, defined as $\pi_{mcpi,t} \equiv \sum_{k=1}^{K} \overline{\partial_e e_k} \pi_{k,t}$.*[19]

We thus recover a version of the "Divine Coincidence" often emphasized in the NK literature. But rather than stabilising the CPI, policy should stabilise the *Marginal* CPI in order to eliminate fluctuations in the output gap. This result follows from the aggregate NKPC, weighted by the Marginal CPI shares, which under the assumptions reduces to:

$$\pi_{mcpi,t} = \kappa \tilde{\mathcal{Y}}_t + \beta(1-\delta)\mathbb{E}_t \pi_{mcpi,t+1}. \tag{21}$$

Note that all remaining wedges drop out of this equation. It follows immediately that when $\pi_{mcpi,t} = 0$ at all times, then $\tilde{\mathcal{Y}}_t = 0$.[20]

The MCPI thus emerges as a natural candidate to be a target for policy. In fact, in this simplified setting the model becomes isomorphic to the standard 3-equation NK model if policy targets MCPI rather than CPI inflation. To see this, suppose that policy follows a simple interest rate rule targeting the MCPI:

$$\hat{R}_t = \phi \pi_{mcpi,t}.$$

[19] When we relax assumption (**A.1**) and (**A.2**), the divine coincidence index is $\pi_{d,t} \equiv \sum_k \frac{\overline{\partial_e e_k}/\lambda_k}{\sum_l \overline{\partial_e e_l}/\lambda_l} \pi_{k,t}$, the divine coincidence index with I-O linkages is derived in appendix F.

[20] Recall that in this section we abstract from I-O linkages. Rubbo (2023) derives a divine-coincidence index for an NK model with such links.

Illustration: responses to negative sectoral productivity shocks.

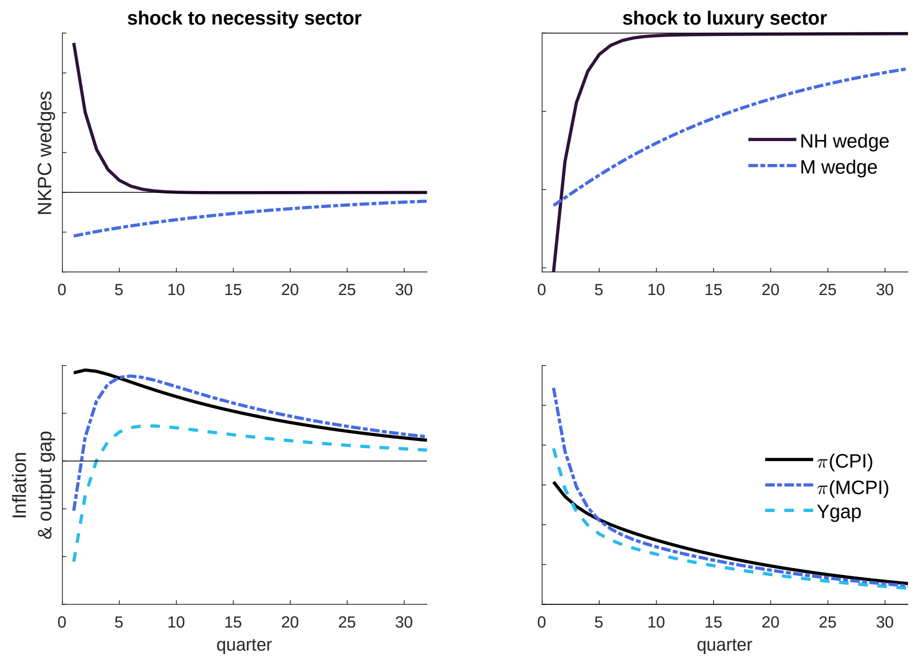

Notes: responses to a negative productivity shock to a necessity sector (left panels) and to a luxury sector (right panels). Simplified version satisfying assumptions (**A.1**)-(**A.2**).

Together with Equation (12), the above two equations form a 3-equation system which take the exact same form as the standard NK model, but with MCPI inflation.

Yet, even when the output gap and MCPI inflation are fully stabilized, there are still fluctuations in the CPI. To see this clearly, consider the NKPC for CPI inflation:

$$\pi_{cpi,t} = \kappa \tilde{\mathcal{Y}}_t + \lambda \mathcal{N}\mathcal{H}_t + \beta(1-\delta)\mathbb{E}_t \pi_{cpi,t+1}. \tag{22}$$

Thus, due to fluctuations in the $\mathcal{N}\mathcal{H}$ wedge, there is a policy trade-off between the output gap and the regular CPI inflation index. Put differently, if monetary policy wishes to neutralise labor market distortions, it must accept fluctuations in CPI inflation. The trade-off between CPI inflation and the output gap depends critically on the sectoral nature of the shock, since the wedge moves in different directions in response to different sectoral shocks:

Result 3 (response of $\mathcal{N}\mathcal{H}$ to a sectoral shocks): *If (**A.1**)-(**A.2**) hold and $\mathcal{M}_t = 0$ then, following a negative productivity shock to a necessity (luxury) sector, $\mathcal{N}\mathcal{H}_t$ rises (falls) on impact.*

The intuition for this result was explained in Section 2.2. In response to a fall in productivity in a necessity sector, prices rise in that sector due to an increase in marginal costs. However, price stickiness prevents the relative sectoral price from rising as much as in the undistorted case, and therefore $\hat{P}_{l,t} < \hat{P}^*_{l,t}$. That is, prices in the necessity sector are distorted *downward*. Since households consume less necessities at the margin than on average, i.e. $\overline{\partial_e e_l} < \bar{s}_l$, this creates an increase in $\mathcal{NH}_t = \sum_l (\overline{\partial_e e_l} - \bar{s}_l)(\hat{P}_{l,t} - \hat{P}^*_{l,t})$.

The implications of the shift in the $\mathcal{NH}_t$ can be further understood by considering an extreme monetary policy which strictly targets the CPI, i.e. $\pi_{cpi,t} = \pi_{cpi,t+1} = 0$. It then follows immediately from Result 3 and Equation (22) that a negative productivity shock to a necessity sector results in a negative output gap.[21] Policy could neutralize this effect by stabilising instead MCPI inflation and the output gap, i.e. by targeting $\pi_{mcpi,t} = \pi_{mcpi,t+1} = \tilde{\mathcal{Y}}_t = 0$, but this would would come at the cost an increase in CPI inflation. Analogously, a negative shock to a luxury sector would *reduce* CPI inflation under this policy.

The response of the output gap thus depends critically on (i) the sectoral nature of the shock (luxury vs necessity), and (ii) the inflation index targeted by the central bank. This remains the case once we consider less extreme policies. To show this, let us first consider an MCPI-based rule $\hat{R}_t = \phi\pi_{mcpi,t}$. In this case, the output gap responds to a change in the natural real interest rate exactly as in the standard 3-equation NK model, as the model is isomorphic. Indeed, following a negative productivity shock, the output gap will increase since the natural rate r^*_t increases, regardless of the sectoral nature of the shock.

However, under a CPI-based rule of the form $\hat{R}_t = \phi\pi_{cpi,t}$, the output gap may actually *decline* following a negative necessity shock. To see this, let us rewrite this rule as $\hat{R}_t = \phi\pi_{mcpi,t} + u^R_t$, where $u^R_t = \phi(\pi_{cpi,t} - \pi_{mcpi,t})$. Together with Equations (12) and (21), we obtain a system that is isomorphic to the standard 3-equation NK model but with an additional, endogenous monetary policy shock, u^R_t. Following a negative shock to a necessity sector, CPI inflation increases by more than MCPI inflation, i.e. u^R_t increases, creating an effect akin to a a monetary contraction, pushing down the output gap. If this additional effect is strong enough, the output gap becomes negative. Intuitively, the CPI overweights necessity sectors relative to the MCPI. Therefore, if the central bank targets the CPI, it increases the interest rate by 'too much' when a negative shock to necessities increases prices in that sector. Following a negative shock to a luxury sector, the opposite effect occurs, i.e. there is an additional expansionary effect.[22]

The figure above illustrates the insights so far, by showing impulse response functions for a simplified version of the model in which (**A.1**)-(**A.2**) apply and the central bank follows a CPI-based rule. Following a negative productivity shock to a necessity sector (left panels), the $\mathcal{NH}$ wedge rises and the output gap falls, whereas CPI inflation rises. With tighter monetary policy,

[21] Intuitively, as explained above, following a negative productivity shocks to necessities, the *relative* price of luxuries falls. This induces households to optimally increase labor supply, since households spend relatively more on luxuries at the margin. However, price rigidities dampen the increase in the relative price of luxuries, hence $\mathcal{NH}_t$rises. Therefore, labor supply is pushed up by less than is optimal, i.e. the output gap falls. Thus, following a negative shock to a necessity (luxury) sector, the non-homotheticity wedge shifts the NKPC upwards (downwards).

[22] The endogenous monetary policy shock is closely related to the non-homotheticity wedge. When λ_k is also homogeneous across sectors, we can express it as $u^R_t = \lambda \sum_{s\geq 0} \frac{1}{R^s} \mathcal{NH}_{t+s}$.

CPI inflation could be reduced, but this would be at the expense of a more negative output gap, i.e. a trade-off arises. By contrast, following a negative shock to luxuries (right panels), both the output gap and CPI inflation increase. In this case, a tightening of policy could bring down both. The figure also shows that MCPI inflation can move rather differently from CPI inflation, and that the former tends to co-move more closely with the output gap. In Appendix C, we provide analytical solutions of the model under simple interest rate rules and the simplifying assumptions.

3.2 The role of the $\mathcal{M}$ wedge

Let us now consider movements in the endogenous markup wedge, $\mathcal{M}$. We can show that the Divine Coincidence breaks down once this wedge is active:

Result 4 (breakdown divine coincidence): *When $\mathcal{M}_t \neq 0$, there generally does not exist an inflation index which can be fully stabilised along with the output gap regardless of the shocks.*

Intuitively, movements the endogenous markup wedge derive from real sources (fluctuations in demand elasticities), which cannot be neutralized with a nominal instrument.

How does the endogenous markup wedge move in respond to shocks? Let us start with an aggregate shock:

Result 5 (dynamics of the endogenous markup wedge): *If (**A.1**)-(**A.2**) hold and $\lambda_k = \lambda\ \forall k$ then $\mathcal{M}_t$ declines following a negative aggregate productivity shock.*

To understand this, it is useful to recall Equation (14) which decomposes the wedge as $\mathcal{M}_{k,t} = \Gamma_k \hat{\mathcal{Y}}_t^* + \mathcal{M}_{k,t}^P + \mathcal{M}_{k,t}^D$. Under the simplifying assumptions, only the first component, $\Gamma_k \hat{\mathcal{Y}}_t^*$ moves in response to aggregate productivity shock, and thus the decline in $\mathcal{M}_t$ is entirely driven by a fall in efficient output. Intuitively, a fall in income creates a decline in aggregate demand, which makes households become more price sensitive, and therefore reduces markups.

Following *sectoral* shocks, the sign of $\mathcal{M}_{k,t}$ is in principle ambiguous, since such shocks bring about relative price changes and redistributions, so that $\mathcal{M}_{k,t}^P$ and $\mathcal{M}_{k,t}^D$ move as well. In other words, the movements in the endogenous markup wedge generally depend on the sectoral source of the shock. To illustrate this point, let us make a further simplifying assumption:

Assumption:
(**A.3**) Outer preferences are of the Stone-Geary form, the superelasticity of the sectoral markup $\gamma_{e,k}(i)\partial_e e_k(i)/E_k$ is equal across sectors, and $\gamma_{e,k}$ is positive and increasing in $e_k(i)$.

We can now derive our final result about the CPI aggregates $\mathcal{M}_{cpi,t}^P = \sum_l \bar{s}_k \mathcal{M}_{k,t}^P$ and $\mathcal{M}_{cpi,t}^D = \sum_l \bar{s}_k \mathcal{M}_{k,t}^D$:

Result 6 (dynamics of the endogenous markup wedge): *If (**A.1**)-(**A.3**) hold, then* $\mathcal{M}^P_{cpi,t}$ *decreases (increases) and* $\mathcal{M}^D_{cpi,t}$ *increases (decreases) following a negative productivity shock to a necessity (luxury) sector.*

Under Stone-Geary preferences, the cross price elasticity of demand is given by $\rho_{k,l}(i) = \overline{\partial_e e_l}\,(1 - \underline{c}_k/c_k(i))$. Thus, expenditure switching in response to necessity price changes is relatively low, since the marginal budget share $\overline{\partial_e e_l}$ is low for necessities. Following a negative necessity shock, the substitution towards luxury goods is therefore relatively weak. As a result, expenditures and markups decline in the necessity sector, but this is not fully compensated by an increase in markups in the luxury sector. Therefore, $\mathcal{M}^P_{cpi,t}$ decreases. Moreover, a negative necessity shock disproportionately reduces the spending power of the poor, i.e. there is a relative redistribution towards the rich. Because the rich are are less price sensitive than the poor, the redistribution puts upward pressure on markups, i.e. $\mathcal{M}^D_{cpi,t}$ increases.[24]

The dynamics of the endogenous markup wedge are illustrated in the figure above. Note that the $\mathcal{M}$ wedge declines following both shocks, as the aggregate demand component $\Gamma_k \hat{\mathcal{Y}}^*_t$ dominates in this illustration. Note further that the decline is relatively modest but very persistent, which drives of the upswing in the output gap several quarters after the necessity shock hits, as well as the persistent increase in the output gap following a negative shock to the luxury sector.

4 Quantitative Analysis

The analytical results presented in the previous section show how shifts in the NKPC, and hence policy trade-offs, can depend critically on the sectoral source of the shock. Our next goal is to study quantitatively the effects of productivity shocks to different sectors. To this end, we revert back to the full model, in which the slope of the NKPC may vary across sectors, there is steady-state heterogeneity in nominal wealth, some households are Hand-to-Mouth and there are Input-Output linkages across sectors. We consider the model with an interest rate rule, targeting CPI inflation. In Section 6, we will consider optimal monetary policy.

4.1 Parameterization

We calibrate the model to the United Kingdom. The model period is set to one quarter. Parameter values are displayed in Tables 3 and 4, and are discussed below in detail. We include eight

[23]Here, $\underline{c}_k$ is the subsistence of sector-k consumption. Under (**A.3**), we have $\mathcal{M}^P_{cpi,t} = \sum_k \int \frac{e_k(i) - p_k \underline{c}_k}{E} \gamma_{e,k}(i) di \sum_l \left(\overline{\partial_e e_l} - \bar{s}_l\right) \hat{P}_{l,t}$

[24]Note that we assumed that $\gamma_{e,k}$ is increasing in expenditures. It may be decreasing when demand elasticities of rich households are relatively insensitive to changes in expenditures, compared to the poor. In that case, a redistribution towards the rich may increase the aggregate demand elasticity, as the compositional effect is overturned.

COICOP sectors in the model: *Food*, *Clothing*, *Electricity & Gas*, *Furniture*, *Transport*, *Recreation*, *Restaurants & Hotels*, and *Miscellaneous*.

Income and wealth distribution. An advantage of the model is that its steady state can be disciplined directly by feeding in observed distributions. To this end, we rely on the Living Costs and Food (LCF) survey, which collects detailed survey data for more than 5600 households in the UK.[25] We think of each household in the survey as a type and we use population weights from the LCF for aggregation.[26]

We construct nominal wealth, $b(i)$, as nominal savings minus mortgage and credit card debt.[27] Total expenditures, $e(i)$, and budget shares by sector, $s_k(i)$, are directly observed in the LCF survey. To ensure consistency with the model, we back out labor income $Wn(i)$ as a residual from the budget constraint.[28] Note that we do not explicitly recover individual labour productivities $\vartheta(i)$, however they are not needed since the sufficient statistics are provided by the labour income share $\frac{Wn(i)}{WN}$ and the Frisch elasticity ψ.

Preferences. We set $\delta = 0.0083$, targeting an adult life expectancy of 60 years. We set $\beta = 0.995$ which implies $R = \frac{1}{(1-\delta)\beta} = 1.0134$ on a quarterly basis. We further set $\psi = \sigma = 1$, in line with conventions in the macroeconomics literature.

Outer utility. In the LCF survey, we directly observe households expenditures on different goods from which we construct the household budget shares for each sector denoted by $s_k(i)$. To recover the marginal budget shares $\partial_e e_k(i)$ and the substitution matrix $\rho_{k,l}(i)$, which are not directly observed, we impose a functional form on the outer utility function and estimate it from the LCF data. Specifically, we parametrize $U_i(\cdot)$ following Comin et al. (2021), who propose a class of non-homothetic CES preferences defined implicitly by:[29]

$$\sum_{k=1}^{K} \mathcal{V}_k(i) \left(\frac{c_k(i)}{g(U(i))^{\zeta_k}} \right)^{\frac{\eta-1}{\eta}} = 1,$$

[25] The UK consumer price index produced by the Office for National Statistics (ONS) is based on expenditure baskets observed in the LCF survey.

[26] We use 2019 data to calibrate the model have 5695 household observations. We think of each of these households as a representative for a particular type. In this sense, our model has about 5695 types of households, with demographic turnover within each type, as households are replaced by steady-state versions of their type at a rate δ.

[27] In the LCF we observe interest income. We convert this into the stock of saving by assuming an interest rate of 1 percent annually. Moreover, to be consistent with zero bond holdings on average, we subtract average wealth for each household.

[28] Note that in the model's steady state, household savings, $b(i)$, are constant at the household level and dividends are zero, so total expenditure equals labor income plus interest income for each household i. In a few cases, implied labor income is negative. We then set labor income to zero and expenditures to asset income.

[29] In our model, the inner layer of preferences contains deviations from CES.

where η is the elasticity of substitution across sectors, ζ_k captures non-homotheticicities in consumption and $\mathcal{V}_k(i)$ are household-specific preference shifters.

As shown by Comin et al. (2021), the non-homothetic CES form implies the following expression for household i's budget share in sector k (relative to some baseline sector $\bar{k}$ whose non-homotheticity parameter $\zeta_{\bar{k}}$ has been normalized to 1):

$$\ln\left(s_k(i)\right) = (1-\eta)\ln\left(\frac{p_k}{p_{\bar{k}}}\right) + (1-\eta)\left(\zeta_k - 1\right)\ln\left(\frac{e(i)}{p_{\bar{k}}}\right) + \zeta_k \ln\left(s_{\bar{k}}(i)\right) + \eta \ln\left(\frac{\mathcal{V}_k(i)}{\mathcal{V}_{\bar{k}}(i)^{\zeta_k}}\right).$$

This class of preferences thus allows the sectoral composition of the consumption basket to vary with total expenditures. In particular, sectors that are more of a luxury than the base sector $\bar{k}$ will have a non-homotheticity parameter ζ_k that is larger than one (as long as $\eta < 1$) and the opposite is true for necessity sectors. In the limit, where the ζ's are all equal across sectors we are back to the homothetic CES case.

We model the household-level preference shifters as $\ln \mathcal{V}_k(i) = \beta_k x(i) + v_k(i)$, where $x(i)$ is a vector of demographic characteristics, such as age or couple status, and $v_k(i)$ captures remaining idiosyncratic preference variation. The latter allows to match the model in steady state precisely to the actual distribution of budget shares observed in the LCF data.

We set the elasticity of substitution between sectors as $\eta = 0.1$ and estimate the ζ_k parameters using a GMM procedure, following Comin et al. (2021) but using household-level data. In Appendix D we show that this specification gives a good fit of the empirical relation between expenditures and budget shares, a key object in our model.[30] Nonetheless, even for the same demographic group and expenditure level, there is still considerable variation in budget shares that is driven by the idiosyncratic shifters $v_k(i)$. In Appendix D, we provide details on the estimation. With the estimated equations at hand, we can compute for each household the implied marginal budget shares $\partial_e e_k(i)$, for each sector k, see Appendix B for the formula.

Figure 2 plots histograms of the distribution of the budget shares and marginal budget shares. In necessity sectors such as *Food* and *Electricity & Gas*, budget shares are decreasing in total expenditures and exceed marginal budget shares. In luxury sectors, such as *Recreation* and *Restaurants & Hotels*, the opposite is true. Table 4 shows the marginal budget shares, averaged across households, $\overline{\partial_e e_k}$ along with the average budget share $\bar{s}_k$, as well as the difference $\overline{\partial_e e_k} - \bar{s}_k$, which matters directly for the $\mathcal{NH}$ wedge.

Inner utility. The distributions of demand elasticities within sectors are not directly observed in the data. However, they do have implications, which we can exploit to impose empirical discipline. Specifically, we assume a HARA form for the inner utility function, $\mathcal{U}_k(\cdot)$, which

[30] The value of η is based on the 10-sector estimation in Comin et al. (2021), Table XII. Our has a relatively short time dimension and, related to this, η does not to appear to be very sharply identified. That said, specifications with low values for η tend to fit the data relatively well.

implies that the elasticity of substitution between goods in sector k, for household i, is then given by:

$$\epsilon_k(i) = a_k + \frac{b_k}{e_k(i)},$$

where $a_k > 0$ and b_k are sector-level constants.[31] When $b_k > 0$, households become less price sensitive as they spend more and it then holds that $\gamma_{e,k}(i) > 0$. It can be shown that the sector-level demand elasticity and super-elasticity are given by, respectively, $\bar{\epsilon}_k = a_k + \frac{b_k}{E_k}$ and $\bar{\epsilon}_k^s = \frac{b_k}{E_k}$. We further assume that intermediate input demand is governed by the same elasticity and superelasticity. Given these objects we can compute the steady-markup at the sector level as $\frac{\bar{\epsilon}_k}{\bar{\epsilon}_k - 1}$ and the long-run pass-through of marginal costs to prices as $\frac{\bar{\epsilon}_k - 1}{\bar{\epsilon}_k - 1 + \bar{\epsilon}_k^s}$. We calibrate a_k and b_k by targeting sector-level markup estimates produced by the Office for National Statistics, following the method of De Loecker and Warzynski (2012). Moreover, we target 70 percent pass-through (in all sectors), based on empirical evidence by Amiti et al. (2019). Table 4 presents the implied sector-level coefficients. Given a_k and b_k and the empirical distribution of expenditures at the sector level, $e_k(i)$, we can compute the distributions of individual demand elasticities, $\epsilon_k(i)$, and super-elasticities, $\epsilon_k^s(i)$, which also gives us $\gamma_{e,k}(i)$ and $\gamma_{b,k}(i)$. Expressions for all relevant objects are provided in Appendix B.

Hand-to-Mouth households. The model is flexible regarding $\varphi(i)$, the fraction of hand-to-mouth households within each household of type i. Our calibration targets empirical evidence for the UK on MPCs for different demographic groups, from Albuquerque and Green (2022). Specifically, we assume that $\varphi(i) = \frac{1}{1+exp(-\Upsilon'\mathbf{X}(i))}$, where $\mathbf{X}(i)$ is a vector consisting of a constant and a number of household characteristics observed in the LCF: age (<40 years, 41-58 years, >58 years), and home ownership status (mortgagor, outright owner, renter). We then use a non-linear least squares procedure to find Υ, targeting the estimated difference in MPC of the young and middle age, relative to the old, and of mortgagors and outright owners relative to outright owners. Here, we limit ourselves to characteristics that are found to have significant effects, according to Albuquerque and Green (2022), see Table 4 column 6. We also target their estimated average quarterly MPC which is 0.11. Since Υ contains four coefficients and we have four targets, the fit is nearly perfect. Figure 1 plots the implied distribution of quarterly MPCs across household types, showing substantial heterogeneity.

Price rigidity. To calibrate the price rigidity parameter in each sector, θ_k, we follow empirical evidence on price adjustment frequencies in the United Kingdom, as documented by Dixon and Tian (2017). We convert these into quarterly Calvo probabilities, see Table 4 for the implied

[31] Note that we are implicitly normalizing the price level in each sector to be 1, which is an innocuous assumption that does not affect our calibration of the HARA utility. The details of why that is so are provided in the Appendix.

values. For *Electricity & Gas*, no direct statistics on price rigidity are available. For this sector, we assume price adjustment probability of 1/6=0.167, corresponding to an energy contract duration of 1.5 years, which is typical in the UK.

Technology. Regarding Input-Output (I-O) linkages, we calibrate the model to UK data using the matrix of industries' intermediate consumption provided by the ONS. One complication is that the categories on which the I-O tables are supplied are based on the CPA (classification of products by activity) method while our sectors are defined from the COICOP classification. We bridge these differences by constructing a mapping between the two, starting from the 10-digit goods classification and using the correspondence tables provided by the UN's Statistics Division. We also check that adjusting for the intermediate flows to the four COICOP sectors excluded from the model does not significantly change the I-O matrix used in the calibration. More details are given in Appendix D.

We further assume an AR(1) process in logs for the shock in the model. For both sectoral and aggregate productivity shocks, we assume an autoregressive coefficient $\rho_A = 0.95$. For the monetary policy shock we assume a coefficient $\rho_R = 0.25$. The monetary policy shock is scaled to correspond to an increase in the annualized nominal interest rate of 100 basis points. The aggregate productivity shock correspond to a decline in productivity of one percent. The sectoral productivity shocks are also negative, and for comparability we scale the magnitude of these shocks such that they all have the same impact on the natural demand index $\mathcal{Y}_t^*$ as the aggregate shock. This is achieved by weighting sectoral shock in sector l by a factor $\frac{\sum_k(\psi\overline{\partial_e e_k}+\bar{s}_k)\sum_m\tilde{\Omega}_{k,m}}{\sum_k(\psi\overline{\partial_e e_k}+\bar{s}_k)\tilde{\Omega}_{k,l}}$, where $\tilde{\Omega}$ is an adjustment for I-O linkages, see Appendix B.

4.2 The full model: results

With the full model at hand, we study to what extent the analytical results of the previous section hold up quantitatively. We also explore the distributional implications of shocks.

Aggregate Responses. Figure 4 plots the responses of the aggregate output gap, the CPI inflation rate, and the MCPI inflation rate, to various shocks. The responses to contractionary monetary policy shocks and negative aggregate productivity shocks, shown in the two top left panels, are typical of the New Keynesian model. Following a monetary contraction, both CPI inflation and the output gap fall, whereas following a negative productivity shock both variables increase.[32] We also observe that, for these two aggregate shocks, CPI and MCPI inflation are closely aligned, although not perfectly, which is due to heterogeneity in the slopes of sectoral NKPCs.

[32]Following a negative aggregate productivity shock, $\mathcal{Y}_t^*$ and $\mathcal{Y}_t$ both decline. But due to price rigidities, the latter falls by less than the former, and hence the output gap, $\tilde{\mathcal{Y}}_t^* = \mathcal{Y}_t - \mathcal{Y}_t^*$ increases.

The responses to negative sectoral productivity shocks are shown in the remaining panels of Figure 4. To the right side of each panel, we display an index which is negative for necessity sectors and positive for luxury sectors. Let us first consider negative productivity shocks in the two necessity sectors: *Food* and *Electricity & Gas*. In line with the analytical results –and in contrast to the aggregate productivity shock– we observe that the aggregate output gap initially *declines* following such shocks. After about a year, the output gap turns positive. Note further that, on impact, the CPI increases by substantially more than the MCPI, underscoring the quantitatively important effects of non-homotheticities.

To the central bank, the shifts of the NKPC create specific trade-offs. Initially, a marginally stronger tightening of policy would help contain inflation, but at the expense of a more negative output gap. Later on, however, this would bring down both the output gap and inflation simultaneously. This suggests that in response to negative supply shocks in necessity sectors, a delayed tightening of policy may be optimal. We will explore this more in the next section.

Let us now turn to productivity shocks in three clear luxury sectors: *Furniture*, *Recreation*, *Restaurants & Hotels*.[33] As expected, the output gap initially *increases* strongly, although quantitatively less so for a *Furniture* shock. Note further that the MCPI increases more on impact than the CPI, again illustrating the importance of non-homotheticities. From a policy perspective, a stronger monetary contraction would both close the output gap and reduce inflation, initially as well as later on.

Finally, we consider shocks to sectors which are neither clear luxuries nor necessities (*Clothing*, *Transport*, *Miscellaneous*). The response of the output gap to these shocks is mixed. This clarifies that quantitative features of the model other than non-homotheticities play a role. In particular, heterogeneity in price rigidity across sectors and I-O linkages matter. In Appendix D.1 we show responses for a version of the model in which we shut down those two features. In that case, the output gap still declines in response to negative productivity shocks in the two necessity sectors (*Food* and *Electricity & Gas*), but not in response to such shocks in any of the other sectors.

In Appendix D.1 we also plot the responses of the NKPC wedges, aggregated across sectors. We show that the movements in the novel wedges are in line with the analytical results: the $\mathcal{NH}$ wedge increases (declines) following negative shocks to necessity (luxury sectors), whereas the $\mathcal{M}$ wedge declines following any negative productivity shocks. Appendix D.1 also shows that, in line with previous literature, the $\mathcal{P}$ wedge plays a quantitatively important role when Calvo probabilities vary across sectors. Moreover, we find that $\mathcal{P}$ and $\mathcal{NH}$ often shift the NKPC in the same direction, as luxury sectors tend to have stickier prices.

[33] To gauge the extent to which a sector is a necessity or a luxury, we show to the right of each panel a luxury index defined as $100(\overline{\partial_e e_l} - \bar{s}_k)$. This index lies between -1 and 1 and is negative (positive) for necessity (luxury) sectors.

Distributional responses. Let us now consider the response of the full distribution of consumption expenditures to different shocks, see Figure 5. Each dot represents a household in the model (and hence in the LCF survey). The horizontal axis denotes the total steady-state income (expenditure) of the household, whereas the vertical axes denotes the real consumption expenditure response of the household to various shocks, averaged over the first four quarters following the shock. The red line represents the linear regression line fitted through these model-generated data.

Following a monetary contraction, consumption falls, and on average more so for low-income households. Strikingly, for any given income level there is a large degree of heterogeneity in the consumption responses. For instance, some lower-income households experience consumption gains. This heterogeneity is due to differences in the composition of labour versus asset income, as well as in steady-state consumption baskets, due to taste heterogeneity. Considering the responses to an aggregate productivity shock, we observe a similar pattern, with low-income households being hit slightly more on average. Again, even conditional on total income there is a large amount of heterogeneity, with some households increasing their consumption, for instance because they benefit from the increase in interest rates following the shock.

When we consider productivity shocks to specific sectors, we again observe that on average the consumption of the poor responds most negatively, and that there is a large amount of heterogeneity, even conditional on income. Moreover, note that the extent to which poorer households are hit varies strongly across shocks, as indicated by the slope of the red line. Indeed, the slope tends to be relatively flat for luxury sectors (*Recreation* and *Restaurants & Hotels*), but relatively steep for necessity sectors (*Food* and *Electricity & Gas*). This is a natural consequence of the fact that price increases in luxury sectors affect the rich relatively more, whereas the poor are more affected by price increases in necessity sectors.

Overall, these results suggest that, if the central bank considers distributional effects, a cost-of-living crisis may present a particularly challenging situation: in addition to the aggregate trade-off described above, an additional tightening of monetary policy may weigh most heavily on the poor, who are strongly affected by the shock to begin with. To gauge the extent to which such considerations should matter for setting interest rates, we now turn our attention to the optimal monetary policy problem.

5 Optimal Policy

Having explored the dynamics of the model under an interest rate rule, let us now analyze the normative implications for monetary policy. Specifically, we study the optimal interest rate policy under commitment.

5.1 The optimal policy problem

We consider a social planner who maximizes, at some initial date 0, a welfare function of the following form:

$$\mathcal{W} = (1-\delta)\int G(V^0(i), i)di + \delta \mathbb{E}_0 \sum_{t_0=0}^{\infty} \beta^{t_0} \int G(V^{t_0}(i), i)di, \tag{23}$$

where the first term on the right-hand side stems from pre-existing households, and the second term from current and future newborns, where the superscript t_0 denotes the period of birth. Moreover, G is a function which captures the social planner's aggregation of welfare levels of different households. The lifetime welfare of household i born at t_0 is given by:

$$V^{t_0}(i) = \mathbb{E}_{t_0} \sum_{s=0}^{\infty} (\beta(1-\delta))^s \left(U_i(\mathbf{c}_{t+s}(i)) - \chi\left(\frac{n_{t+s}(i)}{\vartheta(i)}\right)\right),$$

where setting $t_0 = 0$ gives the value of the pre-existing households. To solve the optimal policy problem, the planner sets the nominal interest rate R_t to maximize the Welfare criterion (23), subject to Equations (10)-(17) holding at all times.

Our setup allows the planner to have an arbitrary social preference function G. But in order to derive concrete policy prescriptions, we need to make further assumptions on this function. We proceed following the literature on inverse optimal taxation. First, we rule out any motive for the central bank to redistribute wealth in the absence of aggregate shocks. That is, the steady-state distribution is treated as efficient. The underlying idea is that long-run wealth redistribution is considered the domain of fiscal rather than monetary policy. We implement this assumption by imposing that:

$$G'\left(V^{t_0}(i), i\right) \partial_e v\left(e(i)\right) = 1,$$

where $v\left(e(i)\right) = \max_{\mathbf{c}(i)} U_i(\mathbf{c}(i))$ s.t. $\sum_k \int_0^1 p_k(j)c_k(i,j)dj \leq e(i)$ is the indirect utility function. Second, we set $G''\left(V^{t_0}(i), i\right) = 0$, which implies that households' percentage fluctuations in marginal utility are weighed equally by the planner. In Appendix E.1 and E.3 we show that the weight of household i in the central bank's welfare loss function (approximated to a second order) can be expressed as $g(i) = \frac{E}{\psi W n(i) + \sigma e(i)}$. Note that poor households are assigned a higher weight, as they are at a point in the utility function with more curvature, i.e. fluctuations in consumption are more costly for them. The equations characterising the optimal policy are derived and presented in Appendix E.

5.2 Analytical results under optimal policy

Before studying the optimal policy quantitatively, we present a number of analytical results in a simplified setting (again without I-O linkages and HtM agents). Proofs are provided in Appendix E.2.

Our first optimal policy result clarifies how heterogeneity and generalized preferences affect the optimal policy problem, relative to a basic NK model. To simplify the problem as much as possible we assume, in addition to (**A.1**)-(**A.2**), that there is sectoral heterogeneity in neither price stickiness nor in demand elasticities or super-elasticities (we relax this in the Appendix).[34] We obtain:

Result 7 (simplified optimal policy problem): *If (**A.1**)-(**A.2**) hold and θ_k , $\bar{\epsilon}_k$ and $\bar{\epsilon}^s_k$ are equal across sectors, then the optimal policy problem can be expressed (to a second-order approximation) as:*

$$\min_{\{\tilde{\mathcal{Y}}_t, \pi_{cpi,t}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{\sigma+\psi}{\sigma\psi} \tilde{\mathcal{Y}}_t^2 + \tilde{\vartheta} \pi_{cpi,t}^2\right)$$
$$s.t.\ \pi_{cpi,t} = \kappa \tilde{\mathcal{Y}}_t + \beta(1-\delta)\mathbb{E}_t \pi_{cpi,t+1} + \lambda(\mathcal{M}_t + \mathcal{NH}_t),$$

where $\tilde{\vartheta} = \frac{\bar{\epsilon}\theta}{(1-\theta)(1-\beta\theta)}$, and where the wedges $\mathcal{M}_t \equiv \sum_{k=1}^{K} \bar{s}_k \mathcal{M}_{k,t}$ and $\mathcal{NH}_t$ evolve independently of monetary policy (Result 1).

Thus, the optimal policy problem closely resembles the one in the basic NK model, see e.g. Galí (2015), Chapter 5. The central bank minimizes a weighted present value of the squared output gap and squared CPI inflation, subject to an aggregate NKPC. However, in our case the NKPC is shifted by the $\mathcal{M}$ and $\mathcal{NH}$ wedges which, as explained previously, are the result of non-CES and non-homothetic preferences, respectively.

Note that even in this simplified setting, household heterogeneity matters for optimal policy, since it shapes the two wedges. This point highlights the interaction between heterogeneity and generalized preferences. Under homothetic CES preferences, the two wedges would vanish and heterogeneity would become irrelevant for optimal policy, as in McKay and Wolf (2023). But once we move beyond such preferences, heterogeneity affects the NKPC and it affects optimal policy even when monetary policy cannot affect distributions (assumption **A.2**) and/or does not consider inequality part of its policy objective.

Let us now study how the optimal policy is shaped by non-homotheticities. In particular, we are interested in the extent to which optimal policy reacts differently to productivity shocks arising in necessity and luxury sectors. We assume that shocks follow AR(1) processes. In

[34] When θ_k , $\bar{\epsilon}_k$ and $\bar{\epsilon}^s_k$ vary across sectors, the inflation index becomes $\sum_k \bar{s}_k \bar{\epsilon}_k \frac{\theta_k/\tilde{\vartheta}}{(1-\theta_k)(1-\beta\theta_k)} \pi_{k,t}$, with $\tilde{\vartheta} = \sum_k \bar{s}_k \bar{\epsilon}_k \frac{\theta_k}{(1-\theta_k)(1-\beta\theta_k)}$, and the NKPC slope becomes $\sum_k \bar{s}_k \bar{\epsilon}_k \frac{\theta_k/\tilde{\vartheta}}{(1-\theta_k)(1-\beta\theta_k)} \lambda_k$.

Appendix E.2 we derive analytically the responses under optimal policy to sectoral shocks and show that the sign of the responses switch at some date t^* (which may vary across variables). Result 8 summarizes these findings:

Result 8 (signs of responses under optimal policy): *If (**A.1**)-(**A.2**) hold and $\mathcal{M}_t = 0$, then the responses of the output gap and inflation to necessity and luxury shocks have opposite signs under optimal policy, in the short, in the medium run, and in present-value terms. The signs of the responses are presented in Table 2.*

Result 8 implies that the sectoral nature of the shock is highly important for the optimal policy response. Table 2 shows that, in response to a negative productivity shock to a necessity sector, the $\mathcal{NH}$ wedge rises, as shown previously. Upon impact, the output gap and MCPI inflation fall, whereas the CPI index increases. Thus, optimal policy does not fully stamp out CPI inflation. Rather, it steers the economy to a point where the corridor between MCPI and CPI inflation includes zero (the former lies below the latter as it down-weights necessity sectors). At the same, optimal policy lets the output gap turn negative. Intuitively, optimal policy strikes a balance between the cost of CPI inflation versus the cost of a negative output gap. After some time, the signs of the responses all switch.[35] However, in present-value terms the short-term effects dominates.

Following a negative productivity shock to luxuries, the exact opposite optimal responses obtain, as shown in the lower half of Table 2. Thus, the optimal policy response critically hinges on the sectoral nature of the shock. To derive Result 8, we have shut down the $\mathcal{M}$ wedge, focusing on the $\mathcal{NH}$ wedge. In Appendix E.2 we derive analytical results on the role of the $\mathcal{M}$ wedge instead.

How does the optimal policy compare to a policy of strict targeting the CPI, i.e. $\pi_{cpi,t} = 0$ at all times? Would it be looser or tighter? Let us define a loose policy as one which targets a higher output gap and higher inflation. We can show the following:

Result 9 (optimal policy versus inflation targeting): *Compared to a strict CPI targeting policy, the optimal policy is initially and in net present value terms relatively loose (tight) following a negative necessity (luxury) shock.*

Intuitively, under a strict CPI targeting policy, the output gap declines initially following a negative necessity shock, as $\mathcal{NH}$ increases. By loosening policy, the decline in the output gap is dampened at the expense of some positive CPI inflation. This improves welfare, since welfare losses are –to a second-order approximation– quadratic in the output gap and CPI inflation.

[35] Intuitively, the prices of goods in the necessity sector which experiences the fall in productivity are initially distorted downward, due to price stickiness. At some point in time, however, the shock has mostly died out while the price level is still elevated, creating an upward distortion in the sectoral price level.

Table 2. **Sign or responses under optimal policy (Result 8)**

	Y gap	CPI	MCPI	$\mathcal{NH}$
negative necessity shock				
short run	-	+	-	+
medium run	+	-	+	-
present value	-	+	-	+
negative luxury shock				
short run	+	-	+	-
medium run	-	+	-	+
present value	+	-	+	-

Note: sign of the responses results assuming $\mathcal{M} = 0$ and (**A.1**)-(**A.2**). All negative productivity shocks. Short run refers to $t < t^*$ and medium run to $t \geq t^*$, where t^* may vary across shocks and variables. Present value discounts the responses with a factor R^{-t}. See Appendix E.2 for the derivations.

Fully stabilising either the output gap or CPI inflation is therefore never optimal.

5.3 Quantitative dynamics under optimal policy

Result 9 suggests that the optimal policy response to cost-of-living crisis (i.e. a negative shock to necessities) can indeed be rather specific. Following the initial shocks around 2021, central banks were seen to be relatively slow in tightening policy. Interestingly, this appears in line with the optimal policy in the model, at least qualitatively.

We now explore quantitatively the optimal policy responses to various sectoral shocks, and study to what extent the optimal policy response to shocks in sectors like *Food* or *Electricity & Gas* is indeed relatively loose, compared to a typical policy rule $\hat{R}_t = \phi\pi_{cpi,t}$ (with $\phi = 1.5$) and compared to the optimal response to other shocks. In order to make this comparison quantitatively, we exploit the fact that one can always implement the optimal interest rate path $\{\hat{R}_t\}_{t=0}^{\infty}$ as a rule $\hat{R}_t = \phi\pi_{cpi,t} + u_t^R$ where $\{u_t^R\}_{t=0}^{\infty}$ is a specific time path for the deviation from the rule ("optimal guidance"), announced when the productivity shock initially hits. We simulate the model both under such an interest rate rule and under optimal policy, and then numerically solve for the guidance path that implements the optimal policy. This path then quantifies how tight or loose the optimal policy is relative to the simple CPI-based rule. In Appendix D.2 we provide the details of this procedure.

The left panel in Figure 6 plots the optimal guidance for the aggregate and sectoral productivity shocks in the full model. In line with the analytical results, optimal policy is initially significantly looser than the rule following a negative necessity shock. Following negative shocks to luxuries, the optimal policy is also initially looser than the rule, but less so than following necessity shocks. We thus find that the sectoral source of the shock indeed has significant quantitative consequences for optimal policy, in line with the analytical results.[36]

How important are redistributive motives in driving the optimal policy? In the right panel of Figure 6 we shut down the redistributive motives of monetary policy.[37] Qualitatively, the results are unchanged, in the sense that the optimal policy response to negative necessity shocks is significantly looser than the response to luxury shocks. Quantitatively however, the redistributive motives push towards a front-loaded accommodative (i.e. looser) policy for all shocks, as this helps to redistribute towards poorer households, who tend to be more heavily affected in utility terms.

Figure 7 shows the response of the output gap and CPI inflation under optimal policy, for an aggregate productivity shock and productivity shocks to *Food* (lowest luxury index) and *Recreation* (highest luxury index). The quantitative responses are again consistent with the analytical findings. Without the redistribution motive, following a negative *Food* shock, the output gap initially is negative, while CPI inflation increases. For a *Recreation* shock, we observe the exact opposite. Once we include a redistribution motive, the responses of the output gap and inflation are both pushed upwards, at least initially.

6 Conclusion

In this paper we addressed the question how monetary policy should respond to sector-specific supply shocks. To this end, we developed a multi-sector New-Keynesian model with household inequality and generalized non-homothetic preferences. An advantage of the framework is that it is relatively tractable, simplifying computations and allowing for analytical results to be derived. Moreover, it can be disciplined directly with data on heterogeneity in income, wealth, MPCs and expenditure baskets.

We showed how, due to non-homothetic and non-CES preferences, two new wedges emerge in the New Keynesian Phillips Curve (NKPC), which directly affect policy trade-offs and which are quantitatively important. In particular, after a negative supply shock to necessity sectors, the NKPC tends to shift upward, creating a policy trade-off between bringing down inflation and avoiding a negative output gap. After studying the optimal policy, we found that –because of this shift in the NKPC– the optimal policy to a negative necessity shock is relatively loose, while later on it tightens.

More generally, the framework developed in this paper offers a rich framework to study the macro effects and distributional consequences of a wide range of shocks and policies. For

[36] Consistent with the quantitative results in the previous section, we find that negative shocks to *Transport* (neither a necessity nor a luxury) call for a relatively loose optimal policy, which is again due to the low degree of price rigidity in this sector and the position of this sector in the I-O matrix. Appendix D.1 repeats the exercise shutting down I-O linkages and heterogeneity in price rigidity. In this case, the optimal policy response to a transport shock is similar to the response to an aggregate productivity shock, i.e. much less loose.

[37] In Appendix E.3 we derive conditions shutting down such motives based on the welfare loss function.

instance, it can be used to study (targeted) fiscal interventions and price controls, which have been used in recent years. We explore such policies in ongoing research.

References

Acharya, S., E. Challe, and K. Dogra (2023) "Optimal monetary policy according to HANK," *American Economic Review*, 113 (7), 1741–1782.

Albuquerque, B. and G. Green (2022) "Financial Concerns and the Marginal Propensity to Consume in Covid Times: Evidence from UK Survey Data," Bank of England staff working paper No.965.

Amiti, M., O. Itshkhoki, and J. (2018) Konings (2019) "International Shocks, Variable Markups and Domestic Prices," *Review of Economic Studies*, 86 (6), 2356–2402.

Argente, David and Munseob Lee (2021) "Cost of Living Inequality During the Great Recession," *Journal of the European Economic Association*, 19 (2), 913–952.

Auclert, A., H. Monnery, M. Rognlie, and L. Straub (2023) "Managing an Energy Shock: Fiscal and Monetary Policy," Working paper.

Auclert, Adrien (2019) "Monetary Policy and the Redistribution Channel," *American Economic Review*, 6, Working Paper.

Baqaee, D., E. Farhi, and K Sangani (2021) "The Supply-Side Effects of Monetary Policy," Working paper.

Baqaee, David and Emmanuel Farhi (2019) "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 87 (4), 1155–1203.

Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas Sargent (2021) "Inequality, Business Cycles and Monetary-Fiscal Policy," *Econometrica*, 6, 2559–2599.

Blanco, C. and S. Diz (2021) "Optimal monetary policy with non-homothetic preferences," mimeo.

Boppart, Timo (2014) "Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences," *Econometrica*, 82, 2167–2196.

Cavallari, L. and F. Etro (2020) "Demand, Markups and the Business Cycle," *European Economic Review*, 127.

Challe, E. (2020) "Uninsured unemployment risk and optimal monetary policy in a zero-liquidity economy," *American Economic Journal: Macroeconomics*, 12 (10), 241–283.

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

Dávilla, E. and A. Schaab (2022) "Optimal Monetary Policy with Heterogeneous Agents: Discretion, Commitment, and Timeless Policy," Working paper.

De Loecker, Jan and Frederic Warzynski (2012) "Markups and Firm-Level Export Status," *American Economic Review*, 102 (6), 2437–71.

Dixon, Huw David and Kun Tian (2017) "What We can Learn About the Behaviour of Firms from the Average Monthly Frequency of Price-Changes:An Application to the UK CPI Data," *Oxford Bulletin of Economics and Statistics*, 79 (6), 907–932.

Engel, Ernst (1857) "Die Productions- und Consumtionsverhältnisse des Königreichs Sachsen," *Zeitschrift des Statistischen Bureaus des Koniglich Sachsischen Ministerium des Inneren*, 8-9.

Galí, Jordi (2015) *Monetary Policy, Inflation and the Business Cycle*, Princeton, N: Princeton University Press.

Hamilton, Bruce (2001) "Using Engel's Law to Estimate CPI Bias," *American Economic Review*, 91 (3), 619–630.

Herrendorf, B., R. Rogerson, and A. Valentinyi (2014) "Growth and Structural Transformation," *Handbook of Economic Growth*, 2, 855–941.

Houthakker, H.S. (1957) "An International Comparison of Household Expenditure Patterns, Commemorating the Centenary of Engel's Law," *Econometrica*, 25, 532–551.

Jaravel, X. (2019) "The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector," *Quarterly Journal of Economics*, 134 (2), 715–783.

Jaravel, X. and A. Olivi (2021) "Prices, Non-homotheticities, and Optimal Taxation The Amplification Channel of Redistribution," Working paper.

Kaplan, G. and S. Schulhofer-Wohl (2017) "Inflation at the Household Level," *Journal of Monetary Economics*, 91, 697–743.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante (2017) "Monetary Policy According to HANK," *American Economic Review*, 108 (3), 697–743.

Kimball, M.S. (1995) "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking*, 27 (4), 1241–1277.

LaO, J. and A. Tahbaz-Salehi (2019) "Optimal Monetary Policy in Production Networks," Working paper.

McKay, A. and C. Wolf (2023) "Optimal Policy Rules in HANK," Working paper.

McKay, Alisdair, Emi Nakamura, and Jon Steinsson (2016) "The Power of Forward Guidance Revisited," *American Economic Review*, 106 (10), 3133–3158.

Melcangi, D. and V. Sterk (2019) "Stock Market Participation, Inequality and Monetary Policy," Working paper.

Moll, Benjamin, Moritz Schularick, and Georg Zachmann (2023) "The Power of Substitution: the Great German Gas Debate in Retrospect," *Brookings Papers on Economic Activity*.

Mongey, Simon and Michael Waugh (2023) "Pricing Inequality," Work in progress.

Nord, Lukas (2023) "Shopping, Demand Composition, and Equilibrium Prices," Working paper.

Nuno, G. and C. Thomas (2022) “Optimal Redistributive Inflation,” *Annals of Economics and Statistics*, 146, 3–64.

ONS (2022) “Inflation and the cost of living for household groups, UK: October 2022,” Office for National Statistics, released 16 November.

Pasten, E., Schoenle R., and M. Weber (2020) “The Propagation of Monetary Policy Shocks in a Heterogeneous Production Economy,” *Journal of Monetary Economics*, 116, 1–22.

Rubbo, E. (2023) “Networks, Phillips Curves and Monetary Policy,” *Econometrica*, 91 (4), 1417–1455.

Sangani, Kunal (2023) “Markups Across the Income Distribution: Measurement and Implications,” Working paper.

Schaab, A. and S.T. Tan (2023) “Monetary and Fiscal Policy According to HANK-IO,” Working paper.

Smets, F. and R. Wouters (2007) “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97 (3), 586–606.

Straub, Ludwig (2017) “Consumption, Savings, and the Distribution of Permanent Income,” Working Paper.

Xhani, D. (2021) “Correcting Market Power with Taxation: a Sufficient Statistic Approach,” Working paper.

Figures

Figure 1. Steady-state distributions.

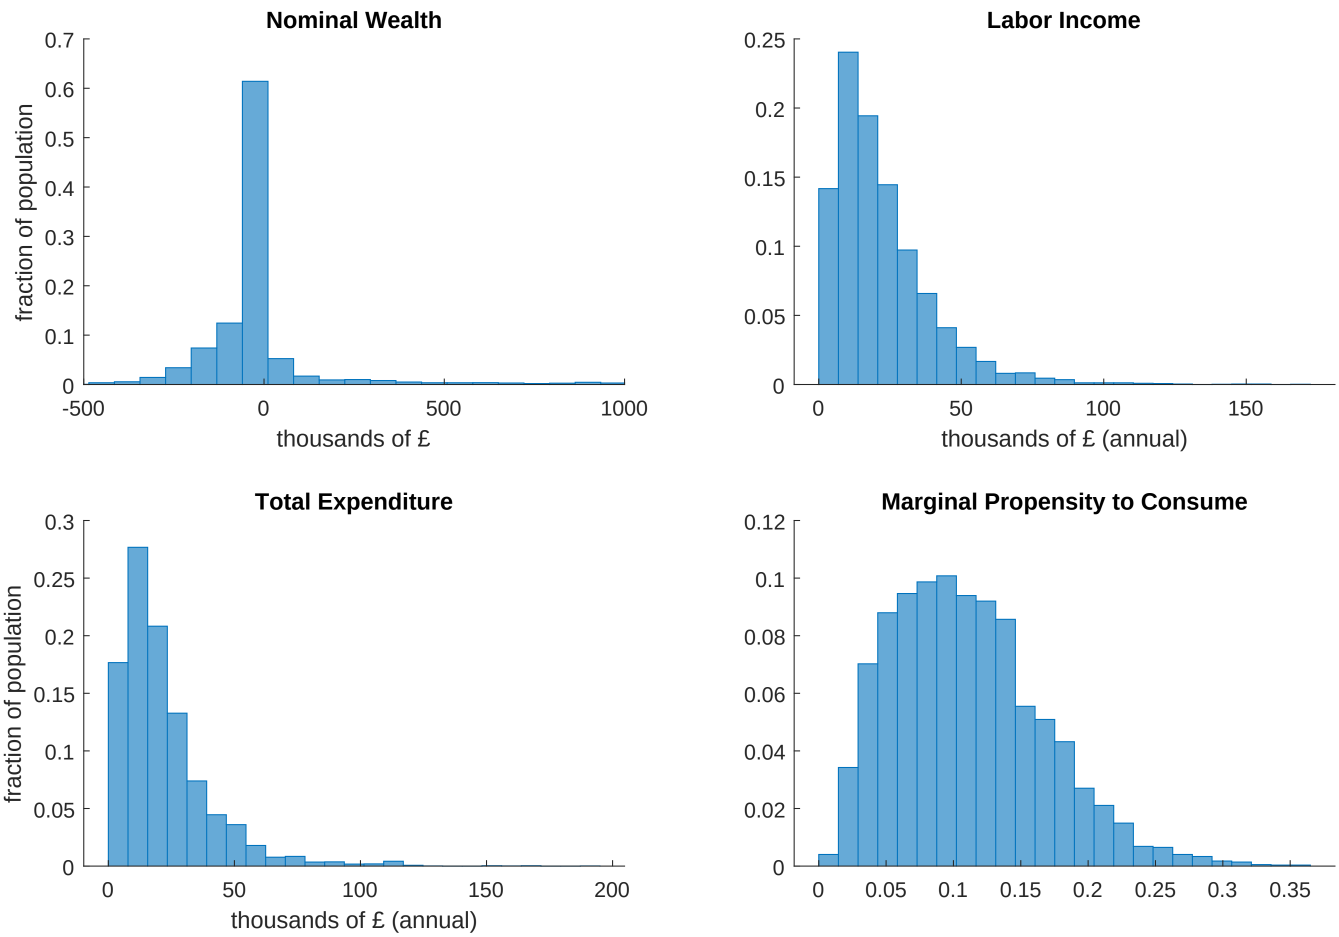

Notes: Data from Living Costs and Food Survey 2019 and authors' calculations, see main text.

Figure 2. Household budget shares by total expenditure decile.

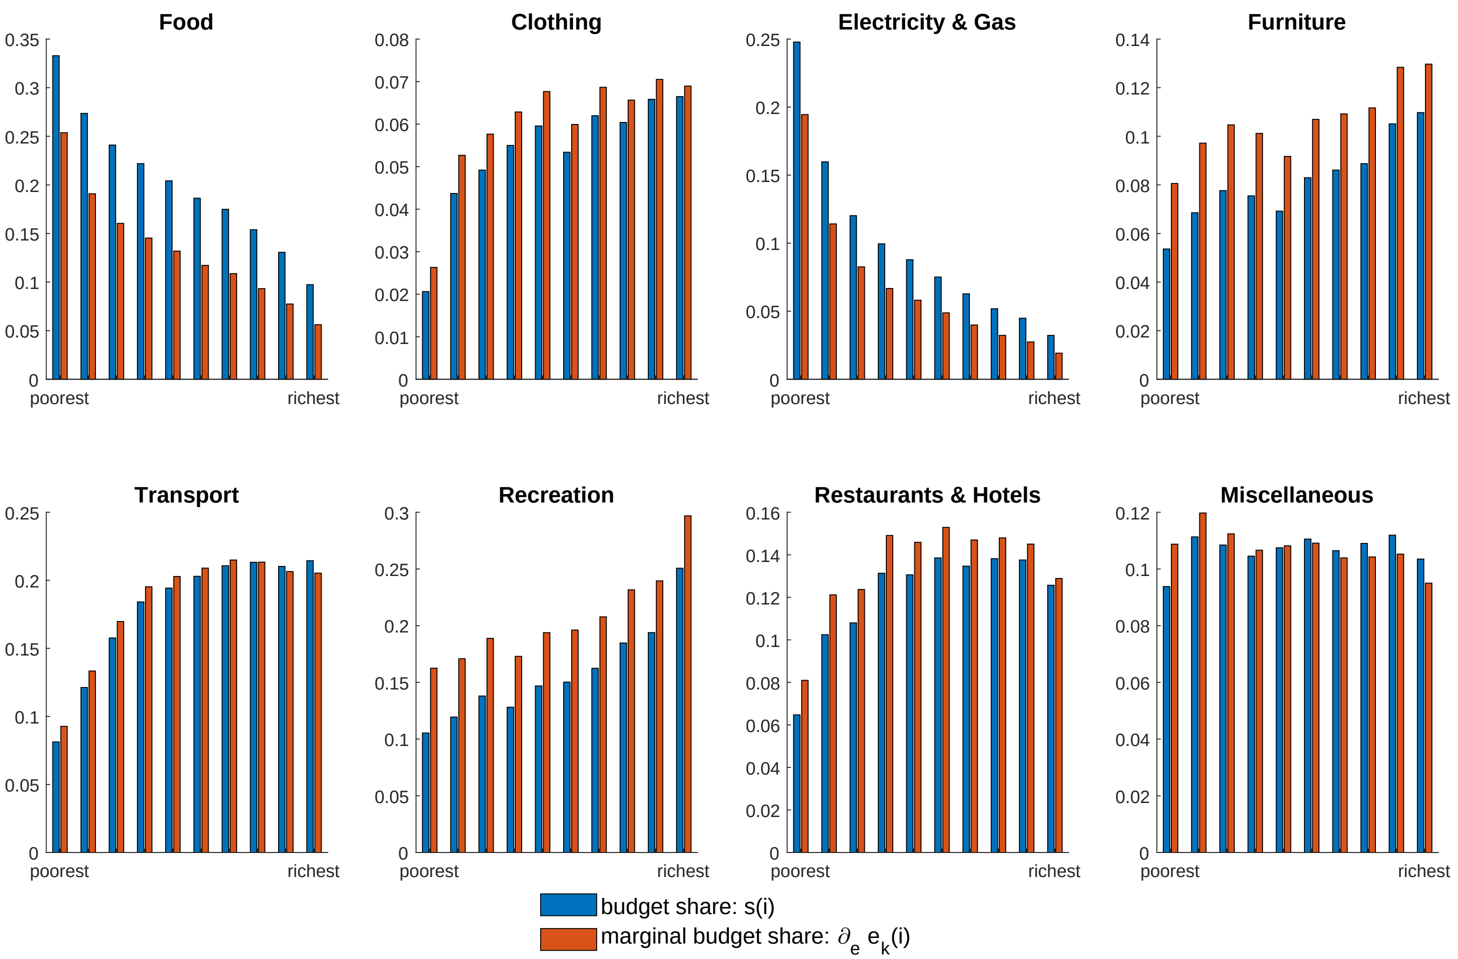

Notes: Budget shares averaged within deciles of total expenditure, ordered from poorest (lowest decile) to richest (highest decile). Source: Living Costs and Food Survey 2019.

Figure 3. Distribution of demand elasticities by sector.

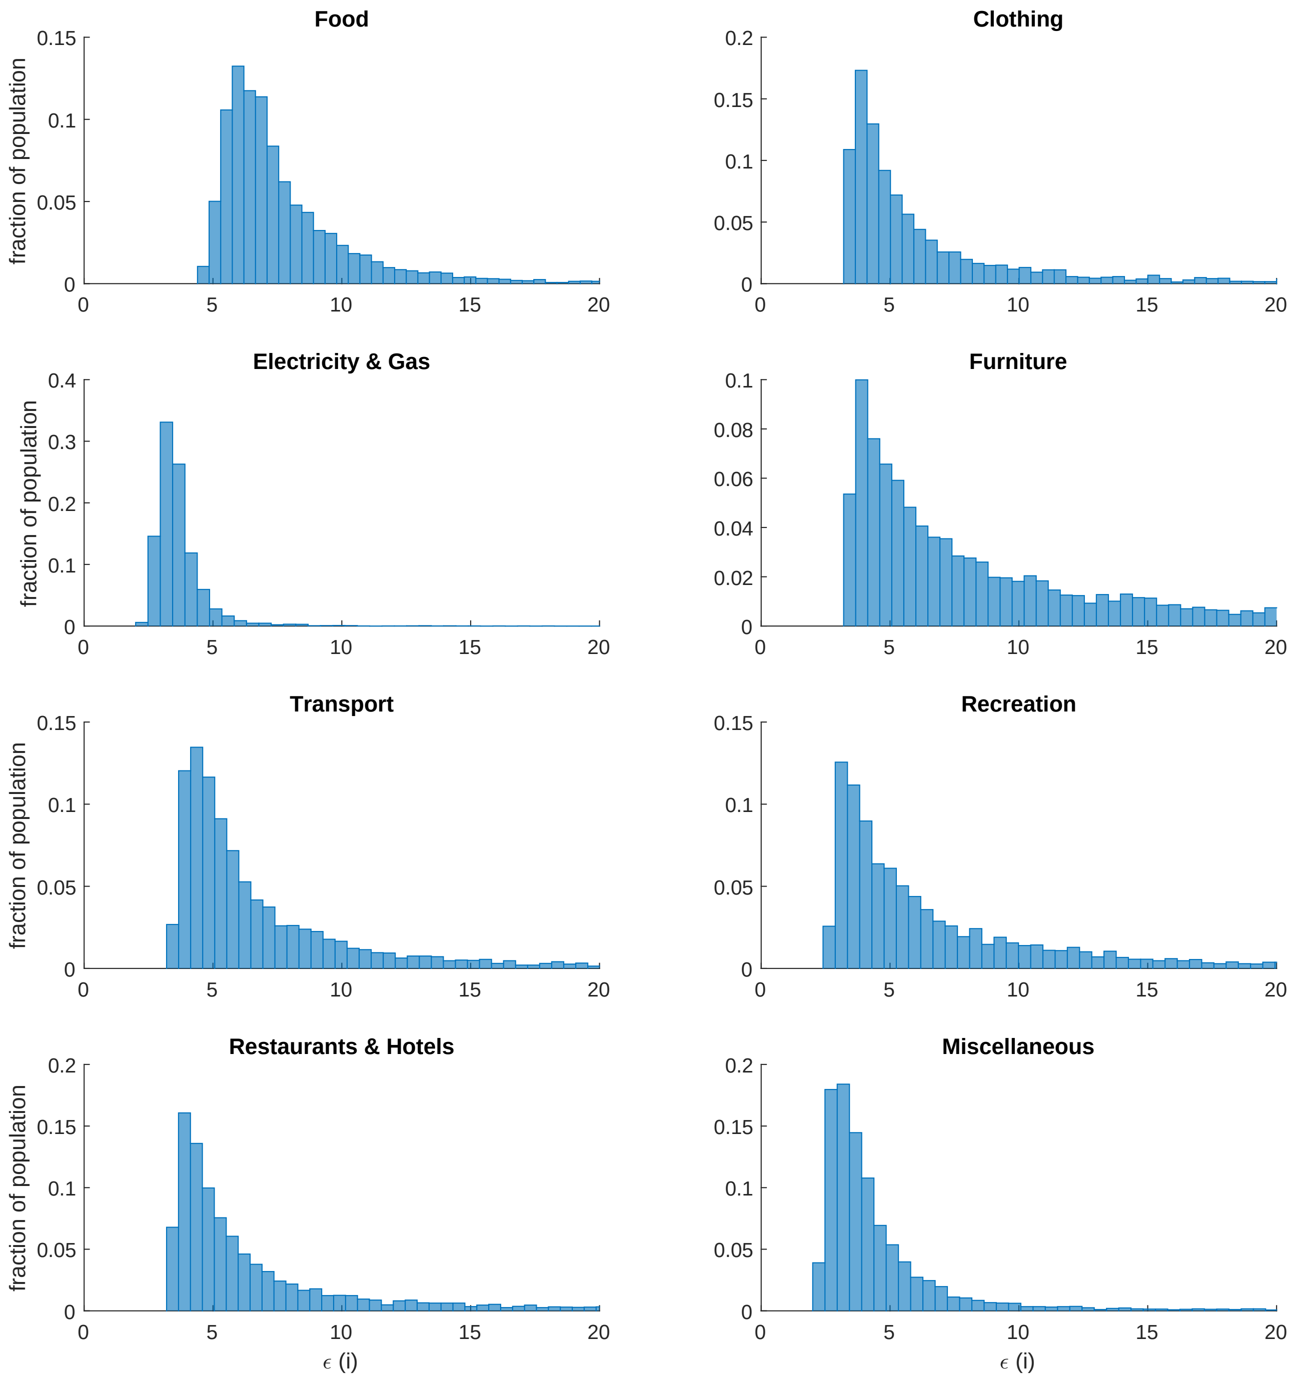

Notes: Histogram of $\epsilon_k(j)$, the price elasticity of demand across households (by sector).

Figure 4. Responses in the baseline model: all shocks.

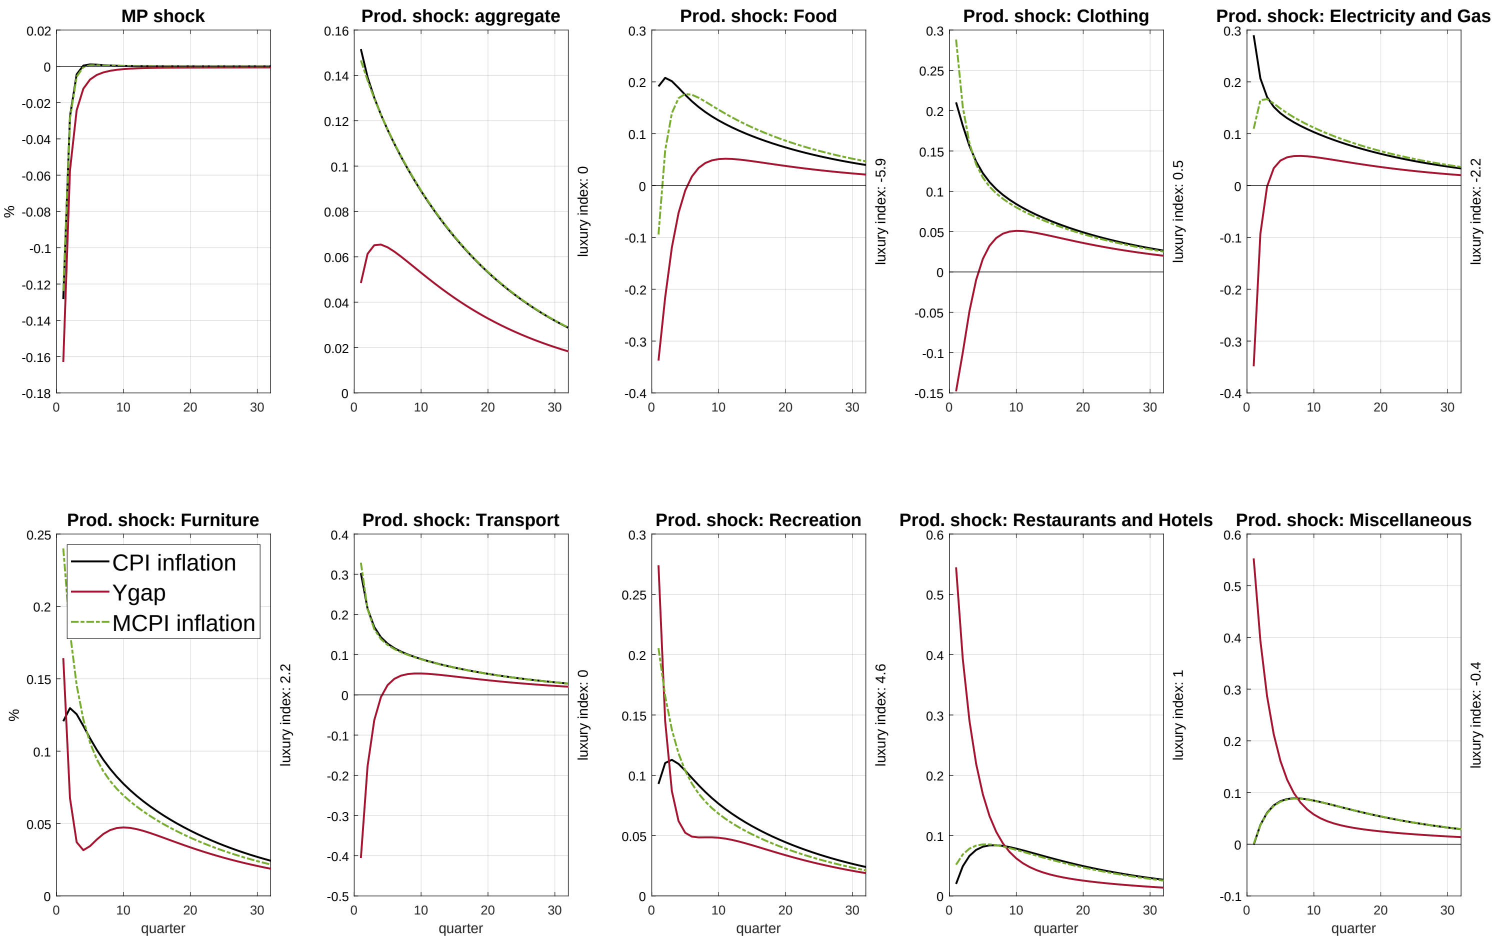

Notes: Impulse Response Functions are generated from the baseline the model, including heterogeneous Calvo probabilities across sectors across sectors, Input-Output linkages, and Hand-to-Mouth households. Responses for productivity shocks are for a 1 percent decline in productivity where scaled for comparability (see main text). On the right axis, the luxury index is defined as $100(\overline{\partial_e e_l} - \bar{s}_k)$.

Figure 5. Heterogeneous consumption responses to aggregate and sectoral shocks.

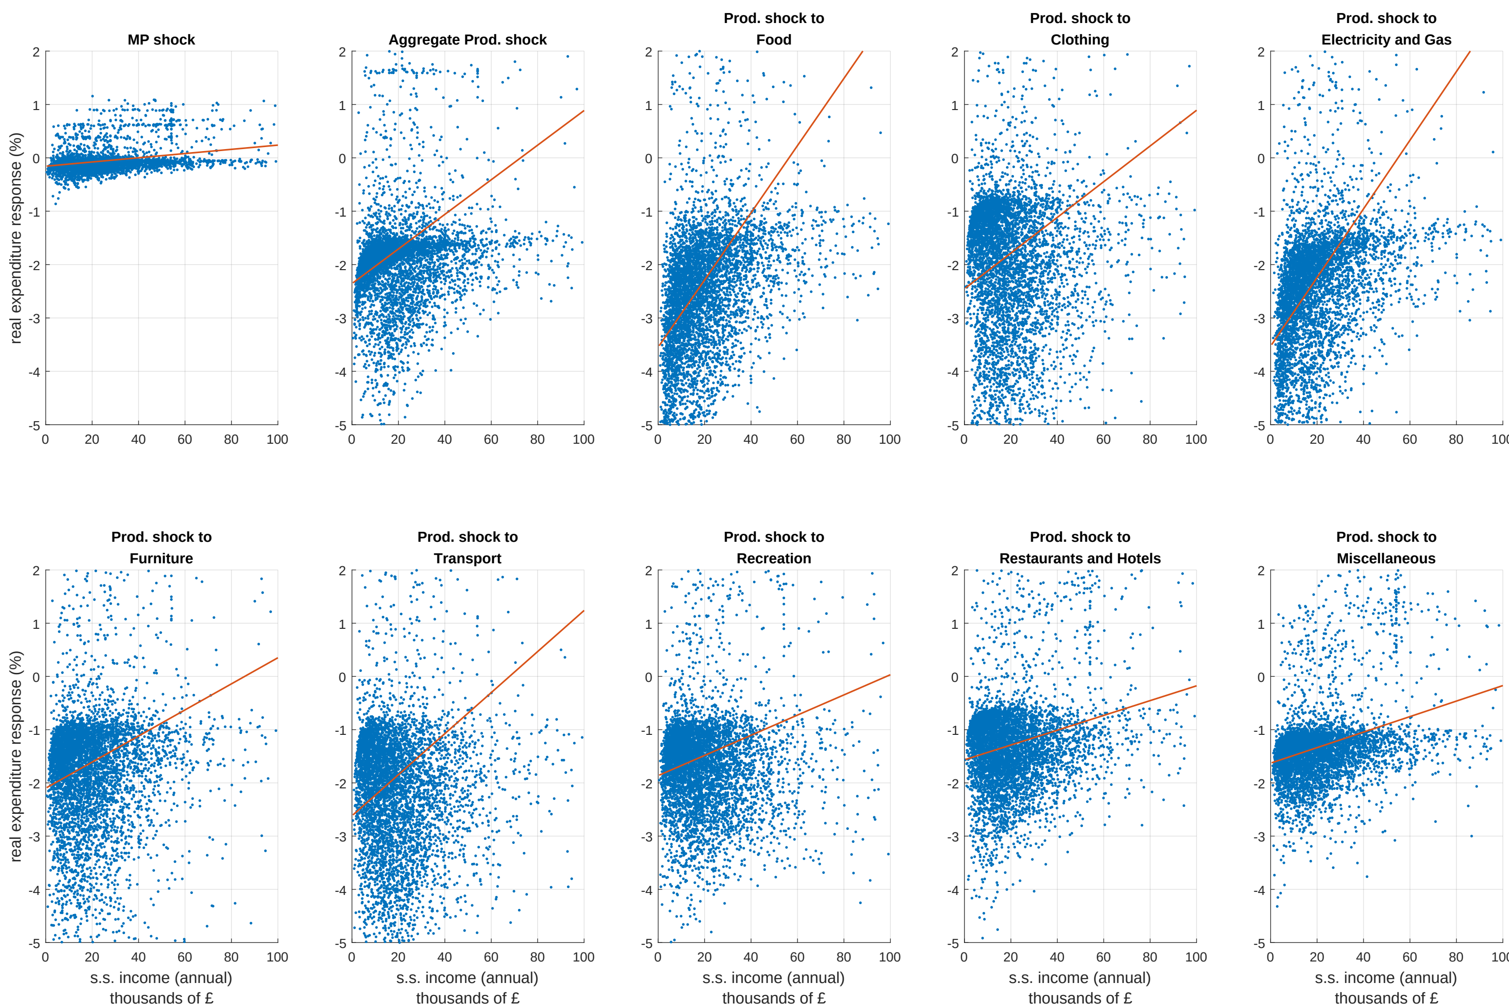

Notes: Response of real expenditures by steady-state income, generated from the baseline model and averaged over first four quarters following the shock. Dots denote individual households. Red lines are fitted 10th order polynomials. All productivity shocks are negative.

Figure 6. Optimal policy relative to Taylor rule.

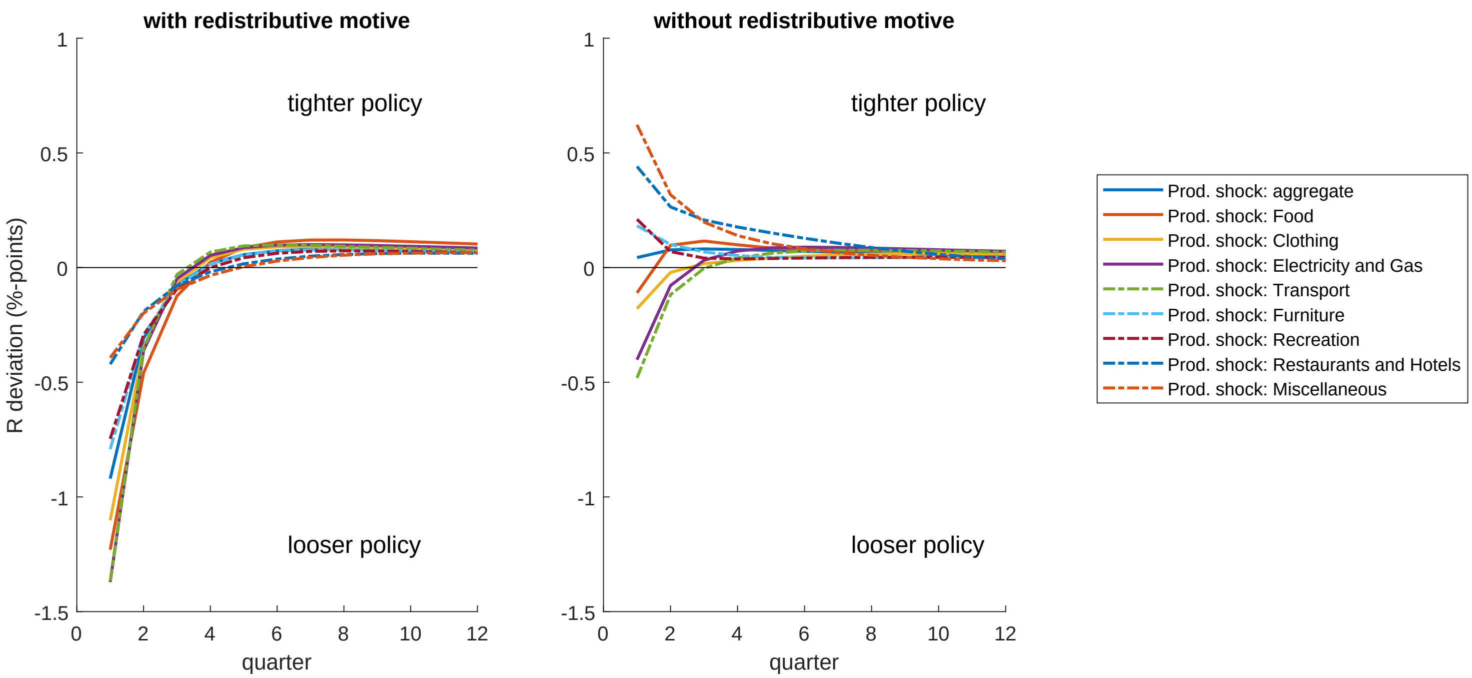

Notes: Deviations from the Taylor rule $\hat{R}_t = 1.5\pi_{cpi,t}$ which implement Optimal Policy ("optimal guidance"). Higher values mean that optimal monetary policy is tight relative to this rule. See the main text for details. All productivity shocks are negative.

Figure 7. Optimal policy relative to Taylor rule.

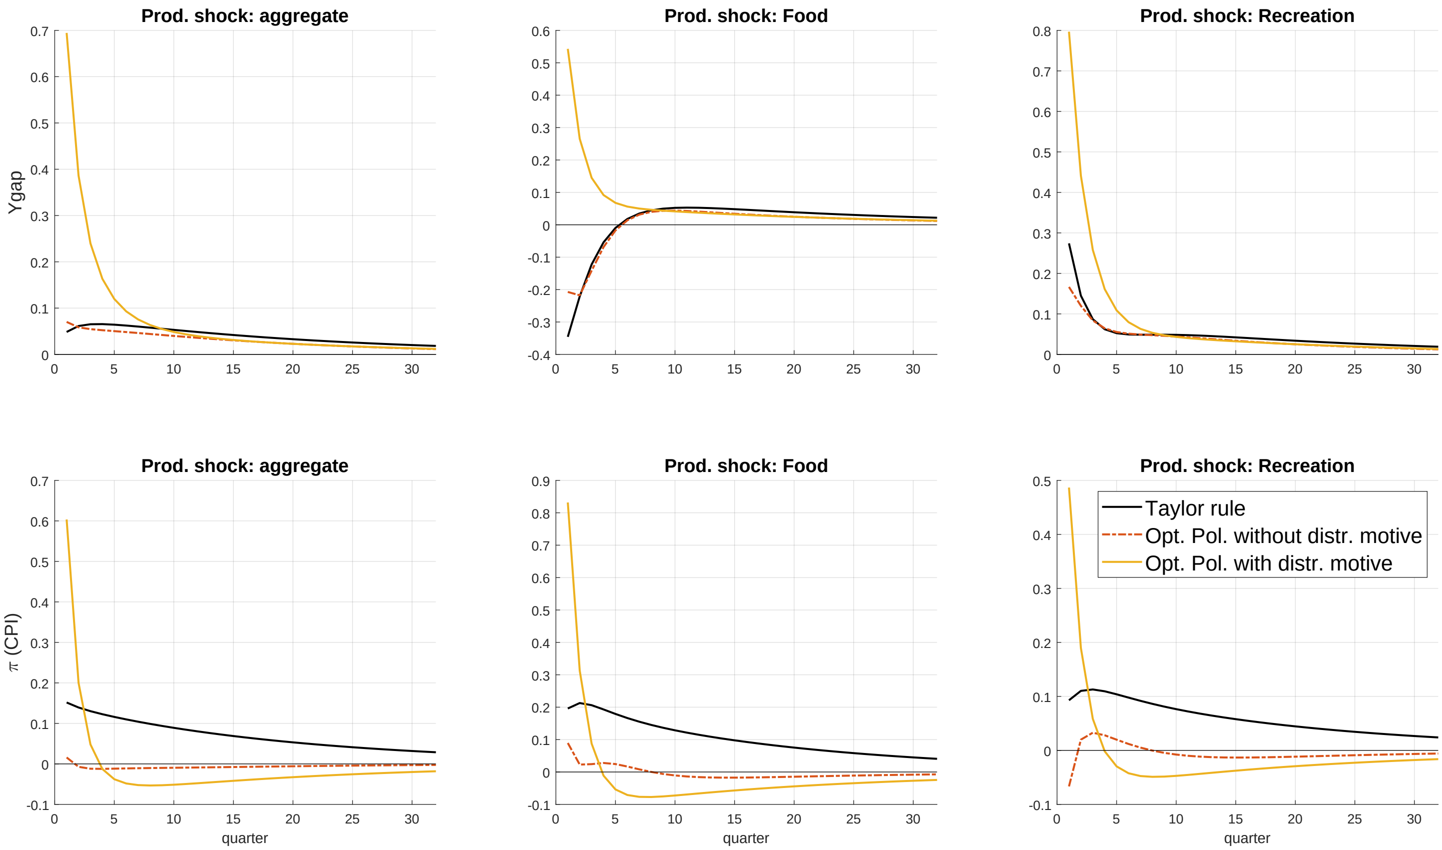

Notes: Responses of the output gap and CPI inflation. All productivity shocks are negative.

Tables

Table 3. Aggregate parameter values.

Parameter	description	value
β	subjective discount factor	0.99
ψ	Frisch elasticity	1
σ	elasticity of intertemporal substitution	1
δ	death probability	0.0083
ϕ	Taylor rule coefficient	1.5
η	cross-sector elasticity of substitution	0.1
ρ_R	persistence monetary policy shock	0.25
ρ_A	persistence productivity shocks	0.95

Table 4. Sector-level parameter values.

Sector	$\bar{\epsilon}_k$	$\bar{\epsilon}_k^s$	$\bar{s}_k$	$\overline{\partial_e e_l}$	θ_k	κ_k	λ_k	Γ_k
Food	6.5775	2.3903	0.1574	0.0988	0.4100	0.7608	0.5998	0.0386
Clothing	4.8259	1.6397	0.0580	0.0631	0.3900	1.1810	0.6735	0.0929
Electricity & Gas	3.2525	0.9654	0.0630	0.0412	0.1667	2.6410	2.9244	0.0807
Furniture	4.9651	1.6993	0.0910	0.1133	0.4600	0.5731	0.4488	0.1018
Transport	5.0243	1.7247	0.2015	0.2018	0.2600	1.5120	1.4812	0.0806
Recreation	3.8950	1.2407	0.1858	0.2318	0.5100	0.3667	0.3341	0.1248
Restaurants & Hotels	4.7313	1.5991	0.1338	0.1440	0.7200	0.1317	0.0788	0.0883
Miscellaneous	3.1534	0.9229	0.1096	0.1061	0.6700	0.1313	0.1168	0.1210

Notes: $\bar{\epsilon}_k$: demand elasticity (household aggregate), $\bar{\epsilon}_k^s$: superelasticity (household aggregate), $\bar{s}_k$: budget share (household aggregate), $\overline{\partial_e e_l}$: marginal budget share (household aggregate), θ_k: Calvo probability, κ_k: slope NKPC w.r.t. output gap, λ_k slope NKPC w.r.t wedges, Γ_k: slope endogenous markup wedge w.r.t efficient demand index. See the main text and Appendix B for the definitions.