

Non-Essential Business-Cycles

Michele Andreolli, Natalie Rickard, and Paolo Surico*

February 2025

Abstract

This paper documents three main regularities on post-WWII data for the United States: (i) spending on non-essentials is more sensitive to the business-cycle than spending on essentials; (ii) earnings in non-essential sectors are more cyclical than earnings in essential sectors; (iii) non-essential industries employ a larger share of hand-to-mouth workers. We develop and estimate a two-sector New-Keynesian model with non-homothetic preferences, hand-to-mouth workers and sectoral labour force heterogeneity that is consistent with these findings. We use the model to revisit the transmission of stabilisation policies. A main finding is that the interaction of cyclical product demand composition and cyclical labour demand composition greatly amplifies the effects of monetary policy, which in our economy can be equivalently implemented by a mix of consumption and labour taxes. However, a VAT change applied only to non-essentials (only to essentials) has a far larger (smaller) impact than a uniform VAT change of the same size.

Keywords: income elasticity, intertemporal substitution, cyclical labour composition.

JEL Classification Codes: E52, D31, E21.

*We are very grateful to Mark Aguiar, Martin Beraja, Florin Bilbiie, Timo Boppart, Andrea Galeotti, Fadi Hassan, Killian Huber, Nir Jaimovich, Peter Karadi, Hasmat Khan, Joseba Martinez, Emi Nakamura, Elias Papaioannou, Christina Patterson, Facundo Piguillem, Giorgio Primiceri, Ricardo Reis, Ayşegül Şahin, Jon Steinsson, Ludwig Straub, Gianluca Violante, Johannes Wieland, as well as conference and seminar participants at PSE Macro Days, Barcelona School of Economics Summer Forum, CREi/UPF, Boston College, LBS, Norges Bank, EWMES, SEA, RES, ECB ChaMP Conference, Midwest Macro, NASM of the Econometric Society, SED, NBER SI 2024 (monetary economics), Oslo Macro Conference, Princeton University and the AEA for very helpful comments and suggestions. Chiara Vergat provided outstanding research assistance. Paolo Surico gratefully acknowledges financial support from ESRC (Grant ref. ES/Y002490/1). Michele Andreolli gratefully acknowledges support from the Wheeler Institute for Business and Development. Address for correspondence: Michele Andreolli (Boston College) andreoll@bc.edu, Natalie Rickard (London Business School) nrickard@london.edu, and Paolo Surico (London Business School and CEPR) psurico@london.edu.

“The poorer a family is, the greater the proportion of total expenditure it must devote to the provision of nourishment.”

– Ernst Engel (1857)¹

1 Introduction

Almost 170 years on and the influence of Engel’s law in economics cannot be overstated. In public finance, it is used to design tax policies and perform welfare comparisons; in growth theory, it is key to rationalize structural transformation; in trade, it contributes to solve the Leontief paradox and explain intra-industry trade among countries with similar levels of income (i.e. Linder’s hypothesis); in development, it is crucial to measure costs of living, poverty and inequality. In business-cycle analysis, however, Engel’s law does not seem to have made any major breakthrough, yet. This is all the more surprising given that a long tradition in macroeconomics advocates a central role for the spending response to changes in household resources to evaluate the effects of stabilisation policies and understand business-cycle fluctuations.

Our analysis starts by connecting Engel’s law with four novel pieces of evidence on post-WWII U.S. data. First, during recessions, spending on non-essentials contracts significantly more than spending on essentials. Second, labour earnings in industries that produce non-necessities is significantly more cyclical than earnings in essential sectors. Third, workers employed in non-essential industries tend to earn lower salaries than workers in essential sectors and are far more likely Hand-to-Mouth (HtM). These data patterns are more pronounced at the bottom of the income distribution, where labour earnings are particularly cyclical. Fourth, we find little evidence of significant price effects across essential and non-essential spending.

The regularities that we document in this paper are pervasive: they hold across three empirical settings that differ markedly by the source of variation that they exploit. Using newly constructed time-series and cross-sectional data on prices, consumption and earnings for essentials and non-essentials, we begin by describing the main features of a typical post-WWII American recession. Then, we move to regional evidence across U.S. states. Finally, we present impulse responses to an identified monetary policy shock that are estimated using the smooth Local Projection (LP) method proposed by [Barnichon and Brownlees \(2019\)](#).

¹“In Engel’s thinking, lowering income acts like a litmus test on the consumer’s priorities: it crowds out expenditures related to wants that are less basic and leaves those expenditures related to more fundamental wants.” ([Chai and Moneta, 2010](#)). We use basic, essentials and necessities, on the one hand, and luxuries, non-essentials and non-necessities, on the other hand, as synonyms for goods and services with, respectively, a lower (e.g. *below one*) and a higher (e.g. *above one*) income elasticity of demand.

Importantly, our empirical findings are not confounded by other consumption categorizations such as durable goods and tradables: the lion’s share of non-essentials is among non-durable goods and services.

We interpret our findings through the lens of a two-sector New-Keynesian model with: (i) non-homothetic preferences, (ii) hand-to-mouth agents, (iii) heterogeneity in sectoral labour composition. Households consume two types of goods that differ by their income elasticity: non-necessities are easier to postpone in the face of a negative shock and to anticipate after a positive income change. Workers have either low productivity and are hand-to-mouth or have high productivity and are unconstrained. Non-essential sectors employ a larger share of the former, as in the data. Because of non-homothetic preferences, workers with high productivity spend a larger budget share on luxuries. This implies that the intertemporal elasticity of substitution is heterogeneous across goods and households. A key model insight is that the discretionary spending of the rich drives the earnings and spending of the poor over the business-cycle.

We estimate the structural model and perform counterfactual simulations that highlight three major results. First, the cyclical demand for non-essentials *combined* with a larger share of HtM workers employed in non-necessity sectors makes the composition of labour demand endogenously cyclical: the labour earnings of the HtM workers in non-essential industries decline significantly more during recessions. Second, the *interaction* between the cyclical spending composition (i.e. Engel’s law) and the heterogeneous sectoral labour composition accounts for a significant share of the effects of monetary policy on aggregate consumption, and greatly amplifies business-cycle fluctuations relative to both the representative agent case and models that feature only heterogeneity in either product or labour demand. Third, we show analytically that, in a representative agent framework, non-homotheticity leads to no amplification at all.

As for fiscal policy, our analysis begins by generalising to a model with heterogenous agents and heterogenous goods, the important theoretical result in [Correia et al. \(2013\)](#) that a suitable combination of consumption and labour income tax changes can exactly replicate the effects of monetary policy. We go further, however, in showing that, in our framework, a VAT change levied only on non-essentials (only on essentials) amplifies (dampens) the effects of fiscal policy relative to a VAT change that raises the same tax revenues but applies uniformly to all sectors. Interestingly, a hike (cut) only in the VAT rate of non-essentials hits (benefits) the earnings of low-income households more than a VAT change applied only to essentials or uniformly.

Altogether, our empirical and theoretical results uncover a novel mechanism for the trans-

mission of business-cycle fluctuations and the dynamic effects of stabilization policies. In response to a temporary decline in household income, consumers face an incentive to postpone non-essential spending, such as dining out, entertainment, hospitality, buying a new car, enrolling in a gym or in an educational course. This intertemporal shift triggers a fall in labour demand from industries that produce non-necessity goods and services. As the labour force in those sectors comprises a significantly larger share of hand-to-mouth workers, the decline in earnings for non-essential employees, such as waiters, ushers, hotel workers, blue collars, fitness trainers and teachers, sets in motion a second round of consumption effects that spillover to essential spending and greatly exacerbate the initial contraction.

Related literature. Our paper contributes to several strands in macroeconomics. Growing research efforts have focused on measuring consumption inequality (Aguiar and Bils, 2015) and the welfare effects of demand and supply shifts along the income distribution (Bergman et al., 2025). Bilbiie (2020) identifies a crucial role for the covariance between the marginal propensity to consume and earning cyclicalities across workers, while Ravn and Sterk (2017, 2021) and Bilbiie et al. (2023) emphasize counter-cyclical income risk. Cloyne et al. (2020) document that the indirect effects of monetary policy on income across households are key to account for the aggregate consumption response, while Holm et al. (2021), Amberg et al. (2022), Andersen et al. (2023) find significant heterogeneity in the earning responses along the income distribution.²

Our work builds on Patterson (2023), who shows that workers with higher MPC experience more cyclical wages. Relative to this important empirical study, we make two innovations. First, we identify two novel sources of heterogeneity: in spending across essential and non-essential categories; in the labor force composition across essential and non-essential industries. Second, we derive potentially significant policy implications: by exploiting the labour market heterogeneity across industries and the spending heterogeneity across consumption categories, governments can design a mix of labour taxes and targeted VATs that can better stabilize the business-cycle. More generally, we propose a novel framework for understanding fluctuations in household expenditure based on the distinction between essentials and non-essentials.

A prominent literature has studied the impact of demand composition on business-cycle dynamics. McKay and Wieland (2019), Zorzi (2020) and Beraja and Wolf (2021) show that

²Another influential literature has emphasized the role of heterogeneity in the marginal propensity to consume. For instance, Kaplan et al. (2018), Auclert (2019), Debortoli and Galí (2017) and Cloyne et al. (2020) separate the direct effects from the indirect effects of monetary policy on consumption, while Auclert et al. (2020) and Bilbiie et al. (2022) highlight the amplifying role of capital investment.

the lumpy nature of durable expenditure can alter the transmission of policy and business cycle shocks; [Jaimovich et al. \(2019\)](#) find that the demand shift towards lower quality products amplified the employment drop during the Great Recession; [Baqae and Farhi \(2018\)](#), [Flynn et al. \(2021\)](#) and [Rubbo \(2023\)](#) revisit the transmission of fiscal and monetary policy in economies with multiple sectors and an input-output network. The distinction between essentials and non-essentials sets us apart from earlier works. First, we document a strong covariance between non-essential spending and non-essential earnings, but find a much lower sectoral comovement between consumption and income in either the durable sector or the non-tradeable sector. Second, non-essential industries are characterized by a much larger share of HtM workers than durable goods producers. Third, we focus on shifts in spending composition towards essentials *across* sectors (e.g. from restaurants to grocery stores), rather than towards low quality products *within* a specific spending category (e.g. from premium to grocery store brands).

Finally, earlier contributions have used non-homothetic preferences to analyse salient features of: (i) structural transformation, such as Kaldor’s facts ([Foellmi and Zweimüller, 2008](#), [Boppart, 2014](#)) and income versus price effects ([Comin et al., 2021](#)); (ii) consumption and saving behaviour, including the intertemporal elasticity of substitution ([Browning and Crossley, 2000](#)), wealth accumulation ([De Nardi and Fella, 2017](#)), price rigidities ([Clayton et al., 2018](#)), costs of living ([Orchard, 2022](#)) and marginal propensity to consume ([Andreolli and Surico, 2021](#)); (iii) skill-premium and labour market polarization ([Jaimovich et al., 2020](#), [Comin et al., 2020](#)).³ We depart from these important works along two main dimensions. First, we develop and estimate a structural model in which the cyclicity of non-essential spending and the uneven sectoral distribution of HtM workers make labour demand in non-essential sectors highly pro-cyclical. Second, we use non-homothetic preferences to quantify the contribution of the interaction between cyclical product and labour demands to business-cycle fluctuations and to the transmission of stabilization policies on aggregate consumption.⁴

Structure of the paper In Section 2, we present our measurement strategy spanning across multiple granular datasets, and provide descriptive evidence about the newly con-

³Non-standard preferences (and non-homotheticity in particular) have been used extensively in finance. Notable examples include [Ait-Sahalia et al. \(2004\)](#), [Wachter and Yogo \(2010\)](#).

⁴Further applications of non-homothetic preferences include [Olivi et al. \(2023\)](#), [Schaab and Tan \(2023\)](#) and [Sonnervig \(2023\)](#). Relative to these studies, we provide: broader empirical evidence based on three different empirical settings and source of variation across post-WWII American recessions, U.S. states and following an identified monetary policy shock; new estimates of a business-cycle mechanism that features *both* good market and *labour market heterogeneity*; counterfactual analyses based on an estimated structural model that not only quantify the contribution of each channel to business cycle fluctuations but also highlight a key complementarity between the cyclical spending composition across goods and the *uneven share of HtM workers across essential and non-essential sectors* to amplify the effects of monetary policy.

structured time series for essentials and non-essentials as well as the share of HtM workers in each sector along the income distribution. In Section 3, we describe the identification of monetary policy shocks, lay out the empirical approach based on local projections and report the responses of consumption and earnings, both across essential and non-essential sectors as well as along the labour earning distributions. In Section 4, we develop a structural business-cycle model that features three main ingredients: (i) hand-to-mouth consumers, (ii) non-homothetic preferences over two consumption goods characterized by different IESs, and (iii) heterogeneity in the skill composition of the labour force across industries. In Section 5, we estimate the structural model by minimizing the distance of its theoretical responses to a monetary policy shock from the IRFs estimated via local projections. In Section 6, we use the estimated structural model to perform counterfactual analyses that allow us to identify and quantify the contribution of each channel, (i) to (iii), to amplifying business-cycle fluctuations. Finally, in Section 7, we compare a set of unconventional fiscal policies to the effects of monetary policy. The Appendices provide a comprehensive set of details and further results about our measurement strategy as well as the empirical and theoretical analyses.⁵

2 Data and Descriptive Evidence

In this section, we outline the construction of novel time series of consumption, prices, earnings, and employment for essentials and non-essentials. We proceed in four steps which involves using multiple (micro) datasets from several sources over different samples. First, we classify spending categories into essentials and non-essentials by estimating Engel curves on CEX data. Second, we apply the categorization above to PCE data and obtain indexes of quantities and prices for essential and non-essential spending. Third, we rely on input-output accounts data and the Leontief inverse to group all industries (including those producing intermediate goods) into essential and non-essential sectors. Fourth, we exploit CPS data to compute monthly time series for employment and for several percentiles of the earning distributions in essential and non-essential industries. At the end of this section, we present descriptive statistics that summarize the cyclical properties of our newly constructed time-series, unconditionally. In the next section, we will explore the responses of consumption and earnings in essentials and non-essentials conditional to identified monetary policy shocks. Finally, we use data from the PSID to compute the share of HTM workers in essential and non-essential sectors.

⁵The Online Appendix with additional details and further results is also available here:

https://mandreolli.github.io/files/AndreolliRickardSurico_NEBC_OnlineAppendix.pdf.

2.1 Measurement

The starting point of our data construction exercise is to pin down a precise definition for essentials and non-essentials. It is important to emphasize, however, that nothing of what we describe below hinges upon any specific definition or income elasticity threshold: our method is general enough to accommodate different user’s choices, including the possibility of allowing some spending categories to move between essentials and non-essentials over time.


I) Consumption classification. For the consumption categorization into essentials and non-essentials, we follow [Aguiar and Bils \(2015\)](#), and use data from the Consumer Expenditure Survey (CEX) over the period 1995-1997 to estimate income elasticities of demand for 24 spending groups. We regress expenditure shares at the household level for each category on total expenditure, using net income as an instrument for total spending. Consumption categories with an elasticity to total expenditure greater (smaller) than one are regarded as non-essentials (essentials). The resulting categorization is reported in [Appendix A](#).⁶

II) Building consumption and price series using PCE data. In this step, we map essential and non-essential spending from the CEX to the Personal Consumption Expenditure (PCE) classification by Type of Product from the U.S. Bureau of Economic Analysis (BEA), following [Aguiar and Bils \(2015\)](#). A main advantage of using PCE data is that the BEA produces nominal expenditure, real consumption, and price indices at a very detailed level of disaggregation, and at monthly frequency, consistently since 1959. This allow us to distinguish between movements in quantities and movements in prices for both essentials and non-essentials. In addition, and unlike the CEX, the BEA reports the flow consumption of housing services (e.g. imputed rents for owner occupiers) rather than the actual spending on housing (e.g. mortgage payments), and the former is consistent with the concept used in theoretical models like the one we develop in [Section 4](#).⁷ Finally, we construct Fisher indices for consumption and prices following the approach outlined in the [BEA NIPA \(2021\)](#)

⁶In [Appendix A](#), we report the estimated elasticities of demand for each spending category and provide details of the method in [Aguiar and Bils \(2015\)](#). We use the same consumption classification into essentials and non-essentials over the entire sample, consistent with the evidence in [Aguiar and Bils \(2015\)](#) that the slope of the Engel curve has been relatively stable over time. As discussed in the main text, however, our method can be easily extended to accommodate time-varying Engel curve slopes, and thus allow spending categories to move between essentials and non-essentials over time. Likewise, users may choose to set the elasticity cutoff that separates essentials from non-essentials to a value different from one.

⁷As in [Aguiar and Bils \(2015\)](#), we either adjust or omit from our essential/non-essential classification, spending categories that are likely poorly measured, such as ‘health expenditure’ (which we scale down by the proportion of healthcare spending made out of pocket) or such as ‘professional and financial services fees’ (which we exclude). These adjustments and omissions are detailed in [Online Appendix F.1](#), and result in our essential and non-essential indices covering an average of about 80% of total expenditure over the sample.

Figure 1: Non-essential consumption share over time



Notes: Personal consumption expenditure shares of non-essentials, constructed from chained (2000\$) spending series, and as a proportion of total classified expenditure. Underlying data are from the BEA PCE by Type of Product tables. Shaded areas in grey represent NBER recession dates.

In Figure 1, we report the outcome of these two initial steps in the form of non-essential consumption as a share of total expenditure. This newly constructed series displays two main regularities. First, the non-essential expenditure share has trended upward, moving from about 45% in the early 1960 to 57% in the late 2010s. Second, the share of spending that goes into non-essentials appears to drop significantly and systematically during (NBER) recessions, which are highlighted as grey areas in Figure 1. We will come back to the cyclical properties of our newly constructed series in the descriptive evidence of next section.

III) Mapping consumption to employment data. In our third step, we construct time series of earnings and employment for essential and non-essential sectors. A main challenge is that a large fraction of workers are employed in intermediate industries, and thus we need a strategy to link industry classifications to downstream consumption categories so as to fulfil our goal of identifying how the cyclicalities of final demand affects labour demand in essential and non-essential sectors. We begin by manually classifying each industry code included in

the Current Population Survey (CPS) to the most closely linked final consumption category. As for industries that primarily produce intermediate goods, we use the BEA Input-Output Accounts Data to construct a Leontief inverse that uncovers the contribution of output produced by intermediate industries to each final consumption category.⁸ We classify an intermediate industry as essential (non-essential) if the downstream final consumption it mostly contributes to is essential (non-essential).⁹

IV) Employment and earnings series using CPS data. Given the industry classification above, in the final step of our data construction, we compute employment and earnings series for workers in essential and non-essential industries using the microdata from the CPS. This covers a representative sample of around 60,000 households surveyed monthly, and includes information on the industry in which each worker is employed. We use the IPUMS harmonized CPS industry codes from Flood et al. (2020), which reduce the inconsistency in the NAIC codes over time. Finally, monthly time series for employment are calculated by summing up the total count of workers in each industry that we have classified as either essential or non-essential in the previous step, using the survey weights and the basic sample from the CPS. The two series for essential and non-essential employment begin in 1976. As for earnings, we use data from the Outgoing Rotation Group (ORG), which is a subsample of roughly a quarter of the main CPS sample, to construct monthly series for average earnings per worker and for median earnings. In each month, we also compute the percentiles of the earnings distribution for essential and non-essential sectors, respectively, based on the weights and the weekly earnings reported by the CPS.¹⁰ The earnings series for essential and non-essential sectors begin in 1982, are deflated using the overall PCE price index in 2012\$, and are seasonally adjusted by the Census Bureau’s X-12-ARIMA Seasonal Adjustment procedure.

It is worth emphasizing that our proposed classification into essentials and non-essentials

⁸Details on the industry classification and the mapping of intermediate goods sectors into final expenditure of essentials and non-essentials are outlined in Online Appendix F.3.

⁹An alternative approach is to construct employment series using the *shares* of downstream production that is essential or non-essential, rather than a binary classification approach. We prefer the binary approach for its simplicity, but have verified that our results are robust to the alternative approach.

¹⁰In our data construction, we combine the mean earnings per worker from the ORG subsample of the CPS with total employment from the full sample of the CPS to calculate monthly earnings for: (i) the whole U.S. economy, (ii) non-essential sectors, and (iii) essential industries. The ORG sample and weights, however, are designed to be representative of the U.S. population at quarterly frequency, and not necessarily at monthly frequency. To ameliorate any possible concerns regarding representativeness, we provide two pieces of evidence. First, in Online Appendix Figure F.2, we show that the response of our newly constructed aggregate earning series from the CPS to an identified monetary policy shock is very similar to the response of total compensation of employees from the BEA. Second, in Online Appendix Figure F.3 we have verified that our results are not overturned if we aggregate overall earnings in the CPS at quarterly frequency, instead.

is conceptually and quantitatively different from the more traditional divide between durables and non-durables. First, on average over our sample, non-essentials account for more than 50% of household expenditure whereas the share of durable purchases is typically around 15% only. Furthermore, not only the vast majority of non-essential spending consists of non-durable consumption but also about half of non-durable consumption is spent on non-essential goods and services. In Appendix B, we present extensive evidence in support of the notion that the essentials/non-essentials classification is very different (and far more pervasive) from the distinctions between durable/non-durable, goods/services and tradeables/non-tradables.

Share of hand-to-mouth workers across essential and non-essential sectors. We use data from the Panel Study of Income Dynamics (PSID) to compute the share of hand-to-mouth households across essential and non-essential industries along the income distribution. The PSID provides comprehensive income and balance sheet data for a representative sample of 17,280 households, surveyed biennially. It also includes information on the industry of employment for each household member. In any given wave, households are identified as Hand-to-Mouth (HtM) if their liquid assets are less than their monthly household income. Following Kaplan et al. (2014), liquid assets are calculated by summing checking and savings accounts and other financial assets (excluding retirement accounts) and subtracting liquid debt. Household income is constructed by summing the labour income of both partners, as well as household income from own business and government transfers, excluding social security.¹¹

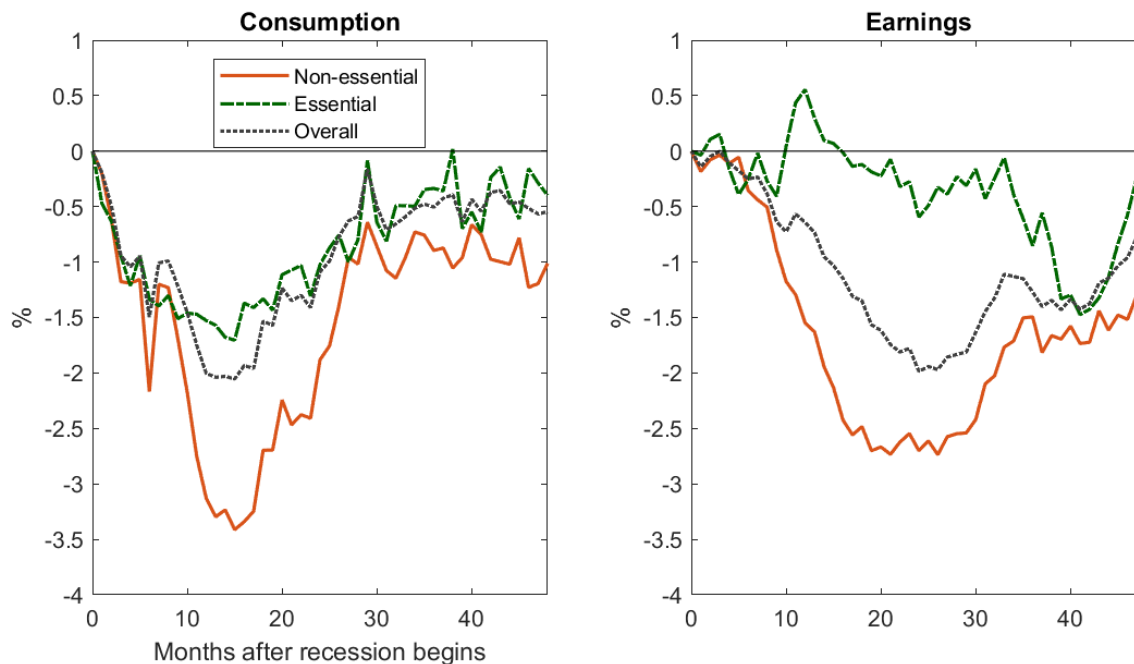
In each PSID wave, we group households by income deciles based on their total family income. Then, we assign households to essential and non-essential industries according to the occupation of the household worker with the highest labour income. The PSID classifies industries using Census 3-Digit Codes, which allows us to link this data with our classification of essential and non-essential industries in a straightforward way. For each wave, we compute the share of hand-to-mouth households in essential and non-essential sectors along the income distribution, and then compute the averages by sector/decile across all waves. We apply the same procedure to construct shares of HtM for durable and non-durable industries.

¹¹The original classification in Kaplan et al. (2014) relies on data from the Survey of Consumer Finances (SCF). We focus instead on the Panel Survey of Income Dynamics because it also reports detailed information on employment and industry occupation per family member. These were first available in the wave of 2003 (Beraldi and Malgieri, 2024). Accordingly, our PSID sample covers the period 2003-2021. In Appendix Section F.8, we provide more details on the specific PSID variables that we use.

2.2 Unconditional correlations

In the previous section, we have noted that the share of non-essential consumption tends to drop systematically during recessions (Figure 1). In this section, we present descriptive statistics that highlight the cyclical properties of consumption and earnings in essentials and non-essentials. For each newly constructed series and for each recession, we compute the percentage change from the local peak to the log-level of each subsequent 48 months. This yields a set of 48 observations after each peak, which we average across all recessions in the sample. The findings of this exercise are reported in Figure 2. The left panel refers to consumption while the right panel to earnings. Solid lines in orange represent non-essentials, broken lines in green stand for essentials while dotted lines in black are for the whole economy.

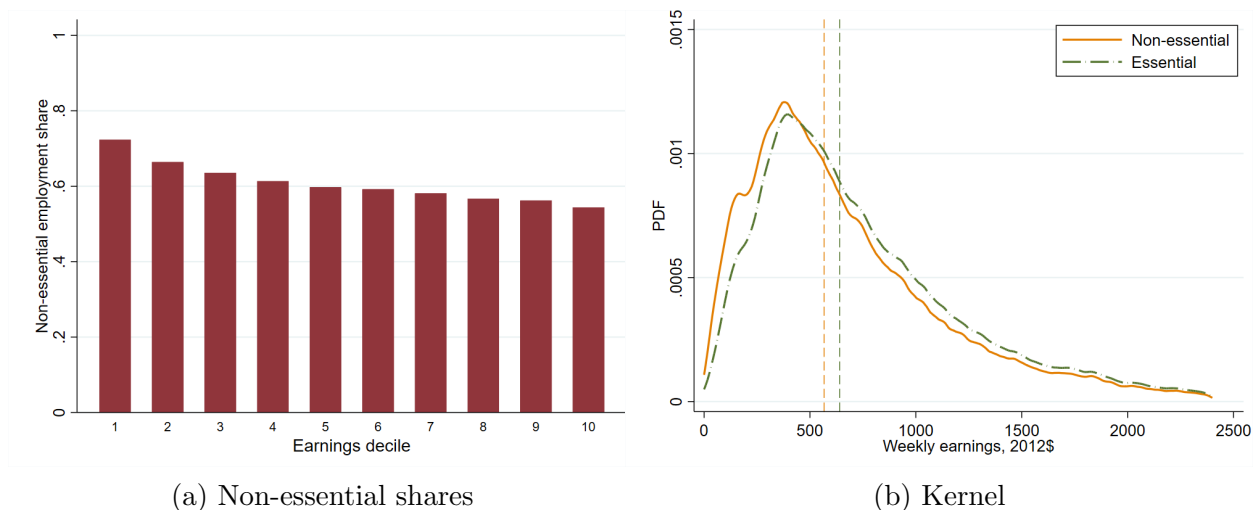
Figure 2: Sensitivity of Essentials and Non-essentials to Recessions



Notes: Response of essential and non-essential series after the start of a recession. To construct these responses, we first log and detrend the series using the HP filter ($\lambda = 14,440$). For the earnings series, we report a 6-month moving average to reduce noise. We then calculate the average decline in the series after all recessions, as defined by the NBER, between 1973-2007, which the data are available for.

Three main take-aways emerge from Figure 2. First, the consumption of non-essentials is far more sensitive to the business-cycle than its essential counterpart (left panel). Non-essential spending drops by almost 3.5% after one year since the inception of the average U.S. recession whereas essential spending only falls by 1.5%. The gap is still significant four years after the peak, with non-essential spending, at -1%, more than doubling the shortfall in essentials. Second, the heterogeneity in earnings is even more pronounced than

Figure 3: Non-essential and essentials across the earnings distribution




Notes: Earnings distributions within essential and non-essential industries. Underlying data is pooled Jan 1982 - December 2020, from the CPS, as described in the text. Panel (a) shows the percent of employees working in non-essential industries for each decile of the income distribution (deciles computed annually). Panel (b) shows the kernel density plot along the median of each distribution

in consumption: two years into a typical recession, and earnings in the non-essential sectors still witnesses a dramatic 2.5% fall against the backdrop of a more gentle -0.3% in essentials (right panel). Third, looking at the aggregate series in dotted black masks the pervasive heterogeneity across essentials and non-essentials, with the latter being a main driver behind the aggregate results. The result in the right panel of Figure 2 chimes with the findings in Ahn et al. (2024), who document that workers in industries that we classify as non-essential have a much higher probability of transiting out of employment, especially during recessions.

Motivated by Figure 2, we zoom into the distribution of earnings within sectors. In the left panel of Figure 3, we report the share of employment in non-essential sectors across the earning distribution deciles. This decays monotonically, from a value shy of 75% in the bottom decile to a number below 55% at the top. The right panel of Figure 2 plots the estimated kernel density of wages across the two sectors. The distribution of earnings in non-essential industries is always to the left of the distribution in non-essential industries, with median earnings recording a 12% gap relative to their essential counterparts. Putting all the pieces together suggests that: (i) low-income workers are more likely employed in non-essential sectors, (ii) non-essential workers tend to be paid less than their essential counterparts. In Appendix Figure F.4, we show that non-essential earnings are first order stochastically dominated by essential earnings.

Finally, in Figure 4, we report the shares of HtM workers in industries that produce

Figure 4: Share of Hand-to-Mouth workers in essential versus non-essential sectors




Notes: Panel Study of Income Dynamics (PSID). Sample: 2003-2021.

essentials (blue bars with dotted outlines) and non-essentials (red bars with solid outlines), by income decile in each sector. The chart reports statistics computed as the average of all PSID waves between 2003 and 2021. Two main findings emerge from Figure 4. First, non-essential industries are characterized by a share of HtM workers that is more than the double the share of HtM in essentials: 23% vs 11%, on average across income deciles. Second, the sectoral *gap* is far more pronounced at the bottom of the income distribution (e.g. 26% and 20% in the first two deciles) and virtually disappears among higher earners (e.g. 6% and 4% in the last two deciles). In sharp contrast, Appendix Figure B.3 reveals that the share of HtM workers in industries that produce durable goods: (i) is significantly smaller than the share of HtM workers in non-durable sectors (i.e. 5% vs 29% on average); (ii) varies little by income deciles. We interpret this as further evidence of a marked difference in the labour force composition of non-essential sectors and durables industries.

State-level analysis. To complement the results above, which are based on variation over time and across deciles of the income distribution, in this section we use state-level data that exploits variation over both time and space to estimate the reduced-form relationship between non-essential spending and employment during recessions. More specifically, we use annual PCE data to construct average non-essential consumption shares for each U.S. state, since

Figure 5: State-level employment during recessions vs non-essential consumption shares



Notes: Seasonally adjusted monthly state level employment series is from the BLS. Employment declines are calculated around state-specific timing of national recessions (see main text), and averaged across all recessions since 1976. Non-essential shares are consumption shares constructed from the BEA’s state-level annual PCE expenditure (by type of product) series, deflated using the corresponding national PCE price indices shares from the BEA’s state-level annual PCE series, averaged over the period 1997-2021.

1997. We then calculate the average decline in employment during all recessions since 1976, using the timing of changes in state-level employment around nationally defined recessions to identify state-level downturns.¹² The analysis in Figure 5 reveals that, on average across all recessions in our sample, U.S. states with a higher share of non-essential consumption experience a larger decline in employment. This holds true not only over the full-sample but also during the most recent post-1990 downturns (independently of Covid), suggesting that the correlation between non-essential spending and employment is stable over time.

3 Empirical Framework

In the previous section, we have used reduced-form correlations to document a number of empirical regularities: non-essential spending and non-essential earnings fall far more than their essential counterparts during recessions; non-essential wages tend to be lower than essential salaries; the share of HtM workers in non-essential is markedly larger than in essentials. In this section, we corroborate these findings by using an identified monetary policy shock and tracing out the responses of consumption, earnings and prices to an unanticipated increase in the interest rate. While focusing on a specific shock has the advantage of allowing us to distinguish correlation from causation, the evidence of the previous section suggests that the

¹²In Online Appendix F.10, we provide details on the data construction for the state-level analysis.

findings in this part of the paper may extend also to other identified business-cycle shocks. In the next section, we will develop and estimate a structural model that can account for the effects of monetary policy on essentials and non-essentials documented in this section.

3.1 Identification and Estimation

Before presenting our empirical results, we briefly discuss the identification strategy and the empirical framework. Both are borrowed from the state-of-the-art and therefore, for full details and motivations, we refer the interested readers to the original contributions by [Gurkaynak et al. \(2005\)](#), [Gertler and Karadi \(2015\)](#) and [Jordà \(2005\)](#).

Monetary policy shocks. To further investigate the dynamic effects of business-cycle fluctuations on essentials and non-essentials, we need to identify plausibly exogenous variation in a macro variable that can affect the entire economy. In our case, over and above any reverse causality concern, the challenge is complicated by the fact that we also need to make sure that the identified shocks do not originate from either the essential or the non-essential sector, otherwise it would be hard to attribute any possible heterogeneous response to differences in demand elasticities across the two types of goods as opposed to asymmetry in the shocks. Monetary policy surprises appear to fulfil both requirements. More specifically, we follow the High-Frequency Identification (HFI) of monetary policy shocks of [Gertler and Karadi \(2015\)](#), who in turn build on [Gurkaynak et al. \(2005\)](#), which measures changes in Fed Funds futures over a short window of time, typically 30 minutes, around monetary policy announcements. This provides plausibly exogenous variation in interest rates under the identifying assumption that any information about macroeconomic conditions that could have potentially affected the endogenous response of monetary policy has actually been already anticipated by financial markets. The exclusion restriction is that any variation in Fed Funds future prices during the short-time window around policy announcements is only due to differences in monetary policy decisions from financial market expectations.¹³

Econometric method. The monetary policy instrument that is used for the high-frequency identification is available for period 1990 to 2016. However, an extended monetary shock series can be produced by estimating a proxy-VAR as in [Mertens and Ravn \(2013\)](#) and [Stock and Watson \(2018\)](#) over a longer sample, and then identifying the monetary policy surprise

¹³A more recent empirical literature has further refined this high-frequency identification by isolating also the ‘information effect’ that may also be contained in meeting announcements if the central bank has private information about the state of the economy relative to financial market participants. In one of the robustness exercises at the end of this section, we obtain very similar results when the analysis uses alternative shocks.

series using the HFI monetary policy instrument over the shorter period over which is available. In practice, we extract the monetary shocks by estimating a proxy-VAR similar to the one of [Gertler and Karadi \(2015\)](#) on the sample January 1973 to December 2020, using the 1y government bond yields, the excess bond premium, the first difference of log industrial production, and the first difference of log PCE price index. We include the monetary policy surprises as an internal instrument. We use twelve lags for the endogenous variables and four lags of the instrument. The extracted monetary policy shocks are reported in Online Appendix Figure [G.1](#).

To check for weak instruments in our specification, we run the weak instruments test proposed by [Montiel Olea and Pflueger \(2013\)](#). The critical value for the test is 12.039, assuming a 5% confidence and worst-case bias of 30%. Using the shocks identified à la Gertler-Karadi, the corresponding robust F-statistic is 13.89, passing the weak instrument test. When using the shocks identified à la Jarocinski-Karadi, however, the F-stat lowers to 10.29, which is below the critical value and thus suggests a possibly weak instrument. Accordingly, we use the Gertler-Karadi shocks as our baseline case and report the results using Jarocinski-Karadi’s refinement in the Appendix as a robustness check for the information effect.

The impulse response functions to a monetary policy shock are computed using the smooth local projection instrumental variable (SLP-IV) estimator of [Barnichon and Brownlees \(2019\)](#) on the following sequence of local projections, as developed by [Jordà \(2005\)](#):

$$y_{t+h} = \alpha_{h,0} + \beta_h 1y \text{ yield}_t + \sum_{l=1}^L Y_{t-l} \gamma_{h,l} + \epsilon_{t,h} \quad (1)$$

The dependent variables y are, in turn, the logs of our newly constructed series for essential, non-essential, and aggregate measures of consumption, prices, employment and earnings. The coefficients β_h s are the object of our interest, as they summarize the impulse responses of the y s at each horizon h to an unanticipated 100bp increase in the one year government bond yields (1y yields). This is instrumented with the series of monetary policy shocks extracted from the proxy-SVAR.

The local projections in [\(1\)](#) are estimated with SLP-IVs over a forecast horizon h of up to four years, using the five-fold cross-validation selection of the smoothing parameter recommended by [Barnichon and Brownlees \(2019\)](#). Standard errors are computed applying the [Newey and West \(1987\)](#) correction. To maximize the number of observations, all samples end in December 2019 (so as to avoid any contamination from Covid) but the starting point varies slightly with the availability of the dependent variable: this is 1973 for consumption; 1978 for prices; 1976 for employment; and 1982 for earnings. In all specifications, we include


as controls the first 12 lags of all variables in the VAR (1y yields, the excess bond premium, log industrial production, and the log PCE price index) as well as 12 lags of aggregate or sectoral consumption and labour market variables, all in logs. Each model features additional model-specific controls, such as 12 lags of the dependent variable, in an effort to balance the trade-off between the benefits of lag-augmentation and the cost of over-fitting. Details of the smoothed local projection IV estimation as well as the full list of controls for each specification are reported in Appendix C and the Online Appendix G.

3.2 Results across spending categories and sectors

In this section, we employ the empirical framework of Section 3.1 to estimate the effects of monetary policy on the newly constructed data of Section 2. The main results are presented in Figure 6. This shows the Impulse Response Functions (IRFs) for consumption (top row) and total earnings (bottom row). The first column refers to the response of the aggregate variable for the whole U.S. economy, and this is what has been typically featured in earlier empirical studies. The second and third columns record the IRFs of essential and non-essential, which is a main contribution of our paper. The fourth column reports the IRFs of the (log-)ratio between non-essentials and essentials in each row, and therefore any significant effect in that column can be interpreted as a rejection of the null hypothesis that the responses of essentials and non-essentials for consumption (in the top row) and for earnings (in the bottom row) are the same. Solid (dashed) blue lines refer to 68% (90%) confidence intervals. Red dots and crosses refer to the IRFs of the estimated structural model of Section 4, which will be discussed in Section 5.

Three main results emerge from Figure 6. First, the aggregate effects displayed in the first column resemble the findings in earlier empirical work. After a 100bp interest rate hike, consumption expenditure falls significantly, up to -0.8% , before the changes become insignificant three years after the shock. The response of income is delayed but larger, peaking shy of -4% , and reverting to values not statistically different from zero by the end of the forecast horizon. Second, the aggregate effects in the first column average (and therefore mask) sizable heterogeneity in the middle two columns, both across sectors and variables. More specifically, the decline in non-essential spending in the third column of the top row is about two times as large and persistent as the decline in essential spending in the second column, echoing the findings in Orchard (2022). But the sharpest heterogeneity emerges in the bottom row: the decline in non-essential earnings peaks significantly, in excess of -4% , whereas the drop in the earnings of the essential sectors is insignificant, and never exceeds -1% . Third, the responses of essential and non-essential, for both consumption and earnings,

Figure 6: IRFs to contractionary 100bp monetary policy shock - Consumption and Earnings



Notes: Blue lines are empirical impulse response functions (IRFs) to a 100bp increase in the 1y year government bond yields estimated by smooth local projections instrumental variable, where the instrument is the monetary policy shocks derived from high-frequency identified monetary policy surprises. Confidence intervals are reported at the 90% (dashed line) and 68% (solid line) level. Sample periods and controls for each column are specified in the main text and Appendix C.1. Red markers refer to quarterly IRFs from the estimated structural model of Section 4. Red “X”s correspond to variables that have been targeted in the structural estimation whereas red “O”s stand for variables that have not been targeted.

are statistically different one from another, as exemplified by the significant IRFs in the last column.

In summary, during a (monetary-policy induced) recession, households are more likely to cut back on non-essential spending. As non-essential industries face a more cyclical demand, these sectors also tend to reduce wage payments significantly, whereas essential industries do not, possibly reflecting the lower sensitivity of their demand to the business-cycle: the responses of non-essentials drive the aggregate results, both for consumption and earnings. In Appendix C.2, we document significant heterogeneity in the responses of both (median) earnings per worker and employment, with a possibly more pronounced contribution of the latter (i.e. the extensive margin) to the effects on total earnings in Figure 6. Finally, we find mild evidence of essential and non-essential prices responding to these relative demand shifts: overall prices fall slightly, as a result of a larger negative movement in non-essential prices and a smaller positive change in essential prices.¹⁴ The sectoral price responses, however, are insignificant at the 90% confidence level, thereby suggesting that the general equilibrium channel through prices is unlikely to be very strong in post-WWII U.S. data.


3.3 Results along the earnings distribution

In the previous sections, we have shown that a (monetary policy-induced) recession causes: (i) households to reduce their non-essential spending more than their essentials, and (ii) firms in non-essential industries to cut their labour demand more than in essentials. In Section 2, we have further shown that (iii) non-essential workers are more likely hand-to-mouth. In this section, we want to go at the heart of the triple interaction emphasized in the introduction and ask: do the labour earnings of low-income (and thus more likely HtM) workers in non-essential (i.e. more cyclical) industries fall more than the wages of (low-income) employees in essential sectors, after a contractionary monetary policy shock? We find that they do, very significantly. In the next section, we show that the triple interaction of hand-to-mouth workers with more cyclical salaries, and employed in more business-cycle sensitive industries provides a powerful, yet overlooked, amplification mechanism in an estimated model for business-cycle analysis.

To build an answer to the question above, in the top row of Figure 7, we display the labour earnings response along the earnings distribution in each sector, across the forecast horizon of up to four years after the monetary policy shock. The bottom row zooms on the

¹⁴The heterogeneity in the price responses across essentials and non-essentials is consistent with the evidence in Stock and Watson (2020) that the slope of the Phillips curve is steeper in sectors with more cyclical demand. In Online Appendix F.5, we also report the IRFs of the other aggregate series included in the proxy-SVAR. These are estimated using SLP-IV and are similar to those in the literature.

Figure 7: IRFs to contractionary monetary policy shock - Earnings distribution



Notes: Empirical impulse response functions (IRFs) to a 100bp increase in the 1y year government bond yields estimated by smooth local projections instrumental variable, where the instrument is the monetary policy shocks derived from the [Gertler and Karadi \(2015\)](#) high-frequency monetary policy surprises. Sample periods and controls for each column are specified in the main text and Appendix C.1. Earnings percentiles are from the CPS, and percentiles are calculated separately for the non-essential and essential earnings distributions. Solid (dashed) blue lines refer to 68% (90%) confidence intervals.

three-year forecast horizon and report point estimates as well as confidence intervals for the earnings response across the earnings distribution at that particular horizon. Three main results can be inferred from this exercise. First, the heterogeneity in the earnings responses across the income distribution within essential industries is modest, both economically (top row) and statistically (bottom row). Second, the earnings responses at the bottom deciles of the earning distribution of non-essentials is significantly larger than the responses at the top deciles. Third, the salaries of low-income workers in non-essentials fall between two and four times more than the salaries of low-income workers in essentials. In other words, Figure 7 reveals that earnings cyclicality is higher for the low-income workers employed in more cyclical industries.

The results in this section may also contribute to account for the counter-cyclical income inequality reported by [Heathcote et al. \(2010\)](#). The higher sensitivity of salaries in non-essential sectors (especially, at the bottom of the earning distribution) and the finding that non-essential wages tend to be lower than in essential industries suggests that earnings inequality may increase during recessions as a result of the labour market responses to the higher cyclicality of non-essential demand. The reason is that low-income workers in non-essentials lose out twice in bad times: not only they are worse paid than essential employees over the business-cycle but also, during recessions, their earnings decline by more.¹⁵

3.4 Additional results

In this section, we show that our results are not driven by confounding factors, such as the cyclicality of durables, tradeables or other types of goods, and that are robust to several sensitivity checks, including a different identification of either monetary policy or business-cycle shocks, different (rectangularized) samples, quarterly frequency of the data and an alternative estimation method. Further details are provided in Appendices [B](#), [C](#) and [H](#).

Alternative expenditure classifications. A possible challenge to our interpretation of non-essentials as highly cyclical industries is that the heterogeneity in income elasticities of demand could be correlated with other product characteristics which may also account

¹⁵The results in this section suggest that the general equilibrium effects of the heterogeneity in the labour market responses may dampen the heterogeneity in the consumption responses. The reason is that the higher cyclicality of non-essential spending has a particularly negative effect on the pay of low-earners in that sector. But those are also the households who not only have a larger MPC but also who spend a higher share of their budget on essentials. Accordingly, the fall in essential spending that we estimate in Section 3.2 may have been indirectly amplified by the decline in non-essential earnings. This chimes with the evidence in [Coibion et al. \(2017\)](#), who document that a contractionary U.S. monetary policy shock causes a larger increase in earning inequality than in consumption inequality.

for the cyclicalities across essentials and non-essentials that we have documented above. In Appendix B, we discuss extensively two popular alternative spending classifications, such as durables versus non-durables and tradeables versus non-tradeables, and show that these are unable to account for the higher sensitivity of non-essential spending to business-cycle fluctuations. In short, while about 78% of durable goods are non-essentials, there are about 50% of non-durables goods and services that are also non-essentials (and indeed display a far higher cyclicalities than the other half of non-durables). Furthermore, durable goods are a relatively small proportion of overall consumption (less than 15%) compared to non-essentials (more than 50%) and they barely contribute to changes in aggregate income. In addition, the positive correlation between income level, being HtM, and the demand elasticity that is crucial for the labour market amplification documented in this paper is weak for durables: non-essential workers are typically paid less than essential workers and are more likely Hand-to-Mouth, whereas durable sectors display higher wages and a lower share of HtM workers than non-durables. Finally, we find that also the distinction between tradeables and non-tradeables bears little correlation with the distinction between essentials and non-essential, and therefore it also seems an unlikely confounding factor behind the evidence in this paper.

Monetary policy shocks and specification. Our baseline specification employs an (updated) series of shocks from Gertler and Karadi (2015), within a proxy-SVAR. In Appendix C.3 we demonstrate that our main findings are dependent on neither this particular shock nor the specification choice. First, we verify that our main results also hold true using a variety of alternative monetary policy shocks proposed in recent literature.¹⁶ Secondly, we also verify our results using a proxy-SVAR specification, which uses the same sample period and a similar set-up to our baseline smooth local projection approach. We report estimated IRFs for essential and non-essential consumption and earnings. The vast majority of these IRFs fall within the 90% confidence interval of our baseline results, although the magnitude of the responses depends somewhat on the choice of monetary policy shock. Moreover, we find the *relative* response of non-essentials is consistently larger than that of essentials, conditional on the choice of shock and specification. This is true of both earnings and consumption.

Business-cycle shocks. To separate correlation from causation, our main results are based on a well-understood and widely used source of business-cycle variation, namely monetary policy shocks. The findings in Section 2.2, however, suggest that our mechanism may apply more broadly to other types of business-cycle shocks, as the heterogeneity across essentials

¹⁶Jarociński and Karadi (2020), Miranda-Agrippino and Ricco (2021), Romer and Romer (2004), Bauer and Swanson (2023) and Aruoba and Drechsel (2024).

and non-essentials, both within and across spending and earnings, emerges also when we look at the ‘unconditional’ impulse responses of Figure 2, which shows the percentage change in consumption and wage bills from the incept of a typical recession in the post-WWII period for the United States. To further investigate the breadth of our mechanism beyond the transmission of monetary policy, in Appendix C.4, we report the impulse responses of consumption, prices and earnings (both in the aggregate and across essential and non-essential sectors) to the business-cycle shocks identified by Angeletos et al. (2020). The idea is to isolate the shock that explains the maximum share of variation in unemployment at business-cycle frequencies between 6 and 32 quarters, in a multi-equation system like a VAR. The estimates of this exercise in Appendix C.4 are qualitatively very similar to the findings in Figure 6, despite the two sets of IRFs are based on very different identification strategies.¹⁷

Further sensitivity. In Online Appendix H.2, we show that our results are robust, if not stronger, when ending the sample in December 2020, and therefore including the effects of the Covid-19 Pandemic. Part of the effects estimated over this extended sample, however, may reflect a mechanical correlation between spending and earnings, simply because low-income workers and non-essential sectors were more likely to be in lockdown (Blundell et al., 2020). Accordingly, we exclude 2020 from our baseline estimates and keep the inclusion of Covid as a robustness check. In Online Appendix H.3, we show that using unsmoothed local projection instrumental variables produce qualitatively similar results, though the point estimates are more jagged and less precise. In a recent contribution, however, Montiel Olea and Plagborg-Møller (2021) recommend using smoothed local projection as an efficient way to shrink a potentially over-parameterized model with many variables and lags, and therefore we keep the latter as our baseline specification. Finally, in Online Appendix F.6, we display the impulse responses of earnings based on quarterly data. This is an important cross-check for our monthly estimates because the Current Population Survey makes clear that representativeness of their sample is ensured only at quarterly frequency. The estimates of this exercise, which are reported in Online Appendix F.6, are qualitatively similar to the estimates in Figure 6, though far less accurate, possibly reflecting the smaller variability and the lower number of observations associated with the quarterly sample. We also find that using a rectangularized sample period across all variables and specifications, which corresponds to beginning all estimation samples in January 1982 (as opposed to maximizing information by using different samples for variables with longer data availability), produces very similar results, which we make available upon request.

¹⁷We also find similarly large declines in non-essential consumption and employment when estimating Generalised IRFs. Results are available upon request.

4 A Model of Cyclical Demand Composition

In the previous sections, we have shown that —during recessions— spending and earnings in non-essentials fall significantly more than their essential counterparts. In this section, we develop a structural model with non-homothetic preferences in consumption and heterogeneity in labour productivity that can account for this evidence. We add three dimensions to an otherwise standard business-cycle model with nominal rigidities and heterogeneous agents. First, households consume two types of goods, which differ for their demand elasticities: essentials and non-essentials. Second, workers are characterized by either high productivity (and hence enjoy high-income and face no financial constraint) or low productivity (and thus have low-income and are HtM). Third, non-essential industries employ a higher share of low productivity/HtM workers, consistent with the evidence in Section 2.2. In the next section, we estimate this structural model and show that contractionary monetary policy encourages households to cut their non-essential consumption. This particularly affects low-income/HtM families, whose workers are more likely employed in non-essential industries. As low-income/HtM households have a higher spending propensity, this non-essential channel amplifies business-cycle fluctuations through a general equilibrium effect. In Section 6, we use the estimated model to run counterfactual simulations that quantify the extent of amplification relative to specifications with either no heterogeneity in spending, no heterogeneity in earnings or no heterogeneity in the composition of the labour force across sectors.

4.1 Non-homothetic preferences

Our starting point are consumers’ preferences which allows us to think about spending categories that may be characterized by potentially different elasticities of demand. For this purpose, we introduce a non-homothetic utility function that builds on the partial equilibrium, finite horizon analysis of [Browning and Crossley \(2000\)](#). In [Andreolli and Surico \(2021\)](#), we study the implication of non-homothetic preferences for heterogeneity in MPCs across households.

Within each period, households with skill/productivity level i ($i = H, L$) receive additively separable flows utility from spending on two categories of consumption: essentials (C^E) and non-essentials (C^N). They also receive disutility from supplying labour (N):

$$U(C_{i,t}^E, C_{i,t}^N, N_{i,t}) = \frac{(C_{i,t}^E)^{1-\frac{1}{\gamma^E}}}{1-\frac{1}{\gamma^E}} + \varphi \frac{(C_{i,t}^N)^{1-\frac{1}{\gamma^N}}}{1-\frac{1}{\gamma^N}} - \xi \frac{N_{i,t}^{1+\chi}}{1+\chi} \quad (2)$$

where χ is the inverse of the macro Frisch elasticity, φ and ξ are scaling constants that

will help calibrate the steady state solution, while γ^E and γ^N are the category-specific Intertemporal Elasticity of Substitution (IES) for essentials and for non-essentials, respectively. [Browning and Crossley \(2000\)](#) show that there is a one-to-one mapping between the spending category-specific IES and the Income Elasticity of Demand (IED) for that type of goods and services in a two-period model; in [Andreolli and Surico \(2021\)](#), we extend that result to an infinite horizon setting. In that paper, we also show that the consumption category with the highest IES also exhibits the highest IED, implying that —by definition— the ranking of γ^s distinguishes essentials from non-essentials:

$$\gamma^E < \gamma^N \tag{3}$$

Intuitively, households are very willing to smooth the consumption of necessities such as groceries, utilities, transportation, personal care or health care. In contrast, consumers are willing to delay spending on food away from home, hospitality, large durables, fitness or education: luxuries are easier to postpone (or move forward) than necessities. Note that the mapping between income elasticity and intertemporal substitution is not an artefact of these preferences but it is a more general property of non-homotheticity. [Browning and Crossley \(2000\)](#) prove that this holds for any time additively separable utility function in good varieties.¹⁸

A main insight of our theoretical model is that the cross-sectional concept of income elasticity of demand and the time-series concept of intertemporal elasticity of substitution can be intimately related and be key to understand non-essential business cycles: non-necessity spending responds more strongly to shocks for both high-income and low-income households. Following a contractionary monetary policy shock, Ricardian agents mainly move along their Euler equation as they want to postpone non-essentials, whereas hand-to-mouth agents adjust their consumption along the Engle curve, reducing non-essential spending in response to income losses.

The specification in (2) has a few other attractive properties. For instance, households with a lower income spend a relatively larger fraction of their budget on essentials, due to the lower income elasticity of demand for essentials. Furthermore, the intertemporal elasticity of substitution is higher for wealthier households, consistent with the evidence in [Crossley and Low \(2011\)](#).¹⁹ Moreover, these preferences are a simple extension of the standard CRRA

¹⁸Note that the non-homothetic CES preferences of [Comin et al. \(2021\)](#), while not separable intratemporally in goods also imply that non-essentials have a higher elasticity of intertemporal substitution than necessities, as they are intertemporally additive. Our preferences incorporate this feature in a straightforward manner, only using two parameters, γ^E and γ^N , which can be identified from macro data.

¹⁹[Stiglitz \(1969\)](#) show how non-homothetic preferences are linked to risk aversion while [Ait-Sahalia et al.](#)

used in business-cycle analyses and therefore they can be easily compared to (and embedded in) existing models. Finally, while these preferences are not aggregable, we do not regard this as a major limitation for our purposes. The reason is that we are primarily interested in modelling spending heterogeneity and in eliciting a mapping from IEDs to IESs, which in turn allows us to quantify the contribution of non-essentials to business-cycle fluctuations. In our view, this benefit exceeds any potential cost of being unable to derive an aggregate Euler equation for the phantomatic representative agent.²⁰

4.2 Households problem

Households i have an instantaneous utility for essentials, $C_{i,t}^E$, and for non-essentials, $C_{i,t}^N$, and instantaneous dis-utility for working $N_{i,t}$ hours. They are also inattentive, as in [Mankiw and Reis \(2007\)](#). Households update their expectations sporadically, with probability λ . Anyone who updates their expectations today has a probability λ of updating them tomorrow, $\lambda(1-\lambda)$ of updating them in two periods, $\lambda(1-\lambda)^2$ in three periods, $\lambda(1-\lambda)^j$ in $j+1$ periods, and so on. As in [Beraja and Wolf \(2021\)](#), household inattentiveness is introduced to match the hump-shape response of consumption (while preserving the differential spending category-specific IES).

As households realise that they might not be able to update, they make plans for future choices in the current period. They choose consumption of a variety, say essentials, for today, $C_{i,t,0}^E$, and for the future if they do not update for j periods ahead, $C_{i,t+j,j}^E$. The same applies to non-essentials. A perfectly competitive union frictionlessly sets wages for households, implying that the choice of hours is not affected by household inattention, as in [Mankiw and Reis \(2007\)](#). Unlike these authors, however, households make plans for two separate consumption goods.

The economy is populated by two types of agents: high-skilled, H , and low-skilled, L . They differ along two dimensions: steady state income levels and whether they can access financial markets. A large empirical literature on survey and administrative data has made the case that low-income households exhibit high marginal propensity to consume (e.g. [Johnson et al., 2006](#)). Accordingly, we assume that H agents have higher income and are Ricardian, whereas L agents have lower income and are hand to mouth.²¹ High-earning agents are paid

(2004) use a version of (2) to rationalise the equity premium puzzle, in a combination of the volatility in the luxury spending of the rich and their consumption-specific risk aversion.

²⁰Other strands of the literature, especially on structural transformation, use aggregable preferences (e.g. the PIGL preferences in [Boppart, 2014](#)), as these are helpful for modelling a balanced growth path.

²¹This simplification would arise endogenously in a heterogeneous agent model with uninsurable income risk and borrowing constraint. The framework can be easily extended to include wealthy hand-to-mouth agents, but we abstract from this here, both for tractability and to highlight the new channel that we propose.

an average wage $W_{H,t}$, while low-earning agents face a salary $W_{L,t}$. Households also obtain profits from firms, $\Pi_{i,t}$, and transfers from the government, $T_{i,t}$. We present the derivation of the household and the union problem in Appendices [D.1](#) and [D.2](#).

4.3 Firms

There are two sets of firms or industries: those that produce essentials and those that produce non-essentials. The two sectors differ in the skill composition of their labour force, with non-essential industries employing a relatively higher share of low-skilled/HtM workers, consistent with the evidence from the PSID in [Section 2](#). Each set of firms consists of three separate entities: a final good producer, a Calvo retailer, and a wholesaler.

Final good producers. Final good producers combine different retail varieties of essentials and non-essentials according to a CES aggregator. This leads to a standard demand facing final good producers for different varieties of either essentials or non-essentials:

$$y_{k,t}^i = Y_t^i \left(\frac{P_{k,t}^i}{P_t^i} \right)^{-\varepsilon} \quad i = \{E, N\}$$

Calvo retailers. Retailers of essentials buy a wholesale essential good at a wholesale price $P_t^{E,w}$ and use it to produce the retail variety $y_{k,t}^E$ with a linear technology that maps one-to-one the wholesale good to the retail variety. As each variety is differentiated, producers have market power and face a Calvo friction to change prices. Their real marginal cost $\mathcal{S}_t^E = \frac{P_t^{E,w}}{P_t^E}$ is the wholesale price relative to its retail average value. Firms receive a subsidy τ^E for each unit of good produced and pay lump sum taxes T_t^E ; these taxes allow them to have zero profit in steady state but do not affect the profit allocation off-steady state. The probability of not being able to reset prices is equal to θ in each period. This leads to a standard non-linear New-Keynesian Phillips Curve. The non-essential retailers problem is fully symmetric.

Wholesalers. These produce one type of good, either essential or non-essential, under perfect competition, and combine high-skill, $N_{H,t}^i$, and low-skill labour, $N_{L,t}^i$, with technology:

$$\begin{aligned} Y_t^E &= A_t^E (N_{L,t}^E)^{\alpha^E} (N_{H,t}^E)^{1-\alpha^E} \\ Y_t^N &= A_t^N (N_{L,t}^N)^{\alpha^N} (N_{H,t}^N)^{1-\alpha^N} \end{aligned}$$

Wholesalers sell goods to retailers at nominal price $P_t^{i,w}$, and pay nominal wage $W_{H,t}$ ($W_{L,t}$) for each unit of high-skilled (low-skilled) household labour. The low-skilled share in produc-

tion is α^i . Consistent with the evidence in Section 2.2, we assume that there are relatively more low-skilled workers in the production of non-essentials than in the essentials production:

$$\alpha^E < \alpha^N$$

As shown in Section 6, this heterogeneity is a main source of amplification in our estimates.

4.4 Rest of the model

The model is closed by two goods market clearing conditions (for essentials and non-essentials), two labour market clearing conditions (for high and low skilled labour), and a bond market clearing condition by which bonds are in zero net supply. Equations are detailed in Appendix D.4. The central bank sets interest rates according to a Taylor rule:

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R}\right)^{\rho_R} \left((\mathbb{E}_t(\pi_{t+1}))^{\phi_\pi} \left(\frac{Y_t}{Y}\right)^{\phi_Y} \right)^{1-\rho_R} \exp(\sigma^{mp} \varepsilon_t^{mp})$$

Fiscal policy ensures that Calvo retailers' profits are zero in steady state. Off-steady state, we specify an allocation rule that assigns profits to Ricardian households, as in Bilbiie (2008) and Debortoli and Galí (2017). In Section 6, we explore alternative profit allocation mechanisms. In Online Appendix I, we present the equilibrium definition, steady state computation, and the log-linearisation of the model around a zero inflation steady state (i.e. $\pi^E = \pi^N = 1$).

5 Structural Estimation

In the previous section, we have developed a novel model of the business-cycle featuring non-homothetic preferences on two consumption goods, heterogeneity in productivity across workers, and uneven skill-composition of the labour force across sectors. In the next section, we will use this structural model to perform counterfactual analyses that elicit the contribution of both non-essential spending and non-essential labour market dynamics to business-cycle fluctuations, and to the transmission of monetary policy in particular. For this exercise to provide a realistic quantification of our amplification, we require our structural model to replicate, as close as possible, the results from the IRFs analysis of Section 3 based on local projections.

In this section, we evaluate the ability of the structural model to produce evidence consistent with the findings of Section 3, by estimating its key parameters via impulse response matching. As customary in the literature, we split the parameters space into two groups.

The first set consists of standard coefficients that we calibrate following earlier studies. The second group refers to key parameters that are specific to our novel mechanism, which therefore we estimate. These are coefficients that govern the heterogeneity across essential and non-essential sectors, and that will be switched off in Section 6 to identify their relative contribution to the effects of monetary policy on aggregate consumption.

Calibration. In Panel A of Table 1, we collect the parameters that are calibrated. Low-skilled households are hand-to-mouth and their share, (μ^L) , is set to 1/3, consistent with the average MPC reported in micro empirical studies such as Johnson et al. (2006). The inverse of the macro Frisch elasticity is 0.1, consistent with the evidence reported by Christiano et al. (2010). The interest rate rule parameters are borrowed from Taylor (1993). We calibrate the standard deviation of the monetary policy shock so as to have an effect of 1% on annualised interest rates, on impact.

We calibrate the steady state consumption shares of essential goods for high-skilled households, (\bar{C}_H^E) , and low-skilled families, (\bar{C}_L^E) , to 0.44 and 0.60 respectively, using the expenditure share data in Online Appendix Table F.1. We match these moments by varying the scaling parameter for the relative utility of non-essential goods (φ) and the relative steady-state productivity between essential good production and non-essential good production ($a^E = A^E/A^N$). With non-homothetic preferences, consumption shares depend on the wages of workers in the two households. These can be affected heterogeneously by varying the relative productivity in the two sectors, given the uneven skill composition of the labour force across industries. We detail the moment matching algorithm and the steady state computation in Online Appendix I.2. The resulting values are: $\varphi = 1.3331$ and $a^E = 1.3372$. Finally, we allocate firms' profits to Ricardian agents.²²

Estimation procedure and prior distributions. We estimate the model parameters by minimizing the distance of the theoretical IRFs from the end-of-quarter impulse responses estimated either with SLPs (for essential and non-essential consumption, prices and earnings), or with the proxy-SVAR (for the 1y yields) in Section 3. We use a maximum a-posteriori approach, with a diagonal weighting matrix that exploits the standard errors of each estimated IRF to construct the likelihood function, following Guerron-Quintana et al. (2017).²³ Moments of the prior distributions are summarized in Table 1 Panel B. Prior means for essentials and non-essentials IESs are chosen so as to imply an aggregate IES of 0.86, which sits in the middle point between the estimate of Smets and Wouters (2007) and the log-utility case.

²²More specifically, we assume that $\phi_H^{\Pi,E} = \phi_H^{\Pi,N} = 1$ and that $\phi_L^{\Pi,E} = \phi_L^{\Pi,N} = 0$.

²³We detail the estimation procedure in Online Appendix J.1.

Table 1: Model Parameters

PANEL A - CALIBRATED PARAMETERS

Description	Parameter	Value
Time preference	β	0.99
Inverse of the macro Frisch elasticity	η	0.1
Dis-utility of working scaling parameter	ξ	1
Interest rate rule coefficient on inflation	ϕ_π	1.5
Interest rate rule coefficient on output gap	ϕ_Y	0.125
Standard deviation of the monetary policy process	σ^{mp}	0.255
Fraction of hand-to-mouth/low-skilled households	μ^L	1/3
Steady state share of essential good consumption by high skilled households	\bar{C}_H^E	0.44
Steady state share of essential good consumption by low skilled households	\bar{C}_L^E	0.60

PANEL B - ESTIMATED PARAMETERS

Description	Parameter	Distribution	Prior		Posterior	
			Mean	SD	Mean	SE
IES for essentials	γ^E	Normal	0.250	0.1	0.216	0.171
IES difference for non-essentials	$\gamma^N - \gamma^E$	Normal	1.000	0.1	0.770	0.299
Low skilled share in essentials	α^E	Beta	0.11	0.004	0.028	0.140
Low skilled share difference in non-essentials	$\alpha^N - \alpha^E$	Beta	0.12	0.004	0.322	0.130
Attentive share of households	λ	Beta	0.9	0.03	0.014	0.056
Interest rate smoothing	ρ_R	Beta	0.900	0.040	0.947	0.011
Price stickiness	θ	Beta	0.900	1.000	0.960	0.0083

Notes: Panel A shows the calibrated parameters and steady-state values. The scaling parameter for the relative utility of non-essential goods (φ) and the relative productivity between essential good production and non-essential good production ($a^E = A^E/A^N$) are computed to with the aid of other parameters to match the steady state share of essential good consumption by high skilled households (\bar{C}_H^E) and low skilled households (\bar{C}_L^E). Panel B shows the estimated parameters. The first column describes the parameter or convolution of parameters being estimated. The second column shows the corresponding symbol. The third column shows the distribution over which we draw the priors, whose mean and Standard Deviation (SD) are reported in columns 4 and 5. The sixth and seventh columns show the posterior mean and posterior standard error.

The priors on the shares of low-skilled workers are centered around 11% for essential industries and 23% for non-essentials, consistent with the evidence from the PSID.²⁴ The prior on the inattentiveness parameter is in line with the point estimate in [Beraja and Wolf \(2021\)](#). Interest rate smoothing and price stickiness coefficients display prior distributions consistent with the available evidence (e.g. [Smets and Wouters, 2007](#), [Justiniano and Primiceri, 2008](#)).

Estimation results. In Panel B of [Table 1](#), we report the posterior distributions of the parameters of the structural model. The IES for Essential goods is $\gamma^E = 0.22$ and the IES for Non-essential goods is $\gamma^N = 0.99$: the difference between the two IESs is economically and statistically very significant, implying an economy-wide IES of 0.64. These estimates

²⁴Unfortunately, PSID waves are only available bi-annually and thus the estimates of the structural model will be based on the quarterly earnings time series for essentials and non-essentials that we have constructed from the CPS. However, the CPS data do not allow us to measure the share of HtM workers in essential and non-essential sectors. Accordingly, we estimate these parameters using CPS quarterly earning sectoral time series and center their prior means around the values obtained in [Section 2](#) on the basis of the PSID.

correspond to an average income elasticity around 0.35 for essentials and about 1.53 for non-essentials. The average IEDs fall well within the range of income elasticities estimated in Appendix Table A.1 on the basis of CEX spending categories data, though those moments have not been targeted in the model estimation.

Our estimates of significant heterogeneity in the intertemporal elasticity of substitution across spending categories are based on macro data and on a classification strategy that covers most of household expenditure and employment in the economy. Still, our evidence is well aligned, and indeed complement, the estimates in [Attanasio et al. \(2002\)](#) and [Calvet et al. \(2021\)](#) which, based on fewer categories, suggest that low-income families have a lower intertemporal elasticity of substitution relative to high-income households. Their finding is consistent with both non-homothetic preferences and IES heterogeneity across essentials and non-essentials.

As for the shares of low-skilled workers in each sector, our posterior distributions move towards an even more unequal labour force skill composition than the priors, with the majority of low-skilled/hand-to-mouth workers employed in non-essential industries. Finally, the posteriors on the coefficients that govern inattentiveness, interest rate smoothing and price stickiness imply a larger inertia than the priors. While this may partially reflect the absence of an internal propagation mechanism in our parsimonious model, we note that the estimates of ρ_R , θ and λ are consistent with the evidence in earlier contributions (e.g. [Smets and Wouters, 2007](#), [Justiniano and Primiceri, 2008](#), [Beraja and Wolf, 2021](#)).

Model impulse response functions. The red crosses in Figure 6 reveal that the IRFs of consumption and earning implied by the estimated structural model track well the corresponding IRFs estimated with SLP-IVs, both for essentials and non-essentials. In Online Appendix Figure J.1, we report the full set of results, including the effects on essential and non-essential prices as well as on the interest rate. Spending and earnings on non-essentials decline by more than for essentials.

The estimates of our structural model are also able to reproduce the dampening of consumer’s spending relative to earnings discussed in Section 3.3: the gap between the decline in essential and non-essential earnings is larger than the gap for consumption. More generally, the estimated model appears able to match not only the qualitative patters of the empirical IRFs for spending and earnings but also the magnitude and the timing of their responses, with all peaks of the model IRFs within the 68% confidence intervals of the IRFs estimated with local projections.

Finally, high price stickiness results in small changes in prices: non-essential prices fall

more than for essentials, but both series are associated with only a small decline, which remains within the 90% empirical confidence bands of the SLPs. Unlike in the local projections IRFs, the estimated structural model suggests that essential prices decline (rather than rise), though neither IRFs appear of any statistical significance.

6 Inspecting the mechanism

The estimated structural model highlights three ingredients that can potentially alter the propagation of business-cycle shocks relative to a representative agent/representative good benchmark. First, non-homothetic preferences imply that non-necessities are easier to shift intertemporally, and thus their demand responds relatively more to income changes. The second ingredient is labour market heterogeneity: low-income workers are more likely employed in non-essential sectors and so they face a labour demand that is relatively more sensitive to the business-cycle. Finally, a large share of low-income workers in non-essential sectors are HtM and, therefore, the relatively stronger decline in their labour earnings during recessions feed back into lower aggregate demand, setting in motion a second round of spending and earnings effects that spillover to both sectors and along the income distribution, exacerbating the initial contraction.

In this section, we isolate the contribution of these channels to the transmission of monetary policy. In the first exercise, we take our estimated full model as benchmark and compute the share of the cumulated consumption response that one can explain using restricted versions that progressively strip down one or more of these dimensions. In the second exercise, we use the representative agent/representative good model as benchmark and quantify how much amplification one can obtain by adding each channel in isolation, and then jointly. The main take away is that the triple interaction between unequal MPC distribution (hand-to-mouth vs. savers), unequal spending composition (essentials vs non-essentials), and unequal sectoral labour composition (low-skilled vs. high-skilled workers) greatly amplify business-cycle fluctuations. In contrast, the contribution of each channel in isolation is much smaller; in fact, we show analytically that non-homothetic preferences lead to no amplification at all in an otherwise standard representative agent model.

6.1 Accounting for the aggregate effects of monetary policy

In Table 2, we seek to decompose the cumulative effects of monetary policy on consumption estimated by our structural model into the contribution of three sources of heterogeneity in: (i) spending composition, (ii) hand-to-mouth agents, and (iii) labour sectoral composition.

The first row focuses on the case of two identical goods under homothetic preferences whereas the second row represents the non-homothetic case in which the two goods exhibit different income elasticities of demand. The first column is for the model with a representative agent, the second column reports the results of versions that feature also hand-to-mouth consumers, while the third column further adds an uneven share of low-skilled workers across sectors.

At the two extremes of the models spectrum, there are the case with a representative agent and a representative good in the top-left corner of Table 2 and the estimated full structural model with hand-to-mouth consumers, unequal spending composition (i.e. non-homothetic preferences) and unequal labour composition (i.e. a higher share of low-skill workers in the non-essential sector) in the bottom-right corner. In all intermediate cases featured in the table, we either consider only one dimension of heterogeneity (i.e. only non-homotheticity in the bottom-left corner or only hand-to-mouth consumers in the top-middle entry) or at most two channels (i.e. non-homothetic preferences and hand-to-mouth households in the bottom-middle entry or hand-to-mouth consumers and unequal labour sectoral composition in the top-right corner).²⁵

For sake of exposition, we normalize all results in Table 2 by the cumulative response of consumption to a monetary policy shock in the estimated structural model (at the bottom right of the table), so that all other entries can be interpreted as the percentage contribution of each channel, either in isolation or in conjunction with another source of heterogeneity, in explaining the estimated overall effects on consumption. For instance, moving from the top-middle entry to the top-right corner (bottom-middle entry), we learn about the marginal contribution of unequal labour (spending) composition. On the other hand, by going diagonally from the top-middle cell to the bottom-right corner, we can evaluate the contribution of the interaction between unequal spending and unequal labour composition to explain the consumption response.

A few findings emerge from this exercise. First, representative agent models account for only 22% of the cumulative effects of monetary policy on consumption estimated using the full structural model, both with and without non-homothetic preferences.²⁶ Second, adding HtM

²⁵To implement the homothetic preference cases in the first row, we set the IES equal to the average IES, γ , implied by the full model and its estimated parameters in Table 1 (i.e. $\gamma = \gamma_E = \gamma_N$). In the representative agent models of the first column, we set the share of constrained agents μ_L to zero and fix the share of low-skilled workers in production to zero: $\alpha_E = \alpha_N = 0$. In the heterogeneous agent models, we set $\mu_L > 0$ and distinguish between two cases: (i) in the second column, we study a model where the share of low-skilled workers is the same in the two sectors (i.e. $\alpha_E = \alpha_N > 0$), with the common α chosen so as to match the relative steady state labour earnings across workers; (ii) in the third column, we use instead the values of α_E and α_N estimated in Table 1 using the full structural model.

²⁶Our representative agent case in the top left of the table could also be seen as an approximation of the direct effects of interest rates on consumption in our full model. Non-homothetic preferences in isolation do not result in additional direct effects (see Section 6.2 for an analytical proof of this result). The amplification

consumers in the second column brings the shares of the explained consumption response to 36% and 38%, respectively with and without equal spending composition.²⁷ Interestingly, the increase recorded when moving from the first to the second column is consistent with the estimates in [Patterson \(2023\)](#) who, as in [Bilbiie \(2020\)](#), emphasizes the ‘unequal incidence’ of recessions on the earnings of high-MPC and low-MPC workers.²⁸ In the top-right corner, we set the labour share of low-skilled workers in the non-essential sector to the value estimated in [Table 2](#), while counterfactually imposing equal spending composition. This raises the share of the explained consumption response to 49%.²⁹

Table 2: Counterfactual exercise: amplification

	Representative Agent	Heterogeneous Agents	
		Equal Labour Composition	Unequal Labour Composition
Equal spending composition	0.22	0.36	0.49
Unequal spending composition	0.22	0.38	1.00

Notes: Each cell display the ratio of the cumulative IRF of the counterfactual model in that cell over the cumulative IRF of estimated model in the bottom-right corner for aggregate consumption. In the homothetic case, which we refer to as ‘equal spending composition’ (i.e. $\gamma^N = \gamma^E$), we set the IES equal to the estimated average IES in the economy. In the representative agent column, we set $\mu^L = 0$ and $\alpha^E = \alpha^N = 0$. Under ‘equal labour composition’, we fix $\alpha^E = \alpha^N > 0$ so as to match the relative steady state labour earnings across the two agents.

The main take away from [Table 2](#), however, is that the interaction between unequal spending composition (i.e. $\gamma^N \neq \gamma^E$) and unequal labour composition (i.e. $\alpha^N \neq \alpha^E$) accounts for the bulk of the estimated consumption response in the full structural model. We conclude this by noticing that the jumps from, respectively, 0.38 and 0.49 to 1 (when the two channels are *jointly* considered) are much larger than the increases from 0.36 to, respectively, 0.38 and 0.49 (when each channel is assessed *individually*). In other words, the interaction of these two sources of heterogeneity provides a far more powerful amplification than each of them in isolation, suggesting a quantitatively important complementarity between unequal spending

in the rest of the table results from indirect, general equilibrium effects. [Kaplan et al. \(2018\)](#) show that direct effects account for almost the entirety of the transmission of monetary policy in representative agent models. In contrast, they find that direct effects in HANK contribute 19% of the total consumption response, which is not dissimilar in magnitude from the 22% explained by our representative agent case.

²⁷The intuition for why unequal spending composition generates further amplification relative to TANK (i.e. $\gamma^N \neq \gamma^E$) is that, under non-homothetic preferences, the unconstrained agents spend more on non-essentials and therefore their IES is higher than the average IES.

²⁸We have verified that even in the special case of equal spending (i.e. $\gamma^N = \gamma^E$) and equal labour composition (i.e. $\alpha^N = \alpha^E$) in the top-middle entry of [Table 2](#), the estimates of [Table 1](#) imply that the income elasticity of HtM agents with respect to aggregate income, $\tilde{\chi}$, is larger than one, which is a necessary and sufficient condition for amplification in this class of models ([Bilbiie, 2020](#)).

²⁹In TANK, constrained workers face more cyclical earnings than unconstrained workers (who are cushioned by the counter-cyclicality of profits) and therefore are less likely to reduce their labour supply in response to negative shocks. In a model that also features unequal labour composition, this produces a relatively larger drop in earnings for the sector that employs a larger share of constrained workers, thereby generating further amplification relative to TANK (i.e. $\alpha^N \neq \alpha^E$).

composition across goods and unequal labour composition across sectors in accounting for business-cycle fluctuations.

A key difference between the representative agent cases and the heterogeneous agents models is the presence of hand-to-mouth consumers only in the latter. This feature interacts with the composition of both cyclical product demand and cyclical labour demand to further amplify the effects of monetary policy. To illustrate this triple interaction, in Appendix Figure J.2, we report the aggregate consumption response in the four heterogeneous agents cases of Table 2 as we vary the share of hand-to-mouth households, μ^L , from 0 to 0.33, a value consistent with the empirical literature on estimating MPCs (e.g. Johnson et al., 2006). Appendix Figure J.2 shows that a higher share of hand-to-mouth consumers leads to a monotonic increase in the extent of amplification across all models. The case with both equal labour composition and equal spending composition, often referred to as Two-Agents New-Keynesian (TANK) model, exhibits a degree of amplification relative to the representative agent/representative that is between 15% and 50%, over the empirically plausible range of [0.15, 0.33] for the average MPC, consistent with the evidence on U.S. data by Patterson (2023) and Bilbiie et al. (2023). The addition of either non-homothetic preferences or unequal labour sectoral composition —on their own— amplifies only marginally more than this baseline TANK model. However, the extent of amplification in the full model is consistently larger the sum of these individual additions. This reveals that the triple interaction between cyclical product demand composition, cyclical labour demand composition and hand-to-mouth households generates a strong complementarity that greatly amplifies business-cycle fluctuations relative not only to the representative agent/representative good case but also to heterogeneous agents models that only feature the double interaction between constrained agents and heterogeneity in either consumers’ spending or workers’ sectoral composition.

6.2 Non-homotheticity alone does not lead to amplification

In the previous section, we have shown that adding non-homothetic preferences to the representative agent version of our estimated model generates no amplification: moving through the rows of Table 2 first column does not change the share of the explained consumption response to monetary policy. In this section, we generalize that result by showing that the distinction between essentials and non-essentials has no impact on aggregate fluctuations in representative agent settings.

In Online Appendix K Proposition 2, we prove analytically this result in a streamlined version of our model in which we remove inattentiveness and employ a simplified Taylor rule, $R_t = \phi_\pi \mathbb{E}_t(\pi_{t+1}) + \varepsilon_t^{mp}$. We begin by demonstrating that in representative agent models where

there are no interactions between sectoral heterogeneity and non-homothetic preferences, the average IES of the economy is a sufficient statistics to measure aggregate fluctuations and, therefore, it fully characterises the impact of monetary policy on total consumption. Next, we consider a homothetic version of the representative agent model where we set the IES to the value of the average IES in the non-homothetic preferences economy. In Corollary 1, we show what the responses of consumption (and inflation) to a monetary policy shock is identical to the non-homothetic version. More specifically, we prove that:

$$\begin{aligned} \text{Non-Homothetic} &\rightarrow \frac{\partial \hat{C}_t}{\partial \varepsilon_t^{mp}} = - \underbrace{(\gamma^E(1 - \bar{C}^N) + \gamma^N \bar{C}^N)}_{\text{Average IES}} \\ \text{Homothetic} &\rightarrow \frac{\partial \hat{C}_t}{\partial \varepsilon_t^{mp}} = -IES \end{aligned}$$

It is worth noting that while the simplifications above allowed us to provide closed-form analytical expressions, this equivalence result is more general than shown in this section. As we have emphasised in Section 6.1, the cumulated response of total consumption in the representative agent models of Table 2 are also identical. This, however, should not be interpreted to mean that the demand for essentials and non-essentials is identical. In fact, in Online Appendix Table J.1, we report that essential consumption still falls less than non-essential spending in the representative agent model. In other words, with non-homothetic preferences alone, the heterogeneity of the effects of monetary policy on essentials and non-essentials is a zero sum-game: the larger fall in non-necessities is perfectly offset by the smaller fall in necessities, so as to generate a decline in aggregate spending that is exactly equal to the aggregate response in the homothetic case that is parameterized using the same average IES. On the other hand, the results in the previous section suggest that non-homotheticity can still be an important source of amplification whenever coupled with other sources of heterogeneity. Our analysis uncovers an important role for an empirically relevant instance of this type of interactions: with cyclical labour demand composition. Other examples for future research may include heterogeneity in price or wage stickiness across sectors.

7 Unconventional Fiscal Policy

Monetary policy aims at stabilizing the business-cycle but is a blunt tool for targeting specific households or sectors. In contrast, fiscal policy can stimulate spending for particular groups or industries. In this section, we illustrate the key mechanism of our model through an unconventional fiscal policy, which we tailor, in turn, to the essential and to the non-essential

sector. We consider three types of temporary VAT shocks, which are levied: (i) uniformly on all goods and services in the economy, (ii) only on non-essentials, (iii) only on essentials.³⁰

A main reason for this choice is that, in our economy, a uniform increase in the VAT rate accompanied by a suitable decline in payroll taxes can perfectly replicate the effects of monetary policy implied by the estimates of the structural model of Section 5. [Correia et al. \(2013\)](#) demonstrate the mapping between monetary policy and this type of unconventional fiscal policy in a one-good, one-agent economy. We generalize their result to a two-goods, two-agents setting.³¹ This allows us to make contacts with the evidence on monetary policy in the rest of the paper.

We extend our model to include payroll taxes, lump-sum transfers and VAT shocks. The latter can be either uniform across consumption categories or specific to essentials and non-essentials, respectively. Our analysis proceeds in two steps. First, in [Appendix Proposition 1](#), we show formally that, in our two-goods-two-agents model, an increase in interest rates is observationally equivalent to a uniform hike in the VAT rate compensated by a decline in payroll taxes. The intuition is that changes in interest rates and changes in the growth rate of a uniform VAT rate enter the Euler equation in exactly the same way. Payroll taxes nullify the effect of the uniform VAT rate both on the labour supply and on the household budget constraint.³²

The second step is the implementation through a tax rule. We specify an AR(1) for the uniform VAT tax rate and we match the response of overall consumption to a monetary policy shock in the estimated structural model. With respect to the goods-specific VAT rates, we specify the same AR(1) process but scale up the size of the tax rate change in a way that is proportional to the economy wide consumption share of the spending category on which the goods-specific VAT is levied. This choice fulfills our desire of keeping the change in gross tax revenues constant across the three cases. We wish to evaluate the following question: suppose the government has an extra dollar that can be allocated to either reduce uniformly the VAT rate across all goods and services or cut the VAT rate only for one consumption category of essentials or non-essentials, which one of these three policies would stimulate

³⁰An example of this type of sectoral policies was the ‘Eat Out to Help Out’ implemented by the British government during the Covid pandemic in 2020. The scheme consisted of a temporary 50% discount, up to GBP 10 per person, on an order for a restaurant meal. The stated objective was to boost spending and help out workers in the hospitality industry. [Gonzalez-Pampillon et al. \(2024\)](#) find that ‘Eat Out to Help Out’ provided a significant boost to sales in the hospitality sector, whereas [Fetzer \(2022\)](#) highlights the unintended consequences of this policy in terms of a significant increase in Covid cases.


³¹[Wolf \(2024\)](#) shows an equivalence between monetary policy and time varying lump-sum transfers. In contrast, we focus on different consumer spending categories and use time-varying distortionary VAT rates, as these can be tailored to specific goods and services.

³²In the extended model, we also allow for changes in lump-sum transfers, though these are not needed to proof the equivalence between monetary policy and unconventional fiscal policy.

aggregate consumption the most? In Appendix Section E.1, we layout the full description of the model whereas, in Appendix Section E.2, we present the equivalence proof and the implementation details.

The unconventional fiscal policy horse race is reported in Figure 8. The left panel refers to the policy rate changes, the middle column is aggregate consumption while the right chart reports aggregate earnings. The blue dashed line summarises the effects of interest rate changes, the purple solid line stands for the uniform VAT shock, the broken orange line exemplifies the impact of a VAT increase levied only to non-essentials whereas the dotted green line refers to a VAT hike applied only to essentials. In the background of Appendix Figure E.1, payroll taxes move in the opposite direction of the VAT tax rates, by design.³³ Consistent with Proposition 1, a uniform VAT tax rate shock can exactly replicate the effects of monetary policy on consumption, earnings, and all other private sector variables. We find that a generalized increase of 5% in the VAT rate is associated with a peak reduction of 0.8% in aggregate consumption, consistent with the evidence in Crossley et al. (2009, 2014) and D’Acunto et al. (2022).

Figure 8: Unconventional fiscal policies and monetary policy equivalence



Notes: Responses of policy rates, aggregate consumption and aggregate earnings to alternative fiscal policies and to monetary policy. The uniform VAT shock is applied evenly to all goods and services in the economy, accompanied by an offsetting income tax (not shown) such as to replicate exactly the response of aggregate consumption and earnings to a monetary policy shock of 100 basis points. The essential and non-essential VAT shocks are sector-specific VAT policies designed to raise the same tax revenues as the uniform VAT change. The size of these policy shocks is shown in the chart on the first column which reports interest rate basis points for monetary policy and percentage change of tax rates for VAT; the resulting aggregate consumption and earnings responses are shown in the middle and right charts respectively.

A main result emerges from Figure 8: a VAT rate hike applied only to non-essentials produces much larger effects than either a uniform VAT changes across all goods and services or a VAT rate increase levied only on essentials. The reason for this result is twofold. First,

³³Appendix Figure E.1 also shows additional variables, including the responses within the essential and non-essential sectors and price responses. Appendix Section E.3 presents further results.

as we have shown in Section 2.2 and implemented in the structural model, many more hand-to-mouth workers are employed in the non-essential sector. This implies that whenever non-essential VAT rates are higher, the lower demand for goods and services produced in this sector triggers a larger decline in labour demand (and therefore income) for HtM workers, which in turn spills over to essential spending and exacerbates the initial contraction. Second, the goods-specific VAT rate change also introduces a relative price effect: immediately after the shock, the post-tax price of non-essentials rises relative to both the contemporaneous price of essentials and the future price of non-essentials. This leads households (both Ricardian and HtM) to substitute toward essentials or postpone consumption (Ricardian households only), resulting in a contraction of non-essential demand that is larger than in the case of either a contractionary monetary policy shock or a VAT increase applied only to essentials. A VAT shock applied only to non-essentials further reduce aggregate consumption by 0.3 percentage points relative to a VAT shock of the same size that is levied uniformly across both sectors.³⁴

In summary, this section exemplifies the inner working of our mechanism. A change in the product demand for non-essentials triggers a relative fall in labour demand, whose burden disproportionately falls on HtM workers. In general equilibrium, this triggers second-round effects on consumption that amplify the impact of the initial shock. This mechanism is at play for the cases of a monetary policy shock and of a VAT rate change on essentials only. But it generates further amplification when the unconventional fiscal policy targets non-essentials, whose VAT rate hike hits the earnings of low-income workers more.

8 Conclusions

This paper proposes a framework to evaluate changes in consumer spending based on the distinction between essentials and non-essentials (luxuries versus necessities). We show that this set up offers novel insights for understanding business-cycle fluctuations and the dynamic effects of stabilization policies. A main finding is that households and workers differ greatly in their exposure to the business-cycle along the income distribution, and that the compositions of product and labour demand across essential and non-essential sectors is crucial to identify and quantify their cyclical exposure to shocks. A key intuition is that the discretionary spending of the rich drives the earnings of the poor. In the face of economic adversities, non-essential purchases are easier to postpone and their fall is driven by affluent households. The latter has a particularly large effect on low-income/HtM workers, who are more likely

³⁴The larger effect is driven by the stronger decline in non-essential spending (see Appendix Figure E.1).

employed in non-essential industries and therefore see their labour demand contracting more: the higher cyclical of non-essential spending leads to a higher cyclical of non-essential earnings. Taken together, the households with less resources in society lose twice as a result of the behaviour of high-income consumers, via: (i) a price effect that makes their necessity-dominated consumption bundle relatively more expensive, (ii) an income effect that lowers their labour earnings and thus the resources that low-income families have available for both types of spending.

Using newly constructed, nationally representative time series from multiple granular micro data, we show that: (i) high-income consumers spend more on non-essentials; (ii) low-earning/HtM workers are more likely employed in non-essential industries; (iii) during recessions, non-essential spending and non-essential earnings fall far more than their essential counterparts. Furthermore, we find that the effects of demand composition on income largely dominates the effects on relative prices across goods. We develop and estimate a structural model with nominal rigidities, non-homothetic preferences, and heterogeneity in both productivity across workers and labour force composition across sectors. We show that the estimated model replicates well the stronger cyclical of both non-essential spending and non-essential earnings after an interest rate change. We use the estimated model to decompose the aggregate effects of monetary policy and find that the triple interaction between cyclical product demand composition (i.e. non-essential spending contracts by more), cyclical labour demand composition (i.e. non-essential earnings fall by more), and heterogeneous sectoral labour composition (i.e. low-income workers/HtM are more likely employed in non-essential industries) accounts for a significant portion of the effects of monetary policy on consumption.

Our findings have potentially interesting fiscal ramifications: government policies aimed at non-essential sectors are likely to work through the mechanism documented in this paper. For instance, a combination of consumption and labour taxes can perfectly replicate the effects of monetary policy. However, a temporary VAT cut applied only to non-essentials stimulates aggregate consumption *and* the earnings of low-income households far more than both monetary policy and a consumption tax cut of the same size levied either only on essentials or uniformly on all sectors. We leave the design of optimal policies for future research. But, we hope that our analysis might stimulate and assist statistical offices and central banks in leveraging the available granular data on households, workers and sectors to construct novel time series for prices, quantities and earnings of essentials and non-essentials.

References

- ACEMOGLU, D., V. M. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2012): “The network origins of aggregate fluctuations,” *Econometrica*, 80, 1977–2016.
- AGUIAR, M. AND M. BILS (2015): “Has Consumption Inequality Mirrored Income Inequality?” *American Economic Review*, 105, 2725–56.
- AHN, H. J., B. HOBIJN, AND A. ŞAHIN (2024): “The Dual U.S. Labor Market Uncovered,” *mimeoographed*, Federal Reserve Board and Princeton University.
- AIT-SAHALIA, Y., J. A. PARKER, AND M. YOGO (2004): “Luxury Goods and the Equity Premium,” *The Journal of Finance*, 59, 2959–3004.
- AMBERG, N., T. JANSSON, M. KLEIN, AND A. R. PICCO (2022): “Five Facts about the Distributional Income Effects of Monetary Policy Shocks,” *American Economic Review: Insights*, 4, 289–304.
- ANDERSEN, A. L., N. JOHANNESSEN, M. JØRGESEN, AND J.-L. PEYDRO (2023): “Monetary Policy and Inequality,” *The Journal of Finance*, 78, 2945–2989.
- ANDREOLLI, M. AND P. SURICO (2021): “LESS Is MORE: Consumer Spending and the Size of Economic Stimulus Payments,” CEPR Discussion Paper.
- ANGELETOS, G.-M., F. COLLARD, AND H. DELLAS (2020): “Business-cycle anatomy,” *American Economic Review*, 110, 3030–3070.
- ARUOBA, S. B. AND T. DRECHSEL (2024): “Identifying monetary policy shocks: A natural language approach,” Tech. rep., National Bureau of Economic Research.
- ATTANASIO, O. P., J. BANKS, AND S. TANNER (2002): “Asset holding and consumption volatility,” *Journal of political economy*, 110, 771–792.
- AUCLERT, A. (2019): “Monetary Policy and the Redistribution Channel,” *American Economic Review*, 109, 2333–2367.
- AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2020): “Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model,” .
- BAQAEE, D. R. AND E. FARHI (2018): “Macroeconomics with Heterogeneous Agents and Input-Output Networks,” Working Paper 24684, National Bureau of Economic Research.
- BARNICHON, R. AND C. BROWNLEES (2019): “Impulse Response Estimation by Smooth Local Projections,” *The Review of Economics and Statistics*, 101, 522–530.
- BARSKY, R. B., C. L. HOUSE, AND M. S. KIMBALL (2007): “Sticky-Price Models and Durable Goods,” *The American Economic Review*, 97, 984–998.
- BAUER, M. D. AND E. T. SWANSON (2023): “A reassessment of monetary policy surprises and high-frequency identification,” *NBER Macroeconomics Annual*, 37, 87–155.
- BERAJA, M. AND C. K. WOLF (2021): “Demand Composition and the Strength of Recoveries,” Working Paper 29304, National Bureau of Economic Research.
- BERALDI, F. AND C. MALGIERI (2024): “Fiscal Multipliers and Phillips Curves with a Consumption Network,” Working paper.
- BERGMAN, N. K., N. JAIMOVICH, AND I. SAPORTA-EKSTEN (2025): “The Impact of Economic Shifts: From Micro-Elasticities to Distributional Welfare Effects ,” .
- BILBIIE, F. O. (2008): “Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic,” *Journal of economic theory*, 140, 162–196.
- (2020): “The New Keynesian Cross,” *Journal of Monetary Economics*, 114, 90–108.

- BILBIIE, F. O., D. R. KÄNZIG, AND P. SURICO (2022): “Capital and income inequality: An aggregate-demand complementarity,” *Journal of Monetary Economics*, 126, 154–169.
- BILBIIE, F. O., G. PRIMICERI, AND A. TAMBALOTTI (2023): “Inequality and Business Cycles,” NBER Working Papers 31729, National Bureau of Economic Research, Inc.
- BLUNDELL, R., M. C. DIAS, R. JOYCE, AND X. XU (2020): “COVID-19 and Inequalities*,” *Fiscal Studies*, 41, 291–319.
- BOPPART, T. (2014): “Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences,” *Econometrica*, 82, 2167–2196.
- BROWNING, M. AND T. F. CROSSLEY (2000): “Luxuries Are Easier to Postpone: A Proof,” *Journal of political Economy*, 108, 1022–1026.
- CALVET, L. E., J. Y. CAMPBELL, F. GOMES, AND P. SODINI (2021): “The cross-section of household preferences,” Tech. rep., National Bureau of Economic Research.
- CHAI, A. AND A. MONETA (2010): “Retrospectives: Engel Curves,” *Journal of Economic Perspectives*, 24, 225–40.
- CHRISTIANO, L. J., M. TRABANDT, AND K. WALENTIN (2010): “DSGE models for monetary policy analysis,” in *Handbook of monetary economics*, Elsevier, vol. 3, 285–367.
- CLAYTON, C., X. JARAVEL, AND A. SCHAAB (2018): “Heterogeneous Price Rigidities and Monetary Policy,” *SSRN Electronic Journal*.
- CLOYNE, J., C. FERREIRA, AND P. SURICO (2020): “Monetary Policy When Households Have Debt: New Evidence on the Transmission Mechanism,” *The Review of Economic Studies*, 87, 102–129.
- COIBION, O., Y. GORODNICHENKO, L. KUENG, AND J. SILVIA (2017): “Innocent Bystanders? Monetary Policy and Inequality,” *Journal of Monetary Economics*, 88, 70–89.
- COMIN, D., D. LASHKARI, AND M. MESTIERI (2021): “Structural Change with Long-Run Income and Price Effects,” *Econometrica*, 89, 311–374.
- COMIN, D. A., A. DANIELI, AND M. MESTIERI (2020): “Income-driven Labor Market Polarization,” Working Paper 27455, National Bureau of Economic Research.
- CORREIA, I., E. FARHI, J. P. NICOLINI, AND P. TELES (2013): “Unconventional Fiscal Policy at the Zero Bound,” *American Economic Review*, 103, 1172–1211.
- CROSSLEY, T., H. LOW, AND C. SLEEMAN (2014): “Using a temporary indirect tax cut as a fiscal stimulus: evidence from the UK,” IFS Working Papers W14/16, Institute for Fiscal Studies.
- CROSSLEY, T. F., H. LOW, AND M. WAKEFIELD (2009): “The Economics of a Temporary VAT Cut,” *Fiscal Studies*, 30, 3–16.
- CROSSLEY, T. F. AND H. W. LOW (2011): “Is the Elasticity of Intertemporal Substitution Constant?” *Journal of the European Economic Association*, 9, 87–105.
- DE NARDI, M. AND G. FELLA (2017): “Saving and Wealth Inequality,” *Review of Economic Dynamics*, 26, 280–300.
- DEBORTOLI, D. AND J. GALÍ (2017): “Monetary Policy with Heterogeneous Agents: Insights from TANK Models,” .
- D’ACUNTO, F., D. HOANG, AND M. WEBER (2022): “Managing Households’ Expectations with Unconventional Policies,” *The Review of Financial Studies*, 35, 1597–1642.
- ENGEL, E. (1857): “Die Productions- Und Consumtionsverhältnisse Der Königsreichs Sachsen,” *Zeitschrift des Statistischen Bureaus des Königlich-Sächsischen, Ministerium des Innern*, [1895 repr. in *Bulletin de l’institute international de statistique*], v. 9, p. 1.

- FETZER, T. (2022): “Subsidising the spread of COVID-19: Evidence from the UK’S Eat-Out-to-Help-Out Scheme,” *The Economic Journal*, 132, 1200–1217.
- FLOOD, S., M. KING, R. RODGERS, S. RUGGLES, AND J. R. WARREN (2020): “Integrated Public Use Microdata Series, Current Population Survey: Version 7.0,” .
- FLYNN, J. P., C. PATTERSON, AND J. STURM (2021): “Fiscal Policy in a Networked Economy,” Working Paper 29619, National Bureau of Economic Research.
- FOELLM, R. AND J. ZWEIMÜLLER (2008): “Structural Change, Engel’s Consumption Cycles and Kaldor’s Facts of Economic Growth,” *Journal of monetary Economics*, 55, 1317–1328.
- GERTLER, M. AND P. KARADI (2015): “Monetary Policy Surprises, Credit Costs, and Economic Activity,” *American Economic Journal: Macroeconomics*, 7, 44–76.
- GONZALEZ-PAMPILLON, N., G. NUNEZ-CHAIM, AND H. G. OVERMAN (2024): “The economic impacts of the UK’s eat out to help out scheme,” *Journal of Urban Economics*, 143.
- GUERON-QUINTANA, P., A. INOUE, AND L. KILIAN (2017): “Impulse response matching estimators for DSGE models,” *Journal of Econometrics*, 196, 144–155.
- GURKAYNAK, R. S., B. SACK, AND E. T. SWANSON (2005): “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements,” *International Journal of Central Banking*, 1, 39.
- HEATHCOTE, J., F. PERRI, AND G. L. VIOLANTE (2010): “Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006,” *Review of Economic Dynamics*, 13, 15–51.
- HOLM, M. B., P. PAUL, AND A. TISCHBIREK (2021): “The Transmission of Monetary Policy under the Microscope,” *Journal of Political Economy*, 129, 2861–2904.
- JAIMOVICH, N., S. REBELO, AND A. WONG (2019): “Trading down and the business cycle,” *Journal of Monetary Economics*, 102, 96–121.
- JAIMOVICH, N., S. REBELO, A. WONG, AND M. B. ZHANG (2020): “Trading Up and the Skill Premium,” *NBER Macroeconomics Annual*, 34, 285–316.
- JAROCIŃSKI, M. AND P. KARADI (2020): “Deconstructing Monetary Policy Surprises—The Role of Information Shocks,” *American Economic Journal: Macroeconomics*, 12, 1–43.
- JOHNSON, D. S., J. A. PARKER, AND N. S. SOULELES (2006): “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 96, 1589–1610.
- JOHNSON, N. N. (2017): “Tradable and nontradable inflation indexes: replicating New Zealand’s tradable indexes with BLS CPI data,” *Monthly Labor Review*, 1–25.
- JORDÀ, Ò. (2005): “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161–182.
- JUSTINIANO, A. AND G. E. PRIMICERI (2008): “The Time-Varying Volatility of Macroeconomic Fluctuations,” *American Economic Review*, 98, 604–641.
- KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): “Monetary Policy According to HANK,” *American Economic Review*, 108, 697–743.
- KAPLAN, G., G. L. VIOLANTE, AND J. WEIDNER (2014): “The wealthy hand-to-mouth,” Tech. rep., National Bureau of Economic Research.
- MANKIW, N. G. AND R. REIS (2007): “Sticky information in general equilibrium,” *Journal of the European Economic Association*, 5, 603–613.
- MCKAY, A. AND J. F. WIELAND (2019): “Lumpy Durable Consumption Demand and the Limited Ammunition of Monetary Policy,” Tech. rep., National Bureau of Economic Research.

- MERTENS, K. AND M. O. RAVN (2011): “Understanding the aggregate effects of anticipated and unanticipated tax policy shocks,” *Review of Economic dynamics*, 14, 27–54.
- (2013): “The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States,” *American Economic Review*, 103, 1212–1247.
- MIRANDA-AGRIPPINO, S. AND G. RICCO (2021): “The Transmission of Monetary Policy Shocks,” *American Economic Journal: Macroeconomics*, 13, 74–107.
- MONTIEL OLEA, J. L. AND C. PFLUEGER (2013): “A Robust Test for Weak Instruments,” *Journal of Business & Economic Statistics*, 31, 358–369.
- MONTIEL OLEA, J. L. AND M. PLAGBORG-MØLLER (2021): “Local Projection Inference Is Simpler and More Robust than You Think,” *Econometrica*, 89, 1789–1823.
- NEWKEY, W. K. AND K. D. WEST (1987): “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703–708.
- NIPA (2021): “NIPA Handbook: Concepts and Methods of the U.S. National Income and Product Accounts,” Tech. rep.
- OLIVI, A., V. STERK, AND D. KHANI (2023): “Optimal Monetary Policy during a Cost of Living Crisis,” Tech. rep., mimeo, UCL. 1.
- ORCHARD, J. (2022): “Cyclical Demand Shifts and Cost of Living Inequality,” *Available at SSRN 4033572*.
- PATTERSON, C. (2023): “The Matching Multiplier and the Amplification of Recessions,” *American Economic Review*, 113, 982–1012.
- RAVN, M. AND V. STERK (2017): “Job uncertainty and deep recessions,” *Journal of Monetary Economics*, 90, 125–141.
- RAVN, M. O. AND V. STERK (2021): “Macroeconomic Fluctuations with HANK & SAM: an Analytical Approach,” *Journal of the European Economic Association*, 19, 1162–1202.
- ROMER, C. D. AND D. H. ROMER (2004): “A new measure of monetary shocks: Derivation and implications,” *American economic review*, 94, 1055–1084.
- RUBBO, E. (2023): “Networks, Phillips Curves, and Monetary Policy,” *Econometrica*, 91, 1417–1455.
- SCHAAB, A. AND S. Y. TAN (2023): “Monetary and Fiscal Policy According to HANK-IO,” Tech. rep., Working paper.
- SMETS, F. AND R. WOUTERS (2007): “Shocks and frictions in US business cycles: A Bayesian DSGE approach,” *American economic review*, 97, 586–606.
- SONNERVIG, M. (2023): “Unequal Business Cycles,” Mimeo, NYU.
- STIGLITZ, J. E. (1969): “Behavior towards risk with many commodities,” *Econometrica: Journal of the Econometric Society*, 660–667.
- STOCK, J. H. AND M. W. WATSON (2018): “Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments,” *The Economic Journal*, 128, 917–948.
- (2020): “Slack and Cyclically Sensitive Inflation,” *Journal of Money, Credit and Banking*, 52, 393–428.
- TAYLOR, J. B. (1993): “Discretion versus policy rules in practice,” *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.
- VERONA, F. AND M. H. WOLTERS (2014): “Sticky information models in Dynare,” *Computational Economics*, 43, 357–370.

- WACHTER, J. A. AND M. YOGO (2010): “Why Do Household Portfolio Shares Rise in Wealth?” *The Review of Financial Studies*, 23, 3929–3965.
- WOLF, C. K. (2024): “Interest Rate Cuts vs. Stimulus Payments: An Equivalence Result,” *Journal of Political Economy*, forthcoming.
- ZORZI, N. (2020): “Investment Dynamics and Cyclical Redistribution,” Working paper.

Appendix

“Non-Essential Business Cycles”

by Michele Andreolli (Boston College), Natalie Rickard (LBS) and Paolo Surico (LBS)¹

¹This appendix contains key supplementary material referenced in the main text. In addition to this, an Online Appendix contains further material and evidence for interested readers, available at:

https://mandreolli.github.io/files/AndreolliRickardSurico_NEBC_OnlineAppendix.pdf.

A Defining non-essentials

We classify consumption categories into essentials and non-essentials by estimating Engel Curves closely following the approach used by [Aguiar and Bils \(2015\)](#). In Table A.1, we report the estimated expenditure elasticities and expenditure shares for the revised goods categories, which is a replication of Table II of [Aguiar and Bils \(2015\)](#). We make two minor alterations to [Aguiar and Bils \(2015\)](#)'s approach; altering slightly their product categories and using data from 1995-1997 rather than 1994-1996. These details and a description of the methodology are included in Online Appendix F.

Table A.1: Engel curves used for Essential/Non-essential classification

Good category	CE share 1995-1997	Elasticity	SE
Rent	5.5	-1.1	0.09
Used car purchases	5.53	0.23	0.16
Communication, telephone contracts	2.59	0.31	0.04
Food at home	11.63	0.4	0.02
Utilities	5.21	0.47	0.02
Children's clothing	0.96	0.65	0.07
Gas and vehicle maintenance	6.14	0.72	0.03
Health expenditures including insurance	4.9	0.81	0.05
Personal care	0.97	0.96	0.05
Shoes and other apparel	1.47	1.07	0.09
Other car spending (leasing, financing, insurance)	5.45	1.14	0.06
Entertainment equipment and subscription television	4.01	1.22	0.07
Alcoholic beverages	0.96	1.22	0.09
Men's and women's clothing	2.47	1.36	0.05
Food away from home	4.53	1.37	0.05
Household appliances	2.3	1.42	0.07
Owner occupied housing consumption	22.25	1.45	0.04
Furniture and fixtures	1.51	1.5	0.11
Education	1.31	1.58	0.18
Domestic services and childcare	1.48	1.61	0.14
New car purchases	3.91	1.74	0.2
Public transport	1.25	1.78	0.13
Entertainment fees, admissions, reading	2.17	1.78	0.07
Cash contributions	2.18	1.78	0.17

Notes: Replication of Table II of [Aguiar and Bils \(2015\)](#), for 1995-1997 and for revised categories. The elasticity is the estimated β_j from Online Appendix (4). (Non-)Essential goods are those with an elasticity less than (greater than) one, above (below) the dashed line. The CE share is the share of expenditure of each category reported in the Consumer Expenditure Survey over the sample period.

B Durables versus non-durables and other classifications

One major alternative characteristic that has been extensively discussed in the literature is the durability of goods ([Barsky et al. \(2007\)](#), [McKay and Wieland \(2019\)](#)). For instance, in our classification furniture and new cars are non-essentials and durable. To explore this, we further break down our essential and non-essential series into durables and non-durables.

Classification. Following the same approach and data sources in the main text, we also construct series from disaggregated data for durables and non-durables. For PCE, we classify goods as durable following the categorisation in the PCE. We only include consumption categories that we also have an essential or non-essential classification. This latter cover the majority (approximately 80%) of overall expenditure, but omits some durables/ nondurables which are categorised in the PCE data. We categorise according the nature of the final good, not the intermediate goods. If a good is not a final good, it is not classified.

For intermediate industries, we classify industry production according to the nature of the final downstream goods it supplies rather than the intermediate industry, following the same approach as for the IO approach to classifying intermediate industries. In addition, we classify construction industries as durable, but in the consumption data this is non-durable because we use the consumption flow of housing, which is a service in the national accounts/ PCE data.

Descriptive statistics. In Table B.1, we show the shares of consumption accounted for by durables and non-durables vs essentials and non-essentials. There is a correlation between the two characteristics; within essentials, there are almost no durables, whereas durables make up a substantial minority of non-essential expenditure and almost half of non-essentials employment. However, we view durables as a separate, significant driver of consumption volatility, rather than an alternative interpretation of our non-essentials vs essentials split.

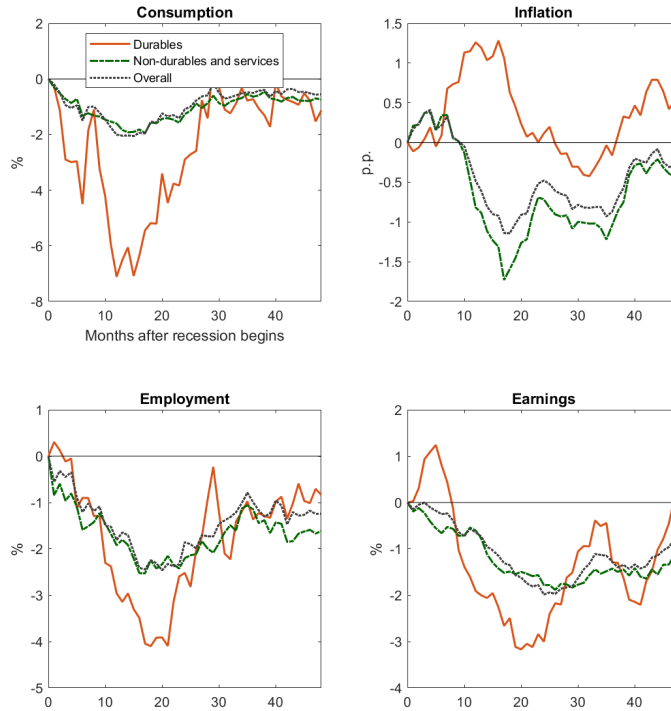
Table B.1: Durables vs Non-essentials

	All	Durable goods	Non-durable goods and services
Essential	45.00	2.57	42.43
Non-essential	55.00	10.86	44.14
Both	100.00	13.43	86.57

Percentages of total consumption. Calculated using PCE expenditure data, consumption shares based on chained 2000 dollar consumption series; only consumption categories in the essentials/ non-essentials classification are included. Averaged over 1973-2020.

To demonstrate the difference between the two categorisations, we first present the equivalent of our main descriptive evidence, split into durables and non-durables. Figure B.1 is the equivalent of Figure 2 for durables and non-durables. It shows that durables are more cyclical than non-durables, as expected, but the fact that durables account for a much smaller proportion of overall consumption and employment means that the response of the aggregate series is extremely close to the non-durable and services series; in contrast, non-essential cyclicity appears to be contributing relatively more to overall cyclicity. Figure B.2 shows the earnings distributions for a four-way split between essentials/non-essentials and durables/non-durables. Non-essential non-durables (rather than the more cyclical durables) are the subsection of industries which generate the lower earnings, so the labour market channel we discuss is distinct from any general equilibrium effects on labour markets of the cyclicity of durables. Both facts suggest that non-essentials' importance for business cycles is distinct from that of durables.

Figure B.1: Response of Durables and Non-durables over the business cycle

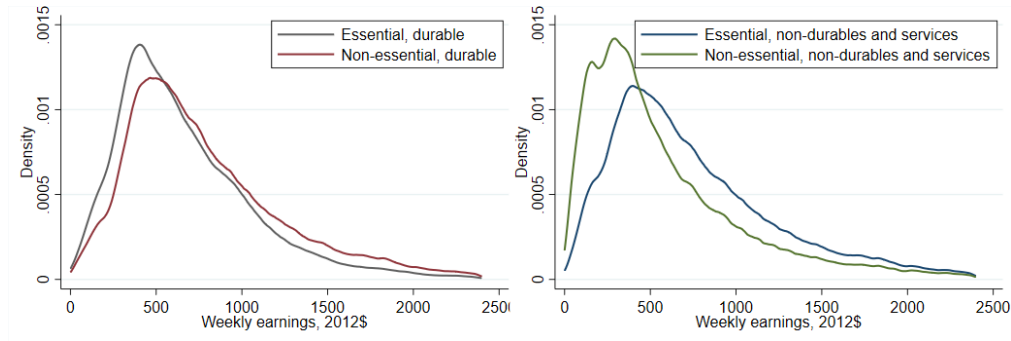


Equivalent of Figure 2 for durables and non-durables. Series starting from the peak of the previous expansion, as defined by NBER. Includes all recession peaks since 1973 where non-essential and essential series for each variable are available for a full 48 months after the peak (peaks in 1973m11, 1981m7, 1990m7, 2001m3, 2007m12 and see sample definitions in text). For consumption, employment median wages, this shows the cyclical component of the logged variable detrended using the HP filter ($\lambda = 14,440$). For median wages, the initial log series is centred 6-month rolling average, to reduce noise. For inflation, the y/y inflation rate is also detrended using the HP filter. All series are normalised to 0 at the initial period by taking the peak observation from all periods.

Share of HtM workers. In Appendix Figure B.3, we report the durable vs non-durable counterpart of Figure 4 in Section 2.2, so as to contrast the labour force compositions of durable sectors and non-essential industries. The main take-away from Figure B.3 is that despite the majority of durable industries being non-essentials (Table B.1), the share of HtM workers that produce durable goods is, on average, only 5% and does not vary much across income deciles. In contrast, the shares of HtM workers in non-durable sectors are markedly higher, varying from 47% at the bottom of the income distribution to 15% in the top decile. Together with Figure 4, these findings suggest that industries producing non-durable goods and services account for the bulk of the share of HtM workers in non-essential sectors.


Impulse response function analysis. We estimate the IRFs for subcategories of consumption and earnings to show that our main results aren't driven by a correlation between durability and non-essentiality. The results are shown in Appendix Figure B.4. Focusing on the heterogeneity between non-durable essentials and non-durable non-essentials, we find

Figure B.2: Earnings distribution - non-essentials and essentials, split into durables vs non-durables and services



Notes: Kernel density of earnings 1982-2020, pooled, for non-essentials and essentials, split by durables vs non-durables and services.

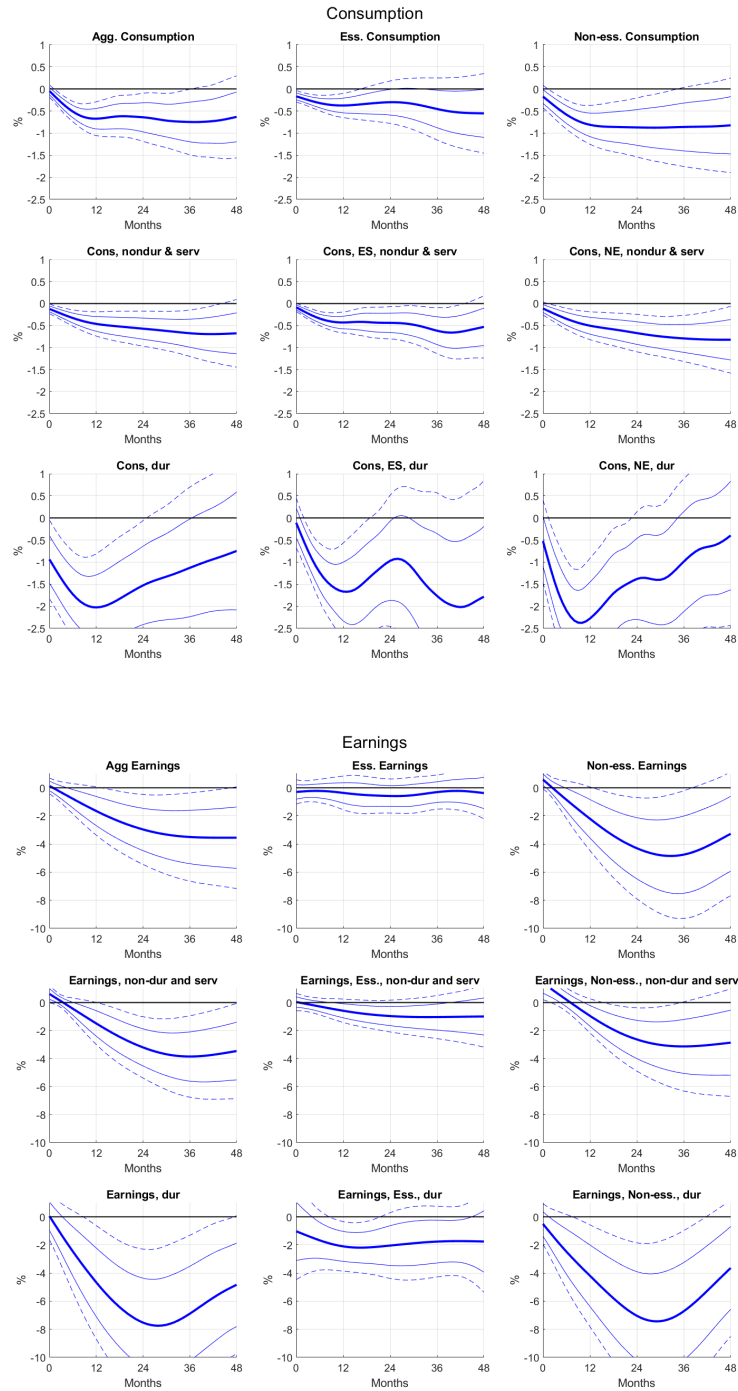
Figure B.3: Share of Hand-to-Mouth workers in durable versus non-durable sectors



Notes: Panel Study of Income Dynamics (PSID). Sample: 2003-2021.

our results remain. For consumption, non-essential non-durables falls by less than the overall non-essentials series, but still falls substantially more than essentials non-durables. For earnings, the results for non-durables are similar to consumption. Durables consumption and employment does fall substantially more than non-durables, consistent with previous findings in the literature, and durables are more cyclical than non-essentials. However, given that durables account for a minority of consumption and employment, whereas non-essentials account for a much larger share, the overall contribution of non-essentials to the cyclicity of these variables is comparable. We also still find substantial heterogeneity in labour market outcomes within durables. Taking these facts together, we view our non-essential channel as separate and distinct from the durables channel discussed previously in the literature.

Figure B.4: IRFs of consumption and earnings - non-essentials and essentials, split into durables and non-durables



Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument. The sample periods and controls are as described in the main text. 68% and 90% confidence intervals.

B.1 Other consumption categorisations

Table B.2: Alternative categorisations

	Durables	Non-durables	Goods	Services	Tradeables	Non-tradeables
Share essential	19.1	49.0	49.2	41.2	36.9	48.2
Share non-essential	80.9	51.0	50.8	58.8	63.1	51.8
Share overall	13.4	86.6	44.6	55.4	29.8	70.2

Notes: Proportion durables/non-durables, services/goods and tradeables/non-tradeables which are non-essential or essential. Calculated using PCE expenditure data; only consumption categories in the essentials/non-essentials classification are included. For durables and non-durables, we show the shares of consumption, based on chained 2000 dollar consumption series. For other categories we show expenditure shares. Proportions are weighted by average expenditure of the categories 1973-2019. Durable/nondurable and services/goods definitions from the PCE by type of product tables, tradeables defined using the classification of consumption categories provided by [Johnson \(2017\)](#).


There are other, alternative categorisations in addition to durables vs non-durables which could potentially confound our results. Two prominent examples include tradeables versus non-tradeables and good versus services. If our non-essential/essential classification strongly correlates with these alternative classifications, then then our empirical results could be caused by the correlated trait of the consumption goods, rather than by the mechanism we suggest. In this section we explore this possibility and show that it is unlikely that our results are confounded by these alternative classifications.

Table [B.2](#) shows the proportion of essentials and non-essentials that are made of up these different categories. The definitions of durables and services are from the PCE by type of product tables, while we define tradeables using the classification of consumption categories provided by [Johnson \(2017\)](#). Unlike durables, which is correlated with our non-essential classification, the proportions of services and tradables are quite similar between non-essentials and essentials. This suggests that it would be hard for a correlation with either of these characteristics to be driving our results.

To make this more concrete, [Figure B.5](#) shows IRFs for consumption of example sub-categories of consumption from the PCE, used to construct our essential and non-essential series. The charts show examples of essential and non-essential series in goods versus services. In both cases, the example non-essential consumption type falls more than the essential example.

[Jaimovich et al. \(2019\)](#) and [Jaimovich et al. \(2020\)](#) analyse the role of quality, how this declines during recessions and is positively correlated with labour intensity and skills. Our mechanism is distinct from this as we focus on consumption shifting between different types of goods, rather than quality of the same category of goods. When comparing across broad categories of goods, we instead find a negative correlation between income elasticity of demand of goods and income of workers, resulting in different implications for business cycle amplification. The two approaches are complementary but distinct. [Jaimovich et al. \(2019\)](#) and [Jaimovich et al. \(2020\)](#) study more granular consumption categories that we do, but are limited to a subset of the consumption bundle and of the labour market. On the other hand, our approach yields a coarser goods categorisation, but with the benefit of analysing the

Figure B.5: Consumption IRFs of example consumption categories, goods versus services



Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument. The sample period and controls are as described in the main text. 68% and 90% confidence intervals.

whole consumption bundle and all the labour market sectors, as we classify sectors which sell intermediate goods and services as well. Given that our data does not display the same level of granularity, we cannot interact the essential/non-essential split with the low/high quality as we did for other possible categorisations.

C Additional empirical results and robustness

This section shows some additional empirical details and checks that our results are robust to alternative choices of specifications. To complement the results in the main body, we show the IRFs of prices, other labour market variables and aggregate series. We next show that our results are robust to using alternative monetary policy shocks and more general business cycles shocks.

C.1 SLP-IV Controls

In our baseline specifications, we control for one year worth of data, with 12 lags of the 1y yields, IP, excess bond premium, log PCE price index, plus aggregate and disaggregated series for consumption and earnings depending on LHS variable. We add model-specific controls,

such as 12 lags of the dependent variable, in an effort to balance the trade-off between the benefits of lag-augmentation discussed in [Montiel Olea and Plagborg-Møller \(2021\)](#) and the cost of over-fitting.

The key idea is to include for every aggregate variable, its lags and the non-essential counterpart lags, and for every disaggregated variable both the non-essential or the essential counterpart lags. As an example, for aggregate consumption, we add aggregate and non-essential consumption. For the spending on essentials, non-essentials, and their ratio, we use both non-essential and essential consumption as controls.

We use aggregate variables as controls in other variables regressions (i.e. when the left hand side is any earning variable, we use aggregate consumption as a control). Moreover, when we control for aggregate earnings in the consumption regressions, we use the BEA NIPA series of total compensation of employees because of its longer time-series availability relative to CPS data we constructed.

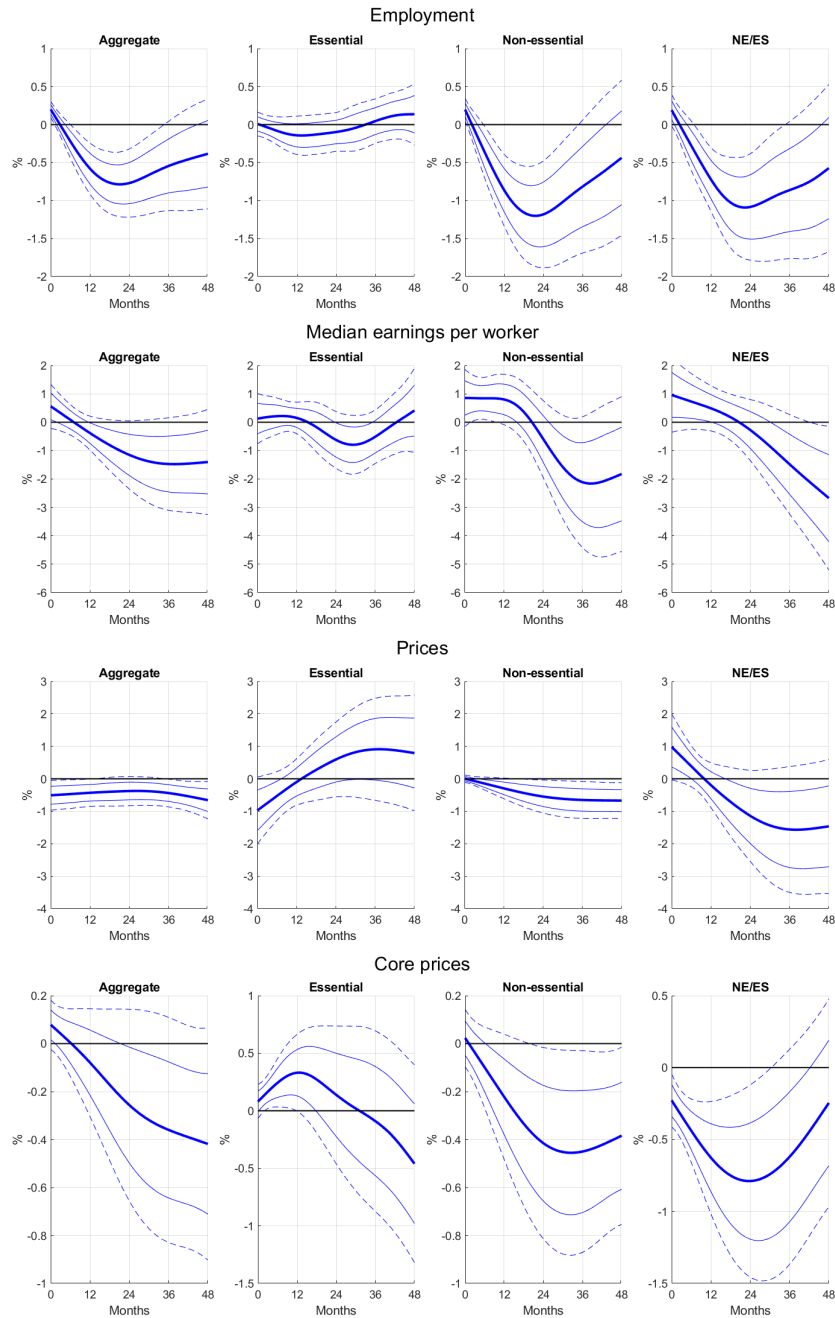
In subsection [C.2](#), we show also results for prices, employment, and median earnings; here, we detail the controls of these regressions. For prices and employment, we use our aggregate and disaggregated series in the same way we have done for other dependent variables above, in addition to aggregate consumption and the BEA NIPA compensation of employees series. For prices, we also add 12 lags of Michigan price expectations to reduce the price puzzle and an interaction of a dummy for 1978-82 with the instrument. For earnings per-worker earnings series in a given sector at different percentiles, we control also for consumption, employment, and median earnings in that sector series plus the LHS variable if not already included.

C.2 Prices, employment and median earnings responses

In addition to the consumption and earnings responses shown in the main text, [Figure C.1](#) shows the price responses and further labour market responses to a contractionary monetary policy shock. We find muted price responses; overall prices decline, with non-essential prices declining marginally more. Meanwhile essential prices appear to rise, albeit insignificantly. The heterogeneity in price responses between non-essentials and non-essentials is insignificant, however. We see this as weak evidence for a price channel of adjustment from the cyclicity of demand for non-essentials.

If we focus only on core prices, excluding food and energy categories in the same way as the BEA's core PCE series, we find somewhat stronger relative price responses. These results are also shown in [Figure C.1](#). Aggregate prices fall, albeit insignificantly and nonessential prices fall significantly, and significantly more than essential prices. These results are consistent with the results of [Orchard \(2022\)](#), who finds stronger relative price responses of luxuries compared to necessities, and who uses that controls for whether a consumption category is an energy product. This exercise excludes a number of important essential consumption categories, however. As a result, we defer to the overall responses in [Figure C.1](#) as capturing the overall, less substantial relative price adjustment mechanism.

Figure C.1: IRFs to contractionary monetary policy shock - Earnings, Employment and Prices

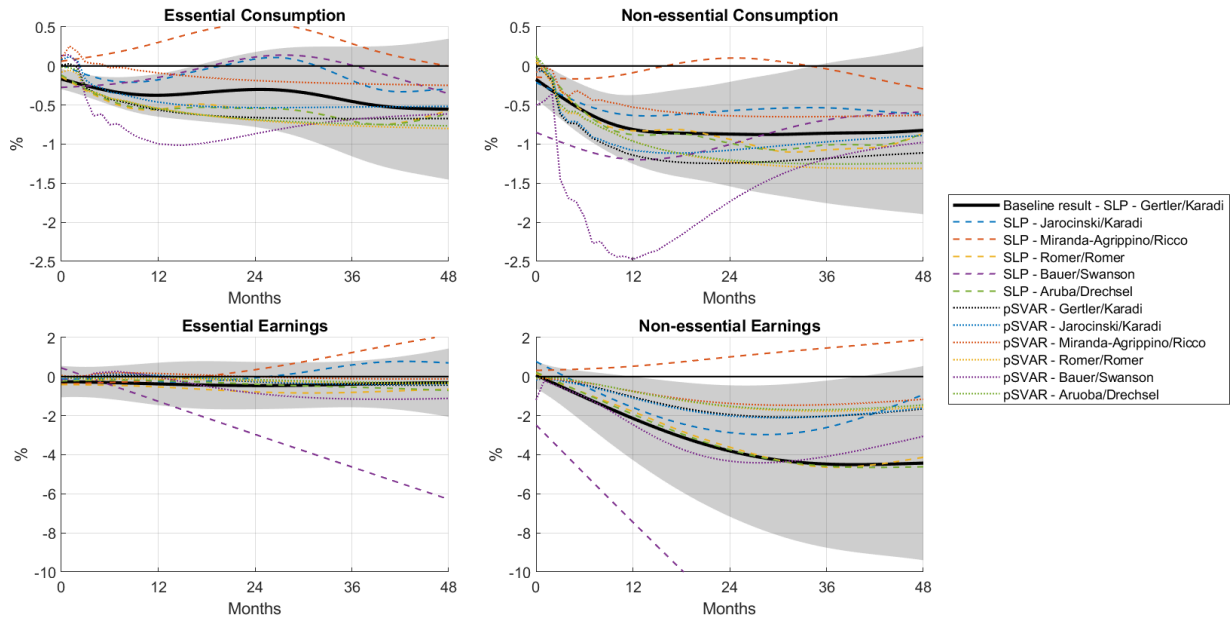


Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument. Sample periods and controls are specified in the main text and Appendix C.1.

In the main text we show the response of overall earnings in each sector to a monetary policy shock. We can deconstruct the overall earnings response separately into employment and per worker earnings responses. Both components of earnings - employment and per worker earnings - are more cyclical in the non-essential sector. The fall in overall employment peaks at just over 75bp, non-essentials employment falls more sharply, peaking at 120bp, while essentials employment does not significant fall over the entire horizon. In the second row of the same figure we can see that median per worker earnings fall more slowly, with peak responses around 3 years after the shock. Again, there is significant heterogeneity between earnings responses by category. Non-essentials earnings have a peak decline of over 2% , whereas essentials earnings only fall slightly and insignificantly. Due to the noise in the earnings data in the CPS - as it draws on the smaller ORG sample described in the main text - the overall and non-essential earnings declines responses are only significant at the 68% confidence level. Nonetheless, at the end of the horizon, the heterogeneity between non-essential and essential median earnings is significant.

C.3 Robustness to other monetary policy shocks and specifications

Figure C.2: IRFs using different of monetary policy shocks and specifications



Notes: IRFs estimated by smooth local projections and proxy-SVARs, to a range of monetary policy shocks. All responses are to a 100bp increase in 1y yields. 90% confidence intervals displayed for the baseline estimates, that of smooth local projections using the Gertler-Karadi monetary policy shock. Sample periods, controls and specification details are outlined in the text.

In our baseline results in the main text, we use the monetary policy shocks of [Gertler and Karadi \(2015\)](#) for identification. A range of alternative shocks have been proposed in recent literature, some of which include approaches to address concerns regarding the ‘information effect’, where the monetary policy announcement also conveys information about the state of


the economy that was privately held by the central bank. To ensure that our main results are robust, we use a number of these alternative shocks; those proposed by [Jarociński and Karadi \(2020\)](#)², [Miranda-Agrippino and Ricco \(2021\)](#), [Romer and Romer \(2004\)](#), [Bauer and Swanson \(2023\)](#) and [Aruoba and Drechsel \(2024\)](#). We use these instruments in the same manner as in our baseline analysis and described in the main text; 1) using these as an instrument within a proxy-SVAR, 2) extracting a monetary policy surprise series from the proxy-SVAR, and 3) using this surprise series as an instrument within a smooth local projection. The details of specifications and sample periods remain the same as the baseline results. The results of this analysis are shown in the first six series of [Figure C.2](#) and [Figure C.3](#). [Figure C.2](#) shows the point estimates of the resulting IRFs for consumption and earnings within essential and non-essentials. Most IRFs lie within the 90% confidence band of the baseline results,³ though the magnitudes of the responses vary depending on the shock used. Moreover, the figures demonstrate that the fact that non-essentials decline more than essentials – both in consumption and earnings – is preserved across a range of monetary policy instruments.

Further to these results, we also verify that our results are not specific to our baseline estimation approach, using smooth local projections. We use the same surprise series used within the smooth local projections as external instruments within a proxy-SVAR. This proxy-SVAR is designed to be as similar as possible to the SLP-IV, in particular using the same sample period. Within the proxy-SVAR, we order the surprise series first, and include 1y yields (which we use to rescale responses to that of a 100bp 1y yield increase), the excess bond premium and the essential and non-essential series of either consumption or earnings. These are estimated on the maximum possible sample, as with the SLP-IV; using monthly data 1973-2020 for consumption and 1982-2020 for earnings. [Figure C.2](#) and [Figure C.4](#) show the results of this exercise. Similarly to the SLP-IV results, the magnitude of the responses vary depending on the shock used. However, the key take-away is that the *relative* response of non-essentials is consistently larger than that of non-essentials. This is true of both earnings and consumption.

²Specifically, from [Jarociński and Karadi \(2020\)](#), we use the shock to the Fed Funds futures (FF4) if there is a negative correlation between the FF4 surprise and the SP500 surprise.


³One notable exception are the earnings responses using the [Bauer and Swanson \(2023\)](#) shocks, but [Figure C.3](#) shows that these are insignificant across the horizon.

Figure C.3: IRFs using different of monetary policy shocks - SLP specification



Notes: IRFs estimated by smooth local projections, to a range of monetary policy shocks. All responses are to a 100bp increase in 1y yields. Sample periods, controls and specification details are outlined in the text.

Figure C.4: IRFs using different of monetary policy shocks - Proxy-SVAR specification




Notes: IRFs estimated by proxy-SVARs, to a range of monetary policy shocks. All responses are to a 100bp increase in 1y yields. Sample periods, controls and specification details are outlined in the text.

C.4 Shocks from Business Cycle Anatomy (Angeletos et al. (2020))

We view our main mechanism as a general property of business cycles, not exclusive to the response to monetary policy shocks. In our main identified empirical results, we focus on responses to monetary policy shocks, both because they are an important and well-identified source of business cycle shocks and because we can draw policy implications for the conduct of monetary policy. However, as a complement to this and our unidentified IRFs in Figure 2, we also estimate the responses to more general business cycle shocks, provided by Angeletos et al. (2020). We use shock directly rather than as an instrument and using a quarterly frequency, similarly to that paper. We specifically use the shock from Angeletos et al. (2020) which maximises the variation in unemployment. Shocks which maximise other key macro-variables are strongly correlated with this shock, and give similar results in our setting. We use the same controls and the same starting sample as in the monetary policy regressions. We end the sample in 2017q4 due to data availability. Results are shown in Figure C.5. As in our main results, we find that the decline in non-essential consumption and earnings is larger than that of essentials, particularly in the first two years after the shock. We corroborate the message that Non-Essential Business Cycles are a broad phenomenon of business cycle amplification, and it is not only specific to monetary policy.

Figure C.5: Response to business cycle shock from Angeletos et al. (2020)



Notes: Based on the shock that maximises the variation in unemployment, which is the main shock used by Angeletos et al. (2020). Results estimated using SLP, same specification as the main results, other than no instrument used and quarterly frequency and sample ending in 2017q4.

D Model description

In this appendix, we provide detailed descriptions for the theoretical model. Further details of the solution method, the steady state computation, and the log-linear equilibrium are included in Online Appendix F.

D.1 Households

We solve separately the Ricardian agent problem and the hand-to-mouth one. Inattention allows us to match the hump shape response in consumption, while maintaining the relative IES across different goods.

D.1.1 Ricardian agent problem

Unconstrained agents can invest in nominal bonds $B_{H,t}$ that earn risk free nominal rate R_t . Their nominal budget constrain is:

$$P_t^E C_{H,t}^E + P_t^N C_{H,t}^N + B_{H,t} \leq W_{H,t} N_{H,t} + \Pi_{H,t} + T_{H,t} + R_{t-1} B_{H,t-1}$$

We can rewrite the budget constraint defining wealth in terms of the essential price $a_{H,t}$:

$$\begin{aligned} a_{H,t} &= b_{H,t-1} \frac{R_{t-1}}{\pi_t^E} + w_{H,t} N_{H,t} + \Pi_{H,t}^r + t_{H,t} \\ a_{H,t+m+1} &= \prod_{k=0}^m \tilde{R}_{t+k+1} a_{H,t} - \sum_{j=0}^m \prod_{k=j}^m \tilde{R}_{t+k+1} (C_{H,t+j}^E + p_{t+j}^N C_{H,t+j}^N) \\ &\quad + \sum_{j=0}^m \prod_{k=j+1}^m \tilde{R}_{t+k+1} (w_{t+j+1} N_{H,t+j+1} + \Pi_{H,t+j+1}^r + t_{H,t+j+1}) \end{aligned}$$

Where $\pi_{t+1}^E \equiv \frac{P_{t+1}^E}{P_t^E}$ is the inflation of essential goods, and similarly for non-essentials, $\tilde{R}_{t+1} \equiv R_t / \pi_{t+1}^E$ is real ex-post rate in terms of the essential price inflation. All lower case variables are the corresponding uppercase variable in terms of the essential price: $p_t^N \equiv \frac{P_t^N}{P_t^E}$, $w_{H,t} \equiv \frac{W_{H,t}}{P_t^E}$, $t_{H,t} \equiv \frac{T_{H,t}}{P_t^E}$, and $b_{H,t} \equiv \frac{B_{H,t}}{P_t^E}$. We define $\Pi_{H,t}^r \equiv \frac{\Pi_{H,t}}{P_t^E}$ as real profits to avoid confusion with inflation.

We now turn to the Bellman equation. Households update their expectations only sporadically. Specifically, they update with probability λ . Somebody who updates today has probabilities λ of updating tomorrow, $\lambda(1 - \lambda)$ of updating in 2 periods, $\lambda(1 - \lambda)^2$ in 3 periods, $\lambda(1 - \lambda)^j$ in $j + 1$ periods, and so on. When they update, the problem is as in year zero, so we use the recursive structure to solve their choice. As they realise that they might not be able to update, households make plans for future choices in the current period. They choose consumption of a variety, say essentials, for today: $C_{i,t,0}^E$ and for the future if they don't update $C_{i,t+j,j}^E$ for j periods ahead, and similarly for non-essential consumption and savings. As households delegate the labour choice to unions, we can ignore the disutility of

labour in the household problem.

$$\begin{aligned}
V(a_{H,t}) &= \max_{\{C_{H,t+m,m}^E, C_{H,t+m,m}^N\}_{m=0}^{\infty}} \left(\sum_{m=0}^{\infty} \beta^m (1-\lambda)^m \left(\frac{(C_{H,t+m,m}^E)^{1-\frac{1}{\gamma^E}}}{1-\frac{1}{\gamma^E}} + \varphi \frac{(C_{H,t+m,m}^N)^{1-\frac{1}{\gamma^N}}}{1-\frac{1}{\gamma^N}} \right) \right. \\
&\quad \left. + \beta \lambda \sum_{m=0}^{\infty} \beta^m (1-\lambda)^m \mathbb{E}_t V(a_{H,t+m+1}) \right) \\
s.t. \ a_{H,t+m+1} &= \prod_{k=0}^m \tilde{R}_{t+k+1} a_{H,t} - \sum_{j=0}^m \prod_{k=j}^m \tilde{R}_{t+k+1} (C_{H,t+j}^E + p_{t+j}^N C_{H,t+j}^N)
\end{aligned}$$

The household makes plans for when they cannot update (first terms) and for when they can update, then the problem becomes the same (second terms). Start by taking the FOC for the essential good:

$$\begin{aligned}
\frac{\partial V(a_{H,t})}{\partial C_{H,t+j,j}^E} &= \beta^j (1-\lambda)^j (C_{H,t+j,j}^E)^{-\frac{1}{\gamma^E}} - \beta \lambda \sum_{m=j}^{\infty} \beta^m (1-\lambda)^m \mathbb{E}_t V'(a_{H,t+m+1}) \frac{\partial V(a_{H,t+m+1})}{\partial C_{H,t+j,j}^E} = 0 \\
\frac{\partial V(a_{H,t})}{\partial C_{H,t+j,j}^E} &= \beta^j (1-\lambda)^j (C_{H,t+j,j}^E)^{-\frac{1}{\gamma^E}} - \beta \lambda \sum_{m=j}^{\infty} \beta^m (1-\lambda)^m \mathbb{E}_t V'(a_{H,t+m+1}) \prod_{k=j}^m \tilde{R}_{t+k+1} = 0
\end{aligned}$$

Rewrite the FOC for the current period and use the expression to express compactly the envelope condition:

$$\begin{aligned}
\frac{\partial V(a_{H,t})}{\partial C_{H,t,0}^E} &= (C_{H,t,0}^E)^{-\frac{1}{\gamma^E}} - \beta \lambda \sum_{m=0}^{\infty} \beta^m (1-\lambda)^m \mathbb{E}_t V'(a_{H,t+m+1}) \prod_{k=0}^m \tilde{R}_{t+k+1} = 0 \\
V'(a_{H,t}) &= \beta \lambda \sum_{m=0}^{\infty} \beta^m (1-\lambda)^m \mathbb{E}_t V'(a_{H,t+m+1}) \frac{\partial V(a_{H,t+m+1})}{\partial a_{H,t}} = (C_{H,t,0}^E)^{-\frac{1}{\gamma^E}}
\end{aligned}$$

This means we can rewrite the FOC with the envelope condition plugged in, so that the choice of the attentive consumer in this period is a function of the expected choices of the attentive consumer in the future:

$$\begin{aligned}
\frac{\partial V(a_{H,t})}{\partial C_{H,t+j,j}^E} &= \beta^j (1-\lambda)^j (C_{H,t+j,j}^E)^{-\frac{1}{\gamma^E}} - \beta \lambda \sum_{m=j}^{\infty} \beta^m (1-\lambda)^m \mathbb{E}_t (C_{H,t+m+1,0}^E)^{-\frac{1}{\gamma^E}} \prod_{k=j}^m \tilde{R}_{t+k+1} = 0 \\
\frac{\partial V(a_{H,t})}{\partial C_{H,t,0}^E} &= (C_{H,t,0}^E)^{-\frac{1}{\gamma^E}} - \beta \lambda \sum_{m=0}^{\infty} \beta^m (1-\lambda)^m \mathbb{E}_t (C_{H,t+m+1,0}^E)^{-\frac{1}{\gamma^E}} \prod_{k=0}^m \tilde{R}_{t+k+1} = 0
\end{aligned}$$

Use the recursive structure to write the FOC as a traditional Euler equation and a condition for the j periods ahead plan.

$$(C_{H,t,0}^E)^{-\frac{1}{\gamma^E}} = \beta \mathbb{E}_t (C_{H,t+1,0}^E)^{-\frac{1}{\gamma^E}} \tilde{R}_{t+1} \quad (C_{H,t+j,j}^E)^{-\frac{1}{\gamma^E}} = \mathbb{E}_t (C_{H,t+j,0}^E)^{-\frac{1}{\gamma^E}}$$

We solved the essential good choice, we now turn to the non-essential good choice. The FOC:

$$\frac{\partial V(a_{H,t})}{\partial C_{H,t+j,j}^N} = \beta^j (1-\lambda)^j (C_{H,t+j,j}^N)^{-\frac{1}{\gamma^N}} \varphi - \beta \lambda \sum_{m=j}^{\infty} \beta^m (1-\lambda)^m \mathbb{E}_t V'(a_{H,t+m+1}) \prod_{k=j}^m \tilde{R}_{t+k+1} p_{t+j}^N = 0$$

For the N good we need to keep track of the relative price. However, the solution is easier as we can use directly the essential good solutions.

$$\begin{aligned} \beta^j (1-\lambda)^j (C_{H,t+j,j}^N)^{-\frac{1}{\gamma^N}} \varphi &= \beta \lambda \sum_{m=j}^{\infty} \beta^m (1-\lambda)^m \mathbb{E}_t V'(a_{H,t+m+1}) \prod_{k=j}^m \tilde{R}_{t+k+1} p_{t+j}^N \\ (C_{H,t+j,j}^N)^{-\frac{1}{\gamma^N}} &= \mathbb{E}_t (C_{H,t+j,0}^N)^{-\frac{1}{\gamma^N}} \end{aligned}$$

We can summarise the problems of the Ricardian agents with four equilibrium conditions, a budget constraint, and two aggregation equations. The equilibrium conditions consist of: an Euler equation for the attentive consumer in terms of the essential good, an intra-temporal condition linking consumption of essential goods to non-essential goods for an attentive consumer,⁴ and two conditions, one for essential goods and one for non-essential goods, linking the consumption plans for consumers who do not update to the expectation of what an attentive consumer would do.

$$\begin{aligned} (C_{H,t,0}^E)^{-\frac{1}{\gamma^E}} &= \beta \mathbb{E}_t \left((C_{H,t+1,0}^E)^{-\frac{1}{\gamma^E}} \frac{R_t}{\pi_{t+1}^E} \right) & \varphi (C_{H,t,0}^N)^{-\frac{1}{\gamma^N}} &= p_t^N (C_{H,t,0}^E)^{-\frac{1}{\gamma^E}} \\ (C_{H,t+j,j}^E)^{-\frac{1}{\gamma^E}} &= \mathbb{E}_t \left((C_{H,t+j,0}^E)^{-\frac{1}{\gamma^E}} \right) & (C_{H,t+j,j}^N)^{-\frac{1}{\gamma^N}} &= \mathbb{E}_t \left((C_{H,t+j,0}^N)^{-\frac{1}{\gamma^N}} \right) \end{aligned}$$

The Ricardian agents budget constraint, which drops out in the equilibrium definition due to Walras law $C_{H,t}^E + p_t^N C_{H,t}^N + b_{H,t} = w_{H,t} N_{H,t} + \Pi_{H,t}^r + t_{H,t} + R_{t-1} b_{H,t-1} \frac{R_{t-1}}{\pi_t^E}$. Consumption aggregation across attentive and non-attentive consumers:

$$C_{H,t}^E = \lambda \sum_{j=0}^{\infty} (1-\lambda)^j C_{H,t-j,j}^E \quad C_{H,t}^N = \lambda \sum_{j=0}^{\infty} (1-\lambda)^j C_{H,t-j,j}^N$$

We plug-in the last FOC to express overall consumption as a function of the expected actions of attentive consumers:

$$C_{H,t}^E = \lambda \sum_{j=0}^{\infty} (1-\lambda)^j \left[\mathbb{E}_{t-j} \left((C_{H,t,0}^E)^{-\frac{1}{\gamma^E}} \right) \right]^{-\gamma^E}$$

⁴Notice that the intra-temporal condition governs the elasticity of substitution across the two goods. Specifically, this elasticity is bounded between γ^E and γ^N . We can see this by manipulating the condition

and show: $-\frac{\partial \left(\frac{C_{H,t,0}^N}{C_{H,t,0}^E} \right)}{\partial p_t^N} \frac{p_t^N}{\left(\frac{C_{H,t,0}^N}{C_{H,t,0}^E} \right)} = \gamma^E - \frac{\gamma^N - \gamma^E}{\gamma^N} \frac{\partial C_{H,t,0}^N}{\partial p_t^N} \frac{p_t^N}{C_{H,t,0}^N} = \gamma^N - \frac{\gamma^N - \gamma^E}{\gamma^E} \frac{\partial C_{H,t,0}^E}{\partial p_t^N} \frac{p_t^N}{C_{H,t,0}^E}$. As $\frac{\partial C_{H,t,0}^N}{\partial p_t^N} \leq 0$ and

$\frac{\partial C_{H,t,0}^E}{\partial p_t^N} \geq 0$ we can see that the elasticity of substitution across the two goods is bounded between γ^E and γ^N .

$$C_{H,t}^N = \lambda \sum_{j=0}^{\infty} (1-\lambda)^j \left[\mathbb{E}_{t-j} \left((C_{H,t,0}^N)^{-\frac{1}{\gamma^N}} \right) \right]^{-\gamma^N}$$

We can define the average marginal utility of consumption aggregated across attentive and inattentive agents, to use as an objective function for the union:⁵

$$\zeta_{H,t} = (C_{H,t}^E)^{-\frac{1}{\gamma^E}} \frac{C_{H,t}^E}{C_{H,t}^E + C_{H,t}^N p_t^N} + (C_{H,t}^N)^{-\frac{1}{\gamma^N}} \frac{\varphi}{p_t^N} \frac{C_{H,t}^N p_t^N}{C_{H,t}^E + C_{H,t}^N p_t^N}$$

D.1.2 Hand-to-mouth agent problem

Constrained agents face the same problem, with the same information friction, but do not have access to bond markets. They make plans for consumption choices in the future, as they can also be inattentive, but do not have saving choices to smooth out inconsistent plans as the Ricardian agents. Therefore, we posit a risk sharing agreement across hand-to-mouth households, to ensure that each household follows ex-post their consumption plans and the overall hand-to-mouth agents budget constraint is satisfied⁶.

First, we show the budget constraint in terms of wealth: $C_{L,t}^E + p_t^N C_{L,t}^N \leq a_{L,t} = w_{L,t} N_{L,t} + \Pi_{L,t}^r + t_{L,t}$. Their maximisation problem, for the periods in which they cannot update:

$$V(a_{L,t}) = \max_{\{C_{L,t+m,m}^E, C_{L,t+m,m}^N\}_{m=0}^{\infty}} \sum_{m=0}^{\infty} \beta^m (1-\lambda)^m \left(\frac{(C_{L,t+m,m}^E)^{1-\frac{1}{\gamma^E}}}{1-\frac{1}{\gamma^E}} + \varphi \frac{(C_{L,t+m,m}^N)^{1-\frac{1}{\gamma^N}}}{1-\frac{1}{\gamma^N}} + \eta_{t+j} \mathbb{E}_t (a_{L,t+m} - C_{L,t+m,m}^E - C_{L,t+m,m}^N p_{t+m}^N) \right)$$

To find the solution, take the FOC for the two goods and equate the marginal utilities to arrive to the three equilibrium conditions as for the Ricardian agents, minus the Euler equation:

$$\begin{aligned} \varphi (C_{L,t,0}^E)^{-\frac{1}{\gamma^E}} &= (C_{L,t,0}^N)^{-\frac{1}{\gamma^N}} \frac{1}{p_t^N} \\ (C_{L,t+j,j}^E)^{-\frac{1}{\gamma^E}} &= \mathbb{E}_t (C_{L,t+j,0}^E)^{-\frac{1}{\gamma^E}} & (C_{L,t+j,j}^N)^{-\frac{1}{\gamma^N}} &= \mathbb{E}_t (C_{L,t+j,0}^N)^{-\frac{1}{\gamma^N}} \end{aligned}$$

We can still aggregate goods consumption across attentive and non-attentive consumers. By assuming risk sharing across consumers, agents can follow through with their plans ex-post. Moreover, we can still define the average marginal utility of consumption.

$$C_{L,t}^E = \lambda \sum_{j=0}^{\infty} (1-\lambda)^j \left[\mathbb{E}_{t-j} \left((C_{L,t,0}^E)^{-\frac{1}{\gamma^E}} \right) \right]^{-\gamma^E}$$

⁵Notice that, for attentive households, this expression simplifies to the standard marginal utility: $\zeta_{H,t,0} = (C_{H,t,0}^E)^{-\frac{1}{\gamma^E}} = (C_{H,t,0}^N)^{-\frac{1}{\gamma^N}} \frac{\varphi}{p_t^N}$.

⁶This is a shortcut for the idea that there is a government in the background that provides insurance across hand-to-mouth agents. With this assumption, we avoid transfers across agent types, which would affect the transmission mechanism.

$$C_{L,t}^N = \lambda \sum_{j=0}^{\infty} (1-\lambda)^j \left[\mathbb{E}_{t-j} \left((C_{L,t,0}^N)^{-\frac{1}{\gamma^N}} \right) \right]^{-\gamma^N}$$

$$\zeta_{L,t} = (C_{L,t}^E)^{-\frac{1}{\gamma^E}} \frac{C_{L,t}^E}{C_{L,t}^E + C_{L,t}^N p_t^N} + (C_{L,t}^N)^{-\frac{1}{\gamma^N}} \frac{\varphi}{p_t^N} \frac{C_{L,t}^N p_t^N}{C_{L,t}^E + C_{L,t}^N p_t^N}$$

D.2 Unions

Unions are perfectly competitive and fully attentive. There are two unions, one to represent each type of consumer, Ricardian and hand-to-mouth. We follow [Mankiw and Reis \(2007\)](#) in separating the consumption choice from the labour supply choice, by employing unions. As each union represents the totality of the family, they take overall marginal utility of consumption to compute the labour supply choice. We end up with two standard intra-temporal equilibrium conditions:

$$\xi \frac{N_{L,t}^X}{\zeta_{H,t}} = w_{L,t} \qquad \xi \frac{N_{H,t}^X}{\zeta_{L,t}} = w_{H,t}$$

D.3 Firms

Final good producers. The final good producers combine different retail varieties of the essential and of the non-essential goods according to a CES aggregator. $Y_t^i = \left(\int_0^1 (y_{k,t}^i)^{\frac{\varepsilon-1}{\varepsilon}} dk \right)^{\frac{\varepsilon}{\varepsilon-1}}$ for $i = \{E, N\}$. This leads to a standard demand that the final good producers have for different varieties of a given good category (essential and non-essentials): $y_{k,t}^i = Y_t^i \left(\frac{P_{k,t}^i}{P_t^i} \right)^{-\varepsilon}$ for $i = \{E, N\}$.

Calvo retailers. Retailers of a given type of good, say essential⁷, buy a wholesale good of the same type at a wholesale price $P_t^{E,w}$ and use it to produce the retail variety $y_{k,t}^E$ with a linear technology that maps one-to-one the wholesale good to the retail variety. As each variety is differentiated they have market power and face a Calvo friction to change prices. Their real marginal cost $\mathcal{S}_t^E = \frac{P_t^{E,w}}{P_t^E}$ is the wholesale price relative to its retail average value. They receive a subsidy τ^E for each unit of good they produce and pay lump sum taxes T_t^E ; these taxes allow to have zero profit in steady state but do not affect the profit allocation off-steady state. The probability of not being able to reset prices is equal to θ in each period. This leads to a standard non-linear New-Keynesian Phillips Curve. The objective function, the discounted present value of profits is:

$$\mathbb{E}_t \sum_{j=0}^{\infty} SDF_{t,t+j} (P_{k,t+j}^E Y_{k,t+j}^E - P_{t+j}^E \mathcal{S}_{t+j}^E Y_{k,t+j}^E)$$

All firms that in period t can reset their price face the same problem (this happens with probability $1 - \theta$), therefore will choose the same price \tilde{P}_t^E , that maximises profits as long

⁷The non-essential retailers one is fully symmetric.

as it remains in place: $\mathbb{E}_t \sum_{j=0}^{\infty} SDF_{t,t+j}(\theta)^j \left(\tilde{P}_t^E Y_{k,t+j}^E - P_{t+j}^E (1 - \tau^E) \mathcal{S}_{t+j}^E Y_{k,t+j}^E - T_{t+j}^E \right)$. We substitute the demand equation, take the first order condition, substitute-in the SDF for the H household⁸, and rearrange to arrive to three equations representing the non-linear New Keynesian Phillips Curve.

$$\begin{aligned} K_t^{E,f} &= (C_{H,t}^E)^{-\frac{1}{\gamma^E}} Y_t^E \mathcal{S}_t^E \frac{\varepsilon^E}{\varepsilon^E - 1} (1 - \tau^E) + \theta \beta \mathbb{E}_t (\pi_{t+1}^E)^{\varepsilon^E} K_{t+1}^{E,f} \\ F_t^{E,f} &= (C_{H,t}^E)^{-\frac{1}{\gamma^E}} Y_t^E + \theta \beta \mathbb{E}_t (\pi_{t+1}^E)^{\varepsilon^E - 1} F_{t+1}^{E,f} \\ \frac{K_t^{E,f}}{F_t^{E,f}} &= \left(\frac{1 - \theta (\pi_t^E)^{\varepsilon^E - 1}}{1 - \theta} \right)^{\frac{1}{1 - \varepsilon^E}} \end{aligned}$$

For the non-essential goods we can solve the same problem and arrive to the same equations with the subscript being N rather than E .

D.3.1 Wholesalers

Wholesalers produce one type of good, essentials or non-essentials, are perfectly competitive and they combine high-skill labour $N_{H,t}^i$ and low-skill labour $N_{L,t}^i$ with technology:

$$Y_t^E = A_t^E (N_{L,t}^E)^{\alpha^E} (N_{H,t}^E)^{1 - \alpha^E} \quad Y_t^N = A_t^N (N_{L,t}^N)^{\alpha^N} (N_{H,t}^N)^{1 - \alpha^N}$$

They sell these goods at nominal price $P_t^{i,w}$ to retailers. They pay nominal wage $W_{H,t}$ for each unit of high-skilled household labour and nominal $W_{L,t}$ for each unit of low-skilled household labour. The low-skilled share in production is α^i . The crucial innovation we present is that $\alpha^E < \alpha^N$: there are relatively more low-skilled workers in non-essential goods production than in essential goods production. As discussed in the main text, this is the source of labour market heterogeneity and the resulting amplification.

The solution to the problem of the essential and the non-essential wholesalers are:

$$\mathcal{S}_t^E \alpha^E \frac{Y_t^E}{N_{L,t}^E} = w_{L,t}, \quad \mathcal{S}_t^E (1 - \alpha^E) \frac{Y_t^E}{N_{H,t}^E} = w_{H,t}, \quad \mathcal{S}_t^N \alpha^N \frac{Y_t^N}{N_{L,t}^N} = \frac{w_{L,t}}{p_t^N}, \quad \mathcal{S}_t^N (1 - \alpha^N) \frac{Y_t^N}{N_{H,t}^N} = \frac{w_{H,t}}{p_t^N}$$

D.4 Market clearing

We close the model with two goods market clearing condition, for essential and non-essential goods, two labour market clearing conditions, for high and low skilled labour, and bond market clearing condition by which bonds are in zero net supply. In this economy the population is divided in the two types of households with total mass equal to one: $1 = \mu_H + \mu_L$.

⁸As we take a first order Taylor approximation to solve the model, the choice of whose SDF we take drops out. However, we need to specify to whom the off-steady state profits are allocated as this affects the propagation mechanism in a heterogeneous agents model: we specify a profit allocation rule directly. Notice that one could use indifferently the SDF for essential and non-essential goods, as the SDFs for the two goods are equal in each state of nature for a given agent. We use the SDF for essentials for the essential good producers and the one for non-essentials for the non-essential retailer.

The market clearing conditions for the two goods markets:

$$Y_t^E = C_t^E = \sum_{i=\{H,L\}} \mu_i C_{i,t}^E \qquad Y_t^N = C_t^N = \sum_{i=\{H,L\}} \mu_i C_{i,t}^N$$

The labour market clearing conditions for the two types of labour:

$$N_{H,t}^E + N_{H,t}^N = \mu_H N_{H,t} \qquad N_{L,t}^E + N_{L,t}^N = \mu_L N_{L,t}$$

The bonds market clearing specifies that bonds are in zero net supply:

$$\mu_H B_{H,t} = 0$$

In this model, with non-homothetic preferences, we cannot construct an ideal price index, so we model CPI inflation as statistical agencies do, with Laspeyres, Paasche, or Fisher price indices. We define CPI inflation to be inflation computed with the Fisher index. When log-linearised, all these indices simplify to inflation being a weighted average of essential inflation and non-essential inflation, weighted by the economy wide steady state consumption shares.

$$\begin{aligned} \pi_{t,Lasp} &= \frac{P_t^E C_{t-1}^E + P_t^N C_{t-1}^N}{P_{t-1}^E C_{t-1}^E + P_{t-1}^N C_{t-1}^N} & \pi_{t,Paasche} &= \frac{P_t^E C_t^E + P_t^N C_t^N}{P_{t-1}^E C_t^E + P_{t-1}^N C_t^N} \\ \pi_{t,Fisher} &= (\pi_{t,Lasp} \pi_{t,Paasche})^{1/2} \equiv \pi_t \end{aligned}$$

We compute real GDP with production in the two sectors weighted by prices in steady state, with P^E being normalised to one: $Y_t = Y_t^E + p^N Y_t^N$.

D.5 Government

The government consists of a central bank that sets interest rates according to a Taylor rule:

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R} \right)^{\rho_R} \left((\mathbb{E}_t(\pi_{t+1}))^{\phi_\pi} \left(\frac{Y_t}{Y} \right)^{\phi_Y} \right)^{1-\rho_R} \exp(\sigma^{mp} \varepsilon_t^{mp})$$

The only role of fiscal policy in the baseline model is to ensure that Calvo retailers profits are zero in steady state. The government sets a lump sum tax on each Calvo retailer such that it pays in a non-distortive way for the subsidy to the same retailer. With this tax, retailers profits are zero in steady state.

$$T_t^E = \tau^E P_t^E \mathcal{S}_t^E Y_t^E \qquad T_t^N = \tau^N P_t^N \mathcal{S}_t^N Y_t^N$$

With $\tau^E = 1/\varepsilon^E$ and $\tau^N = 1/\varepsilon^N$. Taxes to households are zero and there is no government spending. Therefore, the government runs a balanced budget. We specify a profit allocation rule off steady state, where we give profits to Ricardian households in our baseline model in the spirit of [Bilbiie \(2008\)](#) or [Debortoli and Galí \(2017\)](#). We explore alternative profit

allocation mechanism in the counterfactual exercise. The generic transfer policy⁹:

$$\Pi_{k,t} = \phi_{\Pi,k}^E \Pi_t^E + \phi_{\Pi,k}^N \Pi_t^N \quad k = \{H, L\}$$

E Unconventional fiscal policy

In this appendix, we present the derivation of the results on unconventional fiscal policy with heterogeneous goods and heterogeneous households. The objective of this appendix is twofold. First, as a contribution on its own, we extend to a two-agents, two-goods economy the unconventional fiscal policy results that [Correia et al. \(2013\)](#) present in a representative agent economy with only one good. Second, we want to study, in this framework, the effect of changes in goods-specific VAT rates. To achieve these aims, we adapt the model to allow for VAT shocks, both economy-wide and sector specific as well as the corresponding changes in payroll taxes. Next, we show the equivalence between a VAT shock properly compensated with payroll taxes, and possibly transfers changes, and a monetary policy shock.

E.1 A model with distortionary taxation

We extend the model we presented in Section 4 and Appendix Section D to include VAT taxes, payroll taxes, and household transfers. Households pay VAT taxes, which can be either uniform across all goods or goods-specific. To buy one unit of an essential good, households pay $P_t^E(1 + \tau_t^{VAT} + \tau_t^{VATE})$ in nominal terms, and to buy one unit of a non-essential good, they pay $P_t^N(1 + \tau_t^{VAT} + \tau_t^{VATN})$ in nominal terms. τ_t^{VAT} is a standard VAT tax, as it applies uniformly on both consumption goods, whereas τ_t^{VATE} only applies to essential goods and τ_t^{VATN} to non-essential goods. Following [Correia et al. \(2013\)](#), we assume that producer prices are sticky (P_t^N and P_t^E), whereas the prices that consumer face can vary, given the producer price and the consumption taxes. In addition, households face payroll taxes on their labour earnings. Ricardian households receive $(1 - \tau_{H,t}^{Payroll})W_{H,t}N_{H,t}$ and hand-to-mouth households receive $(1 - \tau_{L,t}^{Payroll})W_{L,t}N_{L,t}$, all in nominal terms. As in [Correia et al. \(2013\)](#), the role of payroll taxes is to neutralise the effect that VAT taxes have on the labour supply decision.

$$\begin{aligned} -W_{H,t}N_{H,t}\tau_{H,t}^{Payroll} &= (C_{H,t}^E P_t^E + C_{H,t}^N P_t^N)\tau_t^{VAT} + C_{H,t}^E P_t^E \tau_t^{VATE} + C_{H,t}^N P_t^N \tau_t^{VATN} \\ -W_{L,t}N_{L,t}\tau_{L,t}^{Payroll} &= (C_{L,t}^E P_t^E + C_{L,t}^N P_t^N)\tau_t^{VAT} + C_{L,t}^E P_t^E \tau_t^{VATE} + C_{L,t}^N P_t^N \tau_t^{VATN} \end{aligned}$$

Next, we set up lump sum transfers to ensure that taxation does not affect redistribution across household, which could break the mapping with monetary policy. We note however, that the way we set up the payroll tax response is such that lump sum transfers are always zero.

$$\begin{aligned} \frac{T_{H,t}}{\mu_H} &= (C_{H,t}^E P_t^E + C_{H,t}^N P_t^N)\tau_t^{VAT} + C_{H,t}^E P_t^E \tau_t^{VATE} + C_{H,t}^N P_t^N \tau_t^{VATN} + W_{H,t}N_{H,t}\tau_{H,t}^{Payroll} \\ \frac{T_{L,t}}{\mu_L} &= (C_{L,t}^E P_t^E + C_{L,t}^N P_t^N)\tau_t^{VAT} + C_{L,t}^E P_t^E \tau_t^{VATE} + C_{L,t}^N P_t^N \tau_t^{VATN} + W_{L,t}N_{L,t}\tau_{L,t}^{Payroll} \end{aligned}$$

⁹Recall capital letter Π are profits, small letter π are inflation rates.

The final step is to set up the laws of motion for VAT taxes:

$$\begin{aligned}\tau_t^{VAT} &= \rho^{VAT} \tau_{t-1}^{VAT} + \sigma^{VAT} \varepsilon_t^{VAT} \\ \tau_t^{VATE} &= \rho^{VATE} \tau_{t-1}^{VATE} + \sigma^{VATE} \varepsilon_t^{VATE} \\ \tau_t^{VATN} &= \rho^{VATN} \tau_{t-1}^{VATN} + \sigma^{VATN} \varepsilon_t^{VATN}\end{aligned}$$

In order to have no deviation from the baseline model for a monetary policy shock, we solve the model with a first order Taylor approximation around a zero tax steady state (in addition to zero inflation as in the baseline model), that is $\tau^{VAT} = \tau^{VATE} = \tau^{VATN} = \tau_H^{Payroll} = \tau_L^{Payroll} = T_H = T_L = 0$ in steady state. For this reason, we linearise these tax variables as we do for profits, e.g. $\hat{\tau}_t^{VAT} = \tau_t^{VAT} - \tau^{VAT}$. The full equilibrium is characterised in web Appendix L.

E.2 Equivalence between monetary policy and unconventional fiscal policy

Proposition 1 *Consider the linearised model of Sections E.1 that features VAT taxes, payroll taxes, and lump sum transfers. Changes in interest rates can be replicated in an observationally equivalent way with changes in uniform VAT tax rates, compensated by opposite changes in payroll taxes.*

We now prove this statement. The intuition is that changes in the growth rate of uniform VAT rate enter the Euler equation as interest rates do. Uniform VAT taxes further only affect the labour supply decision by the unions and the budget constraints of the agents, with payroll taxes that nullify both these distortive effects.

Proof of Proposition 1.

The proof of this proposition follows by analysing the private sector choices of this economy. Note that, to simplify the notation, we set changes in goods-specific VAT rates to zero, as they are exogenously specified, so do not react to changes in interest rates or uniform VAT tax rate changes. First, note that the interest rate \hat{R}_t enters the equilibrium conditions of any private sector entity (households, unions, or firms) only in the Euler equation. Moreover, uniform VAT tax rate changes enter the Euler equation in the same way:

$$\frac{1}{\gamma^E} \mathbb{E}_t \left(\hat{C}_{H,t+1,0}^E \right) = \frac{1}{\gamma^E} \hat{C}_{H,t,0}^E - \mathbb{E}_t(\hat{\pi}_{t+1}^E) + \hat{R}_t + \hat{\tau}_t^{VAT} - \mathbb{E}_t(\hat{\tau}_{t+1}^{VAT})$$

Uniform VAT tax rates also enter in the marginal utility that the union uses to set the labour supply. If we substitute out the $\hat{\varsigma}$ s:

$$\begin{aligned}\chi \hat{N}_{H,t} + \hat{\tau}_t^{VAT} + \frac{1}{\gamma^E} \hat{C}_{H,t}^E (1 - \bar{C}_H^N) + \left(\frac{1}{\gamma^N} \hat{C}_{H,t}^N + \hat{p}_t^N \right) \bar{C}_H^N &= \hat{w}_{H,t} - \hat{\tau}_{H,t}^{Payroll} \\ \chi \hat{N}_{L,t} + \hat{\tau}_t^{VAT} + \frac{1}{\gamma^E} \hat{C}_{L,t}^E (1 - \bar{C}_L^N) + \left(\frac{1}{\gamma^N} \hat{C}_{L,t}^N + \hat{p}_t^N \right) \bar{C}_L^N &= \hat{w}_{L,t} - \hat{\tau}_{L,t}^{Payroll}\end{aligned}$$

Changes in uniform VAT are counteracted by corresponding negative changes in payroll taxes, in order to not distort the labour supply choice: $\hat{\tau}_t^{VAT} = -\hat{\tau}_{L,t}^{Payroll} = -\hat{\tau}_{H,t}^{Payroll}$. The budget

constraints of either Ricardian or hand-to-mouth agents are also not affected by the uniform VAT tax rate. In the baseline model, without steady state profits, this is achieved with a zero transfer rule, as steady state consumption equals steady state labour income, so that overall VAT taxes paid $((C_L^E + p^N C_L^N) \hat{\tau}_t^{VAT}$ for a hand-to-mouth agent, equivalently for a Ricardian agent) equals overall payroll taxes received $(-w_L N_L \hat{\tau}_{L,t}^{Payroll})$. In a more general case, lump-sum transfers would ensure no effect of unconventional taxation on the budget constraints.

No other choice of any private sector agent (households, unions, or firms) is affected by the uniform VAT tax, payroll tax, or lump sum transfers. This implies that the government can replicate the effects of monetary policy with unconventional fiscal policy also in this two-agents two-goods economy. ■

This proposition shows the equivalence between the two policy tools.¹⁰ The next step is devising a VAT rule to implement a shock that is equivalent to a monetary policy shock. We do this by specifying an AR(1) for the uniform VAT tax rate. We match the response of overall consumption to a monetary policy shock in the estimated model with the response of overall consumption to the VAT shock, by varying the persistence and the standard deviation of the VAT shock. The resulting coefficients are: $\rho^{VAT} = 0.947$ and $\sigma^{VAT} = 4.84$.

The objective of [Correia et al. \(2013\)](#) is to implement a fiscal policy plan that is equivalent to a monetary policy action but is not subject to the zero lower bound constraint. In this paper, we follow a similar strategy, although the objective is not to study the zero lower bound. Instead, we focus on goods-specific VAT changes. We want to know if the government is better off allocating one dollar to reduce the uniform VAT tax burden, or one dollar to reduce either the non-essential VAT rate or the essential VAT rate, for the purpose of generating a stronger stimulative effect on aggregate consumption, aggregate earnings and potentially the labour income of HtM workers. In order to achieve this, we implement AR(1) rules for the two good specific VAT rates, with the same persistence as in the uniform case ($\rho^{VAT} = \rho^{VATN} = \rho^{VATE}$) and with a standard deviation proportional to the share of non-essential or essentials ($\sigma^{VATN} = \frac{\sigma^{VAT}}{C^N}$ and $\sigma^{VATE} = \frac{\sigma^{VAT}}{1-C^N}$). This formulation makes sure that the changes in gross tax revenues are the same in the three cases. Note that, because of the compensating changes in the payroll taxes, the net change in tax revenue is always zero.

E.3 Experiment results

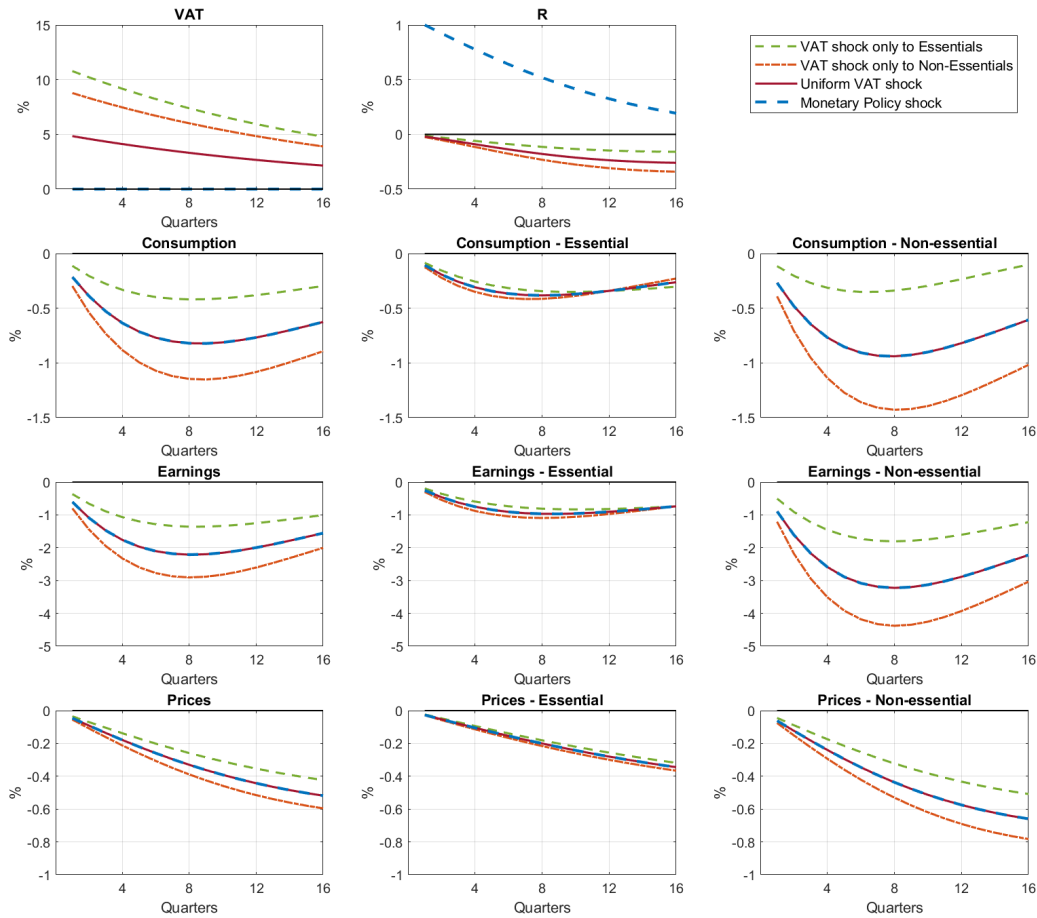
Figures 8 and E.1 show the results of these exercises. First, the uniform VAT shock exactly matches the effect of the monetary policy shocks on all non-policy variables (the red uniform lines overlap with the blue dashed line). Despite having targeted only aggregate consumption, the uniform VAT tax shock replicates the effect of monetary policy on sector specific variables. To replicate a one percentage point increase in interest rates, the VAT rate must increase by almost 5 percentage points, indicating that VAT changes can be equivalent to monetary policy, but at the cost of high VAT changes.

Levying VAT only on non-essentials produces much stronger effects than adjusting overall VAT rates or applying the VAT rate change only on essentials. Overall consumption declines

¹⁰We speculate that the same equivalence result with unconventional fiscal policy would hold in a general class of heterogeneous models with a full distribution of agents. We leave this for future research.

by 0.3 percentage points more at its peak when VAT is applied specifically to non-essentials. This larger effect is entirely due to a stronger reduction in non-essential consumption. There is almost no difference in the response for essential goods, as households shift their spending toward essentials (the orange dashed line should be above the red solid line), but the general equilibrium effect is much smaller (the orange dashed line should be below the red solid line). These opposing effects balance out in this model (indicated by the orange dashed line being roughly equal to the red solid line). A similar pattern is seen in labor earnings and prices.

Figure E.1: Alternative fiscal policies



Notes: Responses to alternative fiscal policies. The equivalent VAT shock is VAT applied uniformly to all goods, accompanied by an offsetting income tax, which replicates the response to monetary policy. The essential and non-essential VAT shocks are sector-specific VAT policies designed to raise the same revenue as the uniform VAT. The size of these policy shocks is shown in the top left subplot chart.