Non-Essential Business-Cycles

Michele Andreolli, Natalie Rickard, and Paolo Surico*

June 2024

Abstract

Using newly constructed time series of consumption, prices and earnings in essential and non-essential sectors, we document three main empirical regularities on post-WWII U.S. data: (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 in essential sectors; (iii) low-earners are more likely to work in non-essential industries. We develop and estimate a structural model with non-homothetic preferences over two expenditure goods, hand-to-mouth consumers and heterogeneity in labour productivity that is consistent with these findings. We use the model to revisit the transmission of monetary policy and find that the *interaction* of cyclical product demand composition *and* cyclical labour demand composition greatly amplifies business-cycle fluctuations.

Keywords: income elasticity, recessions, monetary policy, amplification, inequality.
JEL Classification Codes: E52, D31, E21.

*We are grateful to Martin Beraja, Florin Bilbiie, Timo Boppart, Andrea Galeotti, Peter Karadi, Joseba Martinez, Facundo Piguillem, Giorgio Primiceri, Gianluca Violante, 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, and North America Summer Meeting of Econometric Society for very helpful comments and suggestions. 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.

1 Introduction

A central question in macroeconomics is how business-cycle shocks propagate to the real economy. In the last decade, significant research efforts have convincingly advocated an important role for household heterogeneity and firm heterogeneity in the transmission of monetary and fiscal policies to aggregate consumption and income. Despite these important advances, existing contributions have focused almost exclusively on the demand side or on the supply side in isolation, and little is known about how household heterogeneity and firm heterogeneity may connect, and whether their interaction dampens or amplifies business-cycle fluctuations.

This paper proposes and evaluates a novel source of business-cycle amplification that cuts across household and firm heterogeneity: the interaction between the cyclical composition of product demand and the cyclical composition of labour demand. Our starting point is the Engel curve, namely the observation that the budget share of essential goods (or necessities) decreases with income, while the budget share of non-essentials (or luxuries) increases with household resources. We connect the Engel curve with two further pieces of evidence. First, industries that produce non-essentials employ a larger share of low-earning workers. Second, households at the bottom of the income distribution exhibit a higher Marginal Propensity to Consume (MPC).

In response to a contraction in household resources, driven for instance by a recession or by an increase in interest rates, consumers face an incentive to shift the composition of their spending away from non-necessities, which in turn triggers a fall in the labour demand of non-essential industries. As the labour force in these sectors is characterized by a larger share low-income/high MPC workers, the decline in non-essential earnings sets in motion a second round of consumption effects that amplify the initial contraction. Central to this mechanism is the cyclical composition of labour demand that results from the interaction between the cyclical composition of product demand and the unequal distribution of low-earning workers across industries.

We provide novel empirical evidence and a theoretical foundation for this mechanism, using disaggregated time-series and cross-sectional data on prices, consumption and earnings across essential and non-essential sectors in the United States. We begin with descriptive statistics about a typical post-WWII American recession and regional evidence across U.S. states. Then, we present conditional correlations based on an identified shock. We isolate exogenous variation in monetary policy using the high-frequency approach developed by Gertler and Karadi (2015) and extended by Jarociński and Karadi (2020), which exploit 'surprise' movements in federal funds future rates in the thirty minutes window around FOMC policy announcements. The impulse response functions are estimated using the smooth Local Projection (LP) method proposed by Barnichon and Brownlees (2019), which penalizes the higher variance typically associated with least-square LP estimators (see Li, Plagborg-Møller and Wolf, 2022).

Our empirical analysis uncovers several regularities. During recessions, household expenditure on non-essentials contracts significantly more than spending on essentials. In addition, earnings in industries that produce non-necessities is significantly more cyclical than earnings in essential sectors. And, non-essential industries employ a significantly larger share of low-income workers than sectors producing necessities. These correlations hold true across three empirical settings that differ markedly by the source of data variation: unconditionally, across post-WWII U.S. recessions; unconditionally, across U.S. states; conditional to identified monetary policy shocks. Furthermore, we show that, within non-essential industries, the earning responses to interest rate changes are significantly larger at the bottom of the income distribution. Finally, we find that the relative price effect across essentials and non-essentials is rather modest.

We develop and estimate a structural model that feature three main ingredients: (i) non-homothetic preferences, (ii) hand-to-mouth agents, (iii) heterogeneity in labour productivity. Households consume two types of goods that differ by their income elasticity, following Deaton (1974) and Deaton (1978). Browning and Crossley (2000) show that non-necessities are easier to postpone in the face of a negative shock, which implies that the intertemporal elasticity of substitution is heterogeneous across goods. Workers have either low productivity and are financially constrained or have high productivity and are unconstrained. Non-essential sectors employ a larger share of the former. We show that adding these features to an otherwise standard New-Keynesian model allows us not only to account for our empirical findings but also to offer a microfoundation for both the cyclicality of spending composition documented in this paper and the counter-cyclicality of income inequality emphasized by Bilbiie (2020) and Patterson (2023).

We use the estimated model to perform counterfactual analyses that highlight the drivers of our results. The *interaction* between the unequal spending composition and the unequal labour composition accounts for about half of the effects of monetary policy on aggregate consumption, and amplifies business-cycle fluctuations relative to both the representative agent case and models that feature only heterogeneity in either spending or labour composition. In a representative agent set-up, we show analytically that non-homotheticity leads to no amplification at all.

Another contribution of the paper is to show how to exploit the richness of detailed micro data on consumer prices, labour market and input/output tables from publicly available sources to construct quarterly time series for any categorization of household expenditure, price indices, earnings and employment. Our starting point is the method by Aguiar and Bils (2015), who classify consumer spending into essentials and non-essentials depending on whether their income elasticity is below or above one, respectively. Endowed with this classification from the CEX, we then construct category-specific time series for consumption and prices using the Personal Consumption Expenditure sub-indices. To classify industries into essentials and non-essentials, and thus construct employment and earning measures that reflect the consumption categorization of interest, we distinguish between final and intermediate goods. The final goods sectors are split in two groups depending on whether they produce essentials or non-essentials. Intermediate good industries are classified based on the final goods production they primarily contribute to, using the input-output tables. Finally, we exploit this industry classification to construct time series for earnings and employment in essentials and non-essentials, using the detailed information on labour force sectoral composition along the income distribution provided by the Current Population Survey (CPS).

Related literature. Our paper contributes to several strands of work in macroeconomics. A prominent literature has studied the impact of demand composition on business-cycle dynamics. Barsky, House and Kimball (2007) focus on the transmission of monetary policy when durable goods prices are more flexible than non-durable prices; McKay and Wieland (2019), and Beraja and Wolf (2021) show that the lumpy nature of durable expenditure can alter the transmission of monetary policy and the strength of recoveries; Jaimovich, Rebelo and Wong (2019) find that the shift in demand towards lower quality products amplified the employment drop during the Great Recession. The distinction between essentials and non-essentials differs from the demand composition in these works. First, we document a strong covariance between non-essential spending and non-essential earnings, but find a much lower comovement between consumption and income in either the durable or the non-tradeable sector. Second, non-essential industries witness a much higher concentration of low-income 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 shifts towards low quality products *within* a specific sector or spending category (e.g. from premium to grocery store brands).

Growing research efforts have focussed on the contribution of income inequality and income risk to the amplification of business-cycle fluctuations. Bilbiie (2020) and Patterson (2023) identify the crucial role played by the covariance between the marginal propensity to consume and earning cyclicality across workers, while McKay, Nakamura and Steinsson (2016), Ravn and Sterk (2017, 2021) and Bilbiie, Primiceri and Tambalotti (2023) highlight the contribution of counter-cyclical income risk. Cloyne, Ferreira and Surico (2020) show that the indirect effects of monetary policy on income across households are key to account for the aggregate consumption response, while Holm, Paul and Tischbirek (2021), Amberg et al. (2022), Andersen et al. (2023) document significant heterogeneity in the earning responses along the income distribution.[1] Relative to these studies, we emphasize a novel dimension of heterogeneity: essentials versus non-essentials. By documenting and modelling the interaction between the cyclical composition of product demand and the cyclical composition of labour demand, we show that the uneven distribution of low-skilled and high-skilled workers across industries provides yet another powerful amplification mechanism, which so far has been overlooked.

Finally, earlier contributions have used non-homothetic preferences to analyse salient features of either: (i) consumption and saving behaviour, including the intertemporal elasticity of substitution (Browning and Crossley, 2000), wealth accumulation (De Nardi and Fella, 2017), price rigidities (Clayton, Jaravel and Schaab, 2018), costs of living (Orchard, 2022) and marginal propensity to consume (Andreolli and Surico, 2021); or (ii) structural transformation, such as Kaldor's facts (Foellmi and Zweimüller, 2008, Boppart, 2014) and labor market polarization (Comin, Danieli and Mestieri, 2020).[2] We depart from these important works along two main dimensions. First, we develop and estimate a structural model in which the cyclicality of non-essential spending and the uneven sectoral distribution of low-income workers make labour demand in non-essential sectors highly pro-cyclical. Second, we use our non-homothetic preferences model to quantify the contribution of the interaction between cyclical product demand and cyclical labour demand to business-cycle fluctuations and the transmission of monetary policy on aggregate consumption.[3]

[1] Another important literature emphasizes the role of heterogeneity in the marginal propensity to consume. For instance, Kaplan, Moll and Violante (2018), Auclert (2019) and Debortoli and Galí (2017) separate the direct effects from the indirect effects of monetary policy on consumption, while Auclert, Rognlie and Straub (2020) and Bilbiie, Känzig and Surico (2022) highlight the role of capital investment.

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

[3] Other applications of non-homothetic preferences to business cycle analyses include Olivi, Sterk and Xhani (2023), Sonnervig (2022) and Danieli (2020). Relative to these studies, we provide a broader range of empirical evidence (exploiting both time-series variation across U.S. recessions and geographical variation across U.S. states), new estimates of a novel business cycle model featuring both good market and labour market heterogeneity, counterfactual analyses based on the estimated structural model that not only quantify the contribution of each channel to business cycle fluctuations but also highlight a key complementarity between the unequal spending composition across goods and the unequal workers composition across sectors to amplify the effects of monetary policy on aggregate consumption and income.

Structure of the paper In Section 2, we present our measurement strategy spanning across multiple granular datasets, and provide descriptive evidence about the newly constructed time series for essentials and non-essentials. 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. Finally, 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 altering the transmission of monetary policy and amplifying business-cycle fluctuations.

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.

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.1.[4]

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.[5] Finally, we construct Fisher indices for consumption and prices following the approach outlined in the BEA NIPA (2021) handbook, Chapter 4.

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

[4]In Appendix A.1, 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.

[5]As in Aguiar and Bils (2015), we either adjust or omit from our essential/non-essential classification, spending categories that are likely to be 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 Appendix A.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

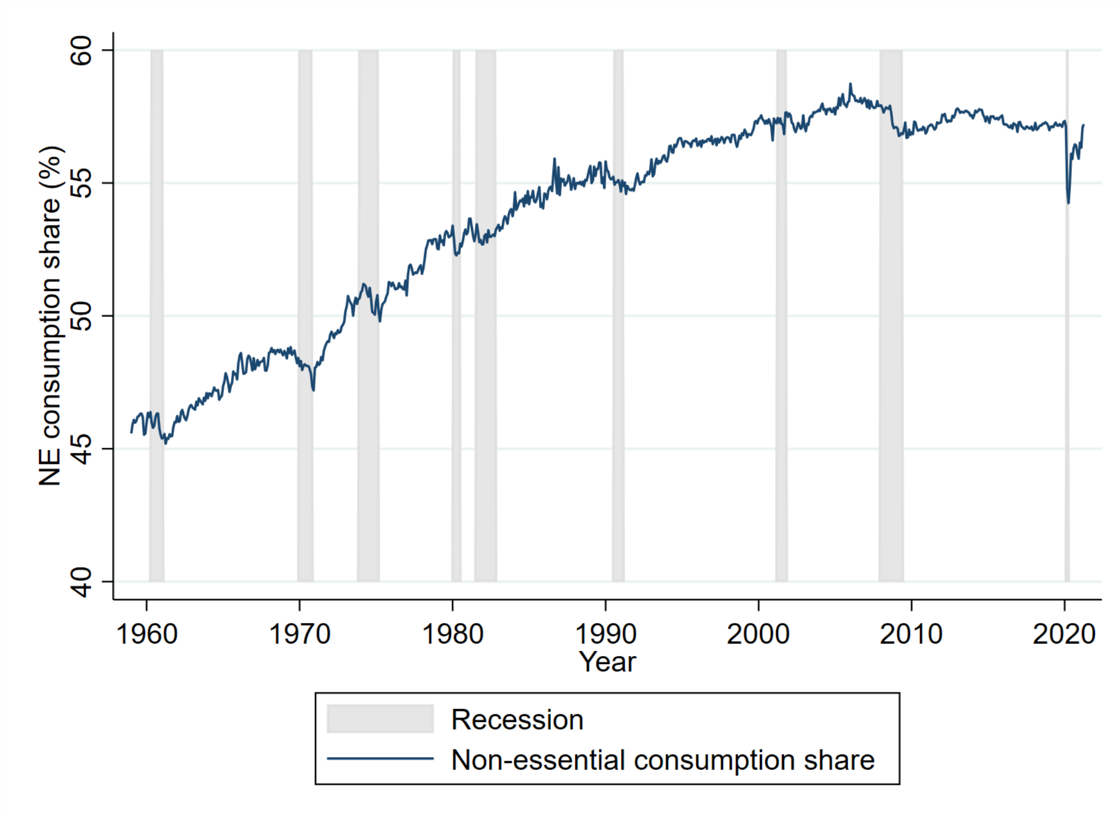

Notes: Personal consumption expenditure shares of non-essentials, constructed from chained (2000$) spending series, and as a proportion of total classified expenditure. Underlying data are from the BEA PCE by Type of Product tables. Shaded areas in grey represent NBER recession dates.

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 cyclicality of final demand affects labour demand in essential and non-essential sectors. We begin by manually classify 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.[6] We classify an intermediate industry as essential (non-essential) if the downstream final consumption it mostly contributes to is essential (non-essential).[7]

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.[8] 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 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

[6]Details on the industry classification and the mapping of intermediate goods industry into final expenditure essential and non-essential sectors are outlined in Appendix A.3.

[7]An 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.

[8]In 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 piece of evidence. First, in Appendix Figure C.4, 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 Appendix Figure C.14 we have verified that our results are not overturned if we aggregate overall earnings in the CPS at quarterly frequency, instead.

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 A.6, 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.

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 is for the whole economy.

Three main take-away 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 from 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 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.

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 deciles of the earning distribution. This decays monotonically, moving from a value shy of 75% in the bottom decile to a number below 55% in the top decile. The right panel of Figure 2 plots the 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,

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

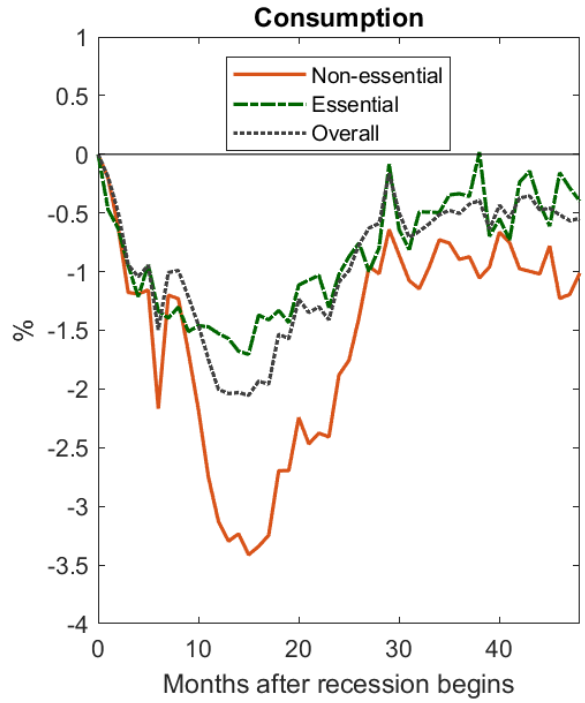

Notes: Response of essential and non-essential series after the start of a recession. To construct these responses, we first log and detrend the series using the HP filter ($\lambda = 14,440$). For the earnings series, we report a 6-month moving average to reduce noise. We then calculate the average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

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

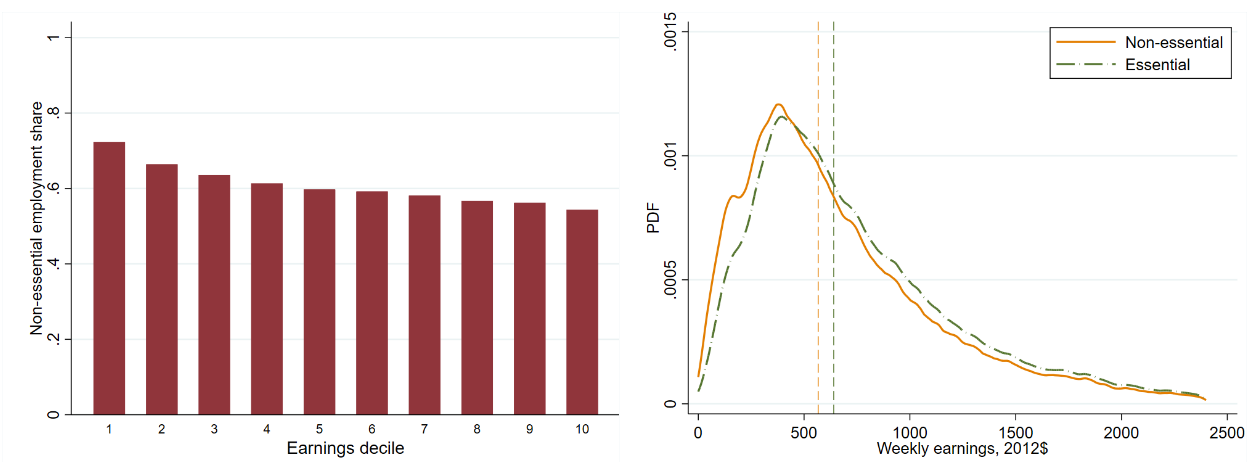

(a) Non-essential shares

(b) Kernel

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 with median earnings recording a 12% gap relative to their essential counterparts.[9] Putting these pieces of evidence together suggests that: (i) low-income workers are more likely to work in non-essential sectors (Panel (a) of Figure 3), and (ii) non-essential workers tend to be paid less than their essential counterparts (Panel (b)).

Figure 4: State-level employment during recessions vs non-essential consumption shares

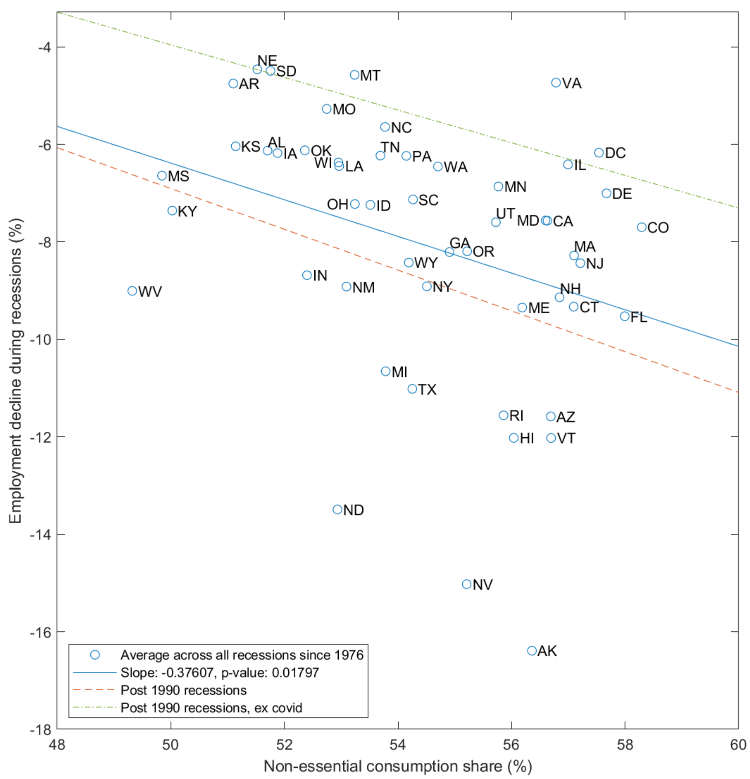

Notes: Seasonally adjusted monthly state level employment series is from the BLS. Employment declines are calculated around state-specific timing of national recessions (see main text), and averaged across all recessions since 1976. Non-essential shares are consumption shares constructed from the BEA's 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.

State-level analysis. To complement the results above, which are based on variation over time, 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 state of the United States, since 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

[9]In Appendix Figure C.5, we show that the CDF of essential earnings first order stochastically dominates the CDF of non-essential earnings. While Figure 3 pools earnings in each sector over the entire sample period, we have verified that the findings in both charts apply also to each individual year over time.

state-level downturns.[10] The analysis in Figure 4 reveals that, on average across all recessions in our sample, U.S. states with a higher share of non-essential consumption have experienced a larger decline in employment. This finding 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 relatively stable over time.

3 Empirical Framework

In the previous section, we have used reduced-form correlations to document a number of novel 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; low-income workers are more likely to be employed in non-essential sectors. 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 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 our identification strategy and our 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, Sack and Swanson (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

[10] In Appendix A.8, we provide details on the data construction for the state-level analysis.

High-Frequency Identification (HFI) of monetary policy shocks of Gertler and Karadi (2015), who in turn build on Gurkaynak, Sack and Swanson (2005). This measures changes in Fed Funds futures over a short window of time, typically 30 minutes, around monetary policy announcements. These provide 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.[11]

Econometric method. The high-frequency identified monetary policy instrument is available for period 1990 to 2016. However, as pointed out by Cloyne et al. (2018), an extended monetary shock series can be produced by estimating a proxy-VAR in the tradition Mertens and Ravn (2013) and Stock and Watson (2018) over a longer sample, and then identifying the monetary policy surprise series using the HFI monetary policy instrument over the shorter period over which is available (as in Gertler and Karadi, 2015). In practice, we extract the monetary shocks estimating a proxy-VAR similar to that 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, in the language of Ramey (2011) and Plagborg-Møller and Wolf (2021). This specification is robust to the non-invertibility of the VAR. We use 12 lags for the endogenous variables and 4 lags of the instrument. The extracted monetary policy shocks are reported in Appendix Figure B.1.

To check for weak instruments in our specification, we run the weak instruments test proposed by 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 and Karadi (2015), the corresponding robust F-statistic is 13.89, passing the weak instrument test. When using the shocks identified à la Jarociński and Karadi (2020), 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

[11] A 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 (see for instance Jarociński and Karadi, 2020, Miranda-Agrippino and Ricco, 2021, Nakamura and Steinsson, 2018). In one of the robustness exercises at this section end, we obtain very similar results when the analysis is based instead on the refined monetary policy surprises constructed by Jarociński and Karadi (2020).

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):

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

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_hs are the object of our interest, as they summarize the impulse responses of the ys 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 discussed in Montiel Olea and Plagborg-Møller (2021) 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 B.

3.2 Results across spending categories

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 5. 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 5. 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 top row is about two times as large and persistent as the decline in essential spending. 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, 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, 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 5. 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.[12] The sectoral price responses, however,

[12]The 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 Appendix C, we also report the IRFs of the other macro aggregate series included in the proxy-SVAR.

Figure 5: IRFs to contractionary 100bp monetary policy shock - Consumption and Earnings

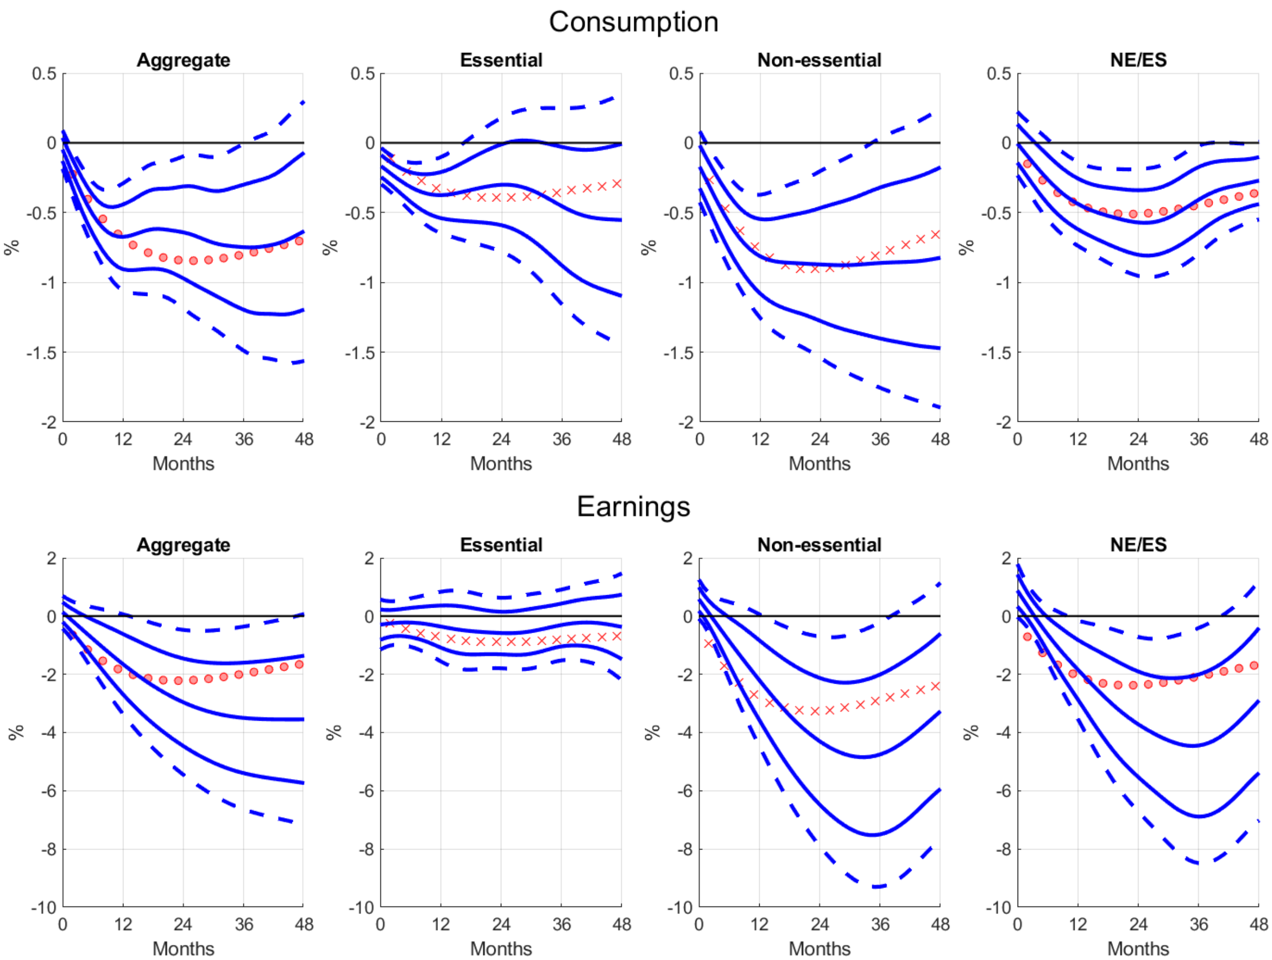

Notes: Blue lines are empirical impulse response functions (IRFs) to a 100bp increase in the 1y year government bond yields estimated by smooth local projections instrumental variable, where the instrument is the monetary policy shocks derived from the Gertler and Karadi (2015) high-frequency monetary 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 B.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 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) low-income workers are more likely to be employed in non-essential sectors. Finally, a long-standing empirical literature on survey and administrative data find that low-income households exhibit a high Marginal Propensity to Consume (MPC), (see for instance Johnson, Parker and Souleles, 2006, Patterson, 2023, among many others). 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 high-MPC) 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 high-MPC 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 6, 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 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 6 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, Perri and Violante (2010). The higher sensitivity of salaries

These are estimated using SLP-IV and are similar to those in the literature (e.g. Gertler and Karadi, 2015).

Figure 6: IRFs to contractionary monetary policy shock - Earnings distribution

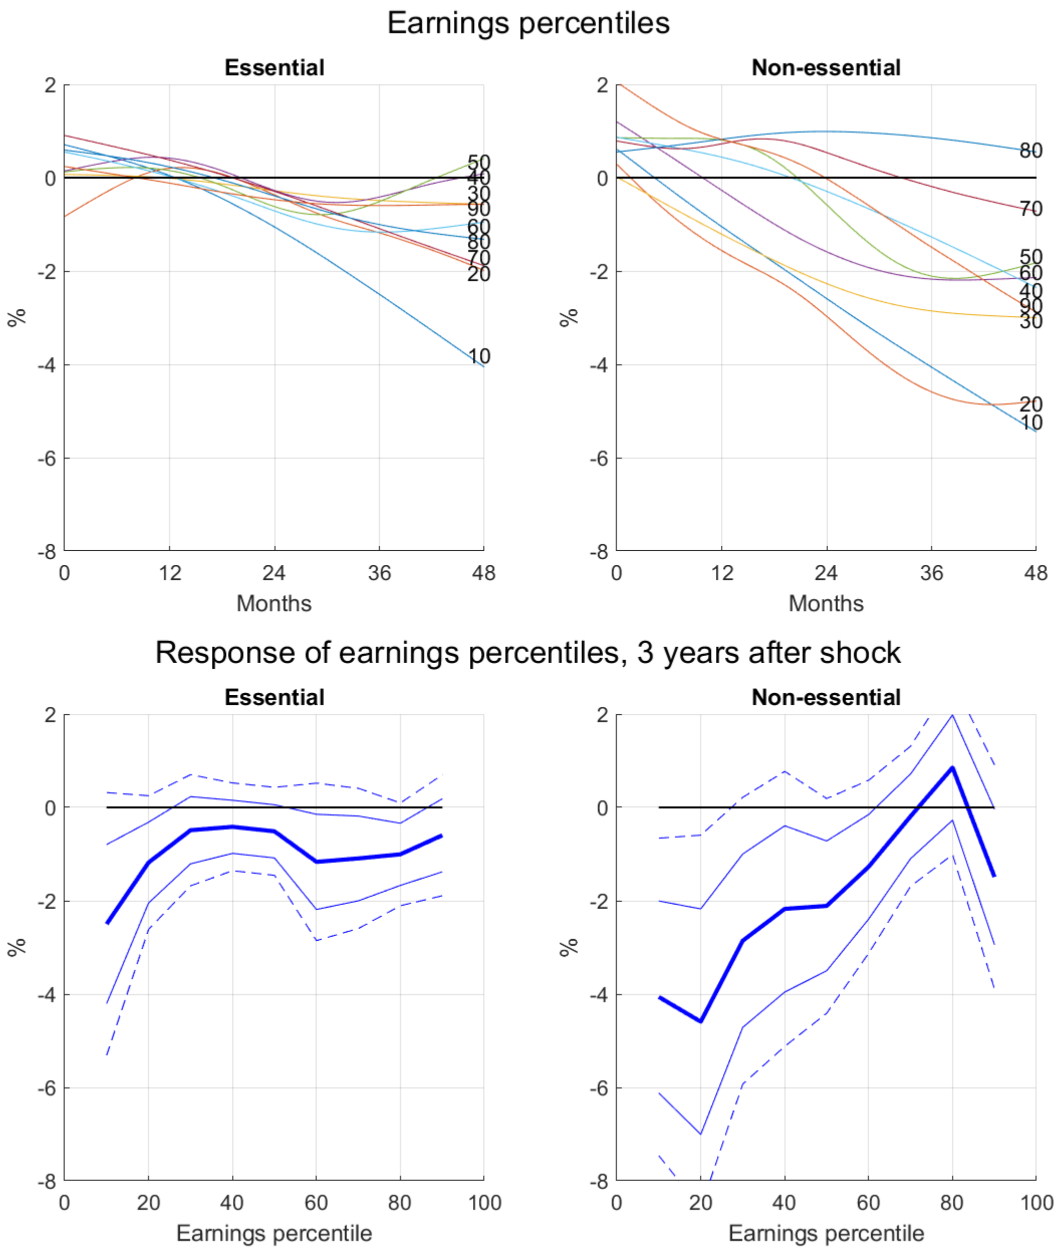

Notes: Empirical impulse response functions (IRFs) to a 100bp increase in the 1y year government bond yields estimated by smooth local projections instrumental variable, where the instrument is the monetary policy shocks derived from the Gertler and Karadi (2015) high-frequency monetary policy surprises. Sample periods and controls for each column are specified in the main text and Appendix B.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.

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.[13]

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 A.6 and C.

Alternative expenditure classifications. A possible challenge for 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 for the cyclicality across essentials and non-essentials that we have documented above. In Appendix A.6, 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 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 cyclicality 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. Finally, the positive correlation between income level and elasticity of demand that is crucial for the labour market amplification discussed in this paper is weak for durables: non-essential workers are typically paid less than essential workers whereas wages in durable industries are not necessarily lower than in

[13]The 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.

non-durable sectors. 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.

Information effect. Our baseline specification employs an (updated) series of shocks from Gertler and Karadi (2015). As outlined by Nakamura and Steinsson (2018), Jarociński and Karadi (2020) and Miranda-Agrippino and Ricco (2021), however, these shocks could be contaminated by the 'information effect', namely the notion that monetary policy announcements may also reflect private information held by the central bank about the state of the economy. Earlier in this section, we have reported that while the standard HFI monetary policy surprise series of Gertler and Karadi (2015) passes the weak instrument test, the refined series that also clean for the information effect does not. Still, in Appendix C.5, we report the impulse responses of all the main aggregate variables and our newly constructed sectoral series on consumption, prices and earnings to monetary policy shocks that remove the information effect, following the strategy proposed by Jarociński and Karadi (2020). We find that the results remain similar to those in Figure 5, suggesting that the information effect does not have any significant influence on our findings.

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 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.6, 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, Collard and Dellas (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.6 are qualitatively very similar to the findings in Figure 5, despite the two sets of IRFs are based on very different identification strategies.[14]

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

Further sensitivity. In Appendix C.7, 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 sample, however, may simply 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 Appendix C.8, we show that using unsmoothed local projection instrumental variables produce qualitatively similar results, though the point estimates are more jaggered 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 kep the latter as our baseline specification.

Finally, in Appendix C.9, 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 Appendix C.9, are qualitatively similar to the estimates in Figure 5, 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.

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 stuctural 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 face a financial constraint). Third, non-essential industries employ a higher share of low productivity 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 families, whose workers are disproportionately employed in non-essential industries. As low-income households have a higher MPC, 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 potetially different elasticites 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):

$$U(C^E_{i,t}, C^N_{i,t}, N_{i,t}) = \frac{\left(C^E_{i,t}\right)^{1-\frac{1}{\gamma^E}}}{1-\frac{1}{\gamma^E}} + \varphi \frac{\left(C^N_{i,t}\right)^{1-\frac{1}{\gamma^N}}}{1-\frac{1}{\gamma^N}} - \xi \frac{N^{1+\chi}_{i,t}}{1+\chi} \tag{2}$$

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 γs distinguishes essentials from non-essentials:

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

Intuitively, households are very unwilling to delay consumption of necessities such as groceries, and prefer instead to smooth their consumption. In contrast, households are more willing to delay spending on large durables or on hospitality and food away from home: 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. Browning and Crossley (2000) prove that this holds for any additively separable utility function in good varieties.[15]

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 varies over time and is higher for wealthier households, consistent with the evidence in Crossley and Low (2011).[16] Moreover, these preferences are a simple extension of the standard CRRA 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.[17]

4.2 Households problem

Households i have an instantaneous utility for essentials, $C^E_{i,t}$, and for non-essentials, $C^N_{i,t}$, 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).

[15]Note that the non-homothetic CES preferences of Comin, Lashkari and Mestieri (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.

[16]Stiglitz (1969) show how non-homothetic preferences are linked to risk aversion while Ait-Sahalia, Parker and Yogo (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 an their consumption-specific risk aversion.

[17]Other stands 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.

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, $C^E_{i,t,0}$, and for the future if they do not update for j periods ahead, $C^E_{i,t+j,j}$. 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, Parker and Souleles, 2006). Accordingly, we assume that H agents have higher income and are Ricardian, whereas L agents have lower income and are hand to mouth.[18] High-earning agents are paid 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 workers. As these shares turn out to be an important dimension of heterogeneity, both in the model and in the data, in Section 5 we estimate these parameters. 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:

$$y^i_{k,t} = Y^i_t \left(\frac{P^i_{k,t}}{P^i_t}\right)^{-\varepsilon} \quad i = \{E, N\}$$

[18] This 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.

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 production 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:

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

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 Appendix D, 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, using standard values in the literature. 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, Parker and Souleles (2006). The inverse of the macro Frisch elasticity is 0.1, consistent with the evidence reported by Christiano, Trabandt and Walentin (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, ($\bar{C}^E_H$), and low-skilled families, ($\bar{C}^E_L$), to 0.44 and 0.60 respectively, using the expenditure share data in Appendix Table A.2. 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 Appendix D.7. The resulting values are: $\varphi = 1.4225$ and $a^E = 1.4749$. Finally, we follow Bilbiie (2008) and Debortoli and Galí (2017) by allocating firms' profits to Ricardian agents.[19]

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	s_{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	0.050	0.197	0.111
IES difference for non-essentials	$\gamma^N - \gamma^E$	Normal	1.000	1.000	0.686	0.183
Low skilled share in essentials	α^E	Beta	0.100	0.004	0.019	0.078
Low skilled share difference in non-essentials	$\alpha^N - \alpha^E$	Beta	0.100	0.004	0.310	0.084
Inattentiveness	λ	Normal	0.050	1.000	0.013	0.028
Interest rate smoothing	ρ_R	Beta	0.900	0.040	0.952	0.007
Price stickiness	θ	Beta	0.900	1.000	0.958	0.010

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.

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),

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

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 impulse to construct the likelihood function, following Guerron-Quintana, Inoue and Kilian (2017).[20] Key moments of the prior distributions are summarized in Panel B of Table 1. The prior means for essentials and non-essentials IESs are chosen so as to imply an aggregate IES of 0.86, which represents the middle point between the point estimate of Smets and Wouters (2007) and the log-utility case. The priors on the shares of low-skilled workers are diffuse and centered around 10% for essential industries and 20% for non-essentials, consistent with the evidence from the CPS. The prior on the inattentiveness parameter is relatively uninformative and its mean corresponds to the point estimate in Beraja and Wolf (2021). Interest rate smoothing and price stickiness coefficients display prior distributions in line with the available evidence (see for instance Smets and Wouters, 2007, Justiniano and Primiceri, 2008).

Estimation results. In Table 1 Panel B, we report the estimated parameters of the structural model. The IES for Essential goods is $\gamma^E = 0.20$ and the IES for Non-essential goods is $\gamma^N = 0.88$: the difference between the two IESs is economically and statistically very significant, implying an economy-wide IES of 0.54. These estimates correspond to an average income elasticity around 0.48 for essentials and about 1.39 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 IESs across spending categories draw 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, Banks and Tanner (2002) and Calvet et al. (2021) which, based on fewer categories, suggest that low-income families have a smaller intertemporal elasticity of substitution relative to high-income households. This latter finding is consistent with both non-homothetic preferences and the IES heterogeneity across essentials and non-essentials that we uncover.

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

[20]We detail the estimation procedure in Appendix E.1.

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 5 reveal that the IRFs of consumption and earning implied by the estimates of our structural model track well the corresponding IRFs estimated with SLP-IVs, both for essentials and non-essentials. In Appendix Figure E.1, we report the full set of results, including also the effects on essential prices, non-essential prices, and on the interest rate. Spending and earnings on non-essentials decline by more than for essentials. The estimated structural model is also able to reproduce the dampening of consumption 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 spending. 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 anticipate or postpone, and therefore their demand responds relatively more to income changes. The second ingredient is labour market heterogeneity: low-income workers are more likely to be employed in non-essential sectors and thus they face a labour demand that is relatively more sensitive to the business-cycle. Finally, low-income households have a high MPC (i.e. hand-to-mouth) and hence the relatively stronger decline in their labour earnings during recessions strongly feed back into lower aggregate demand, setting in motion a second round of spending and earnings effects, across sectors and along the income distribution, that further exacerbate the 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 (between hand-to-mouth and savers), unequal spending composition (between essentials and non-essentials), and unequal labour market sectoral composition (between low-skilled and 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) marginal propensity to consume, 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 refers to 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).[21]

[21] To 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 $\mu_L > 0$

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 conjuction with another source of heterogeneity, to 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,[22] both with and without non-homothetic preferences.[23] Second, adding hand-to-mouth consumers in the second column brings the shares of the explained consumption response to 35% and 38%, respectively with and without equal spending composition.[24] 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.[25] 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 47%.[26]

---

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.

[22] Our 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 further discussion). The amplification in the rest of the table results from indirect, general equilibrium effects. Kaplan, Moll and Violante (2018) show that direct effects account for almost the entirety of the transmission of monetary policy in representative agent models. In addition, 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.

[23] In the next section, we show analytically that non-homothetic preferences do not lead to any amplification at all in a representative agent model, thereby generalizing the result in the rows of Table 2 first column.

[24] The 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.

[25] We 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 hand-to-mouth 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).

[26] In 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

Table 2: Counterfactual exercise: amplification

	Representative Agent	Heterogeneous Agents	
		Equal Labour Composition	Unequal Labour Composition
Equal spending composition	0.22	0.35	0.47
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 0.47 to 1 (when the two channels are *jointly* considered) are much larger than the increases from 0.35 to, respectively, 0.38 and 0.47 (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 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 both cyclical product demand composition and cyclical labour demand composition to further amplify the effects of monetary policy. To illustrate this triple interaction, in Figure 7, 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, Parker and Souleles, 2006).[27]

In each simulation, the cumulated consumption response to monetary policy is normalized by the cumulated effect in the representative agent/good case. This implies that each point of Figure 7 can be interpreted as the extent of amplification of that model (and for that value of μ^L) relative to the representative benchmark. The blue line refers to the full structural

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. $\alpha^N \neq \alpha^E$).

[27] To ensure that the economic significance of hand-to-mouth agents reflects their relative size, for any value of μ^L, we adjust the labour income shares accrued to hand-to-mouth households in Figure 7 such that $\alpha^J = \bar{\alpha}^J \frac{\mu^L}{\bar{\mu}^L}$, where $\bar{\mu}^L, \bar{\alpha}^J$ are the values taken by these parameters in the estimated full structural model.

Figure 7: On the Sources of Amplification

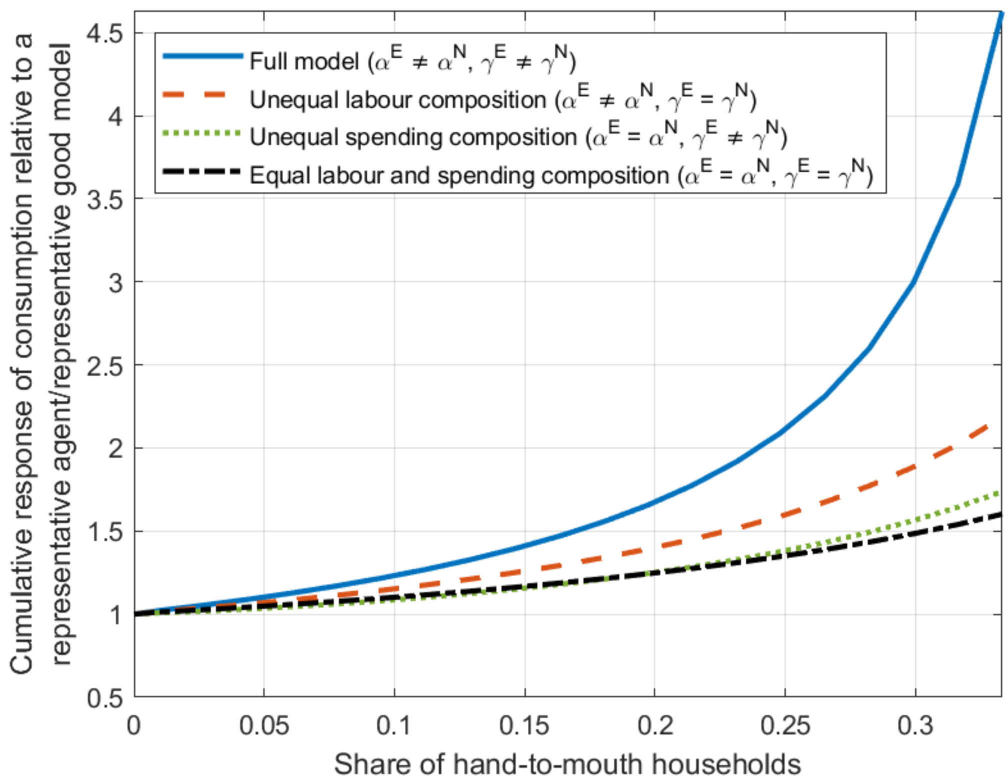

Notes: Amplification is measured by the cumulative IRF of consumption of each model, divided by the cumulative response of consumption in the restricted model with no hand-to-mouth agents. The figure depicts four scenarios: (i) the unrestricted full model as blue solid line, (ii) unequal labour sectoral composition (i.e. $\gamma^E = \gamma^N$) as orange dashed line, (iii) unequal spending composition (i.e. $\alpha^E = \alpha^N$) as green dotted line, and (iv) equal labour and equal spending composition (i.e. $\alpha^E = \alpha^N$ and $\gamma^E = \gamma^N$) as black broken line. The latter is often referred to in the literature as Two-Agents New-Keynesian (TANK) model. As in Table 2, whenever $\alpha^E = \alpha^N = \tilde{\alpha}$, we set $\tilde{\alpha}$ so as to match the relative steady state labour earnings across the two agents. Whenever $\gamma^E = \gamma^N$, we set the IES to equal the average IES in the estimated full structural model.

model that features both cyclical product demand composition and cyclical labour demand composition, whereas the black broken line summarizes the results of the restricted model with neither of the two. The dashed orange line and the dotted green line stand for the two intermediate cases of only unequal labour composition or only unequal spending composition, respectively.

Four main results emerge from this exercise. First, in all models, a higher share of hand-to-mouth consumers leads to a monotonic increase in the extent of amplification, though the nonlinearity of this relationship is very heterogeneous across models. Second, 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 in earlier studies on U.S. data such as Patterson (2023) and Bilbiie, Primiceri and Tambalotti (2023). Third, non-homothetic preferences seem to add little amplification over TANK, whereas the marginal contribution of the unequal labour sectoral composition appears relatively larger. Fourth, the extent of amplification in the full model (depicted as blue line) is consistently larger than the sum of the dashed orange line and the green dotted line over the whole range of values for μ^L. 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 distrinction between essentials and non-essentials has no impact on aggregate fluctuations in representative agent settings.

In Appendix F Proposition 1, 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:

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

$$\text{Homothetic} \rightarrow \frac{\partial \hat{C}_t}{\partial \varepsilon_t^{mp}} = -IES$$

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 Appendix Table E.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 Conclusions

What drives business-cycle fluctuations? Demand or supply? Consumer spending or workers' earnings? The top or the bottom of the income distribution? And what are the mechanisms through which shocks propagate to the rest of the economy? These questions have fascinated macroeconomists for centuries; yet, they are still hotly debated. In this paper, we take a fresh look at these important issues by uncovering a novel transmission mechanism that cut across demand and supply, household expenditure and wage payments, high-income and low-income consumers.

The main idea is based on the observation that households and workers differ greatly in their exposure to the business-cycle along the income distribution, and that the composition of product and labour demand, in particular the divide between essentials and non-essentials, is crucial to identify and quantify their cyclical exposure to shocks. The intuition is that the consumption of the rich drives the income of the poor. In the face of economic adversities, non-essential purchases are easier to postpone and their contraction is dominated by the behaviour of affluent households. The latter has a particularly large effect on low-income families, whose workers are more likely to be employed in non-essential industries, and thus their labour demand suffers more from the drop in product demand: the higher cyclicality of non-essential expenditure leads to a higher cyclicality of non-essential earnings. Taken together, the households with less resources in society lose twice because of the spending 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, we show that: (i) high-income consumers spend more on non-essentials; (ii) low-earning 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 effect 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 cyclicality 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 the unequal incidence of recessions in the goods markets (i.e. non-essential spending contracts by more), the unequal incidence of recessions in the labour markets (i.e. non-essential earnings fall by more), and the uneven labour force composition across sectors (i.e. low-income workers are more likely employed in non-essential industries) accounts for about half of the effects of monetary policy on aggregate consumption.

Our findings have potentially interesting ramifications for normative and policy analyses. Inflation and output dynamics in essential and non-essential sectors are likely to vary significantly across a number of dimensions —such as consumption smoothing, price elasticity and price stickiness— that are crucial for the design of optimal policies. In addition, the sectoral heterogeneity we have documented may prove relevant for the design of fiscal interventions: industrial policies and labour market policies targeted to non-essential sectors are more likely to short-circuit the amplification mechanism documented in this paper, as a higher share of low-income workers is employed in those sectors. We leave these fascinating topics for future research. But, we hope that our analysis might stimulate and assist national statistical offices and central banks in levering the available granular data on households, workers, sectors and input-output accounts to construct nationally representative series for the prices, consumption and earnings of essential and non-essential industries.

References

Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. "The network origins of aggregate fluctuations." *Econometrica*, 80(5): 1977–2016.

Aguiar, Mark, and Mark Bils. 2015. "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review*, 105(9): 2725–56.

Ait-Sahalia, Yacine, Jonathan A. Parker, and Motohiro Yogo. 2004. "Luxury Goods and the Equity Premium." *The Journal of Finance*, 59(6): 2959–3004.

Amberg, Niklas, Thomas Jansson, Mathias Klein, and Anna Rogantini Picco. 2022. "Five Facts about the Distributional Income Effects of Monetary Policy Shocks." *American Economic Review: Insights*, 4(3): 289–304.

Andersen, Asger Lau, Niels Johannesen, Mis JØrgesen, and Jose-Luis Peydro. 2023. "Monetary Policy and Inequality." *The Journal of Finance*, 78(5): 2945–2989.

Andreolli, Michele, and Paolo Surico. 2021. "LESS Is MORE: Consumer Spending and the Size of Economic Stimulus Payments." CEPR Discussion Paper.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2020. "Business-cycle anatomy." *American Economic Review*, 110(10): 3030–3070.

Attanasio, Orazio P, James Banks, and Sarah Tanner. 2002. "Asset holding and consumption volatility." *Journal of political economy*, 110(4): 771–792.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review*, 109(6): 2333–2367.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2020. "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model." Download paperRevise and resubmit at American Economic Review.

Barnichon, Regis, and Christian Brownlees. 2019. "Impulse Response Estimation by Smooth Local Projections." *The Review of Economics and Statistics*, 101(3): 522–530.

Barsky, Robert B., Christopher L. House, and Miles S. Kimball. 2007. "Sticky-Price Models and Durable Goods." *The American Economic Review*, 97(3): 984–998.

Beraja, Martin, and Christian K Wolf. 2021. "Demand Composition and the Strength of Recoveries." National Bureau of Economic Research Working Paper 29304.

Bilbiie, Florin O. 2008. "Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic." *Journal of economic theory*, 140(1): 162–196.

Bilbiie, Florin O. 2020. "The New Keynesian Cross." *Journal of Monetary Economics*, 114: 90–108.

Bilbiie, Florin O., Diego R. Känzig, and Paolo Surico. 2022. "Capital and income inequality: An aggregate-demand complementarity." *Journal of Monetary Economics*, 126(C): 154–169.

Bilbiie, Florin O., Giorgio Primiceri, and Andrea Tambalotti. 2023. "Inequality and Business Cycles." National Bureau of Economic Research, Inc NBER Working Papers 31729.

Blundell, Richard, Monica Costa Dias, Robert Joyce, and Xiaowei Xu. 2020. "COVID-19 and Inequalities*." *Fiscal Studies*, 41(2): 291–319.

Boppart, Timo. 2014. "Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences." *Econometrica*, 82(6): 2167–2196.

Browning, Martin, and Thomas F. Crossley. 2000. "Luxuries Are Easier to Postpone: A Proof." *Journal of political Economy*, 108(5): 1022–1026.

Calvet, Laurent E, John Y Campbell, Francisco Gomes, and Paolo Sodini. 2021. "The cross-section of household preferences." National Bureau of Economic Research.

Christiano, Lawrence J, Mathias Trabandt, and Karl Walentin. 2010. "DSGE models for monetary policy analysis." In *Handbook of monetary economics*. Vol. 3, 285–367. Elsevier.

Clayton, Christopher, Xavier Jaravel, and Andreas Schaab. 2018. "Heterogeneous Price Rigidities and Monetary Policy." *SSRN Electronic Journal*.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary Policy When Households Have Debt: New Evidence on the Transmission Mechanism." *The Review of Economic Studies*, 87(1): 102–129.

Cloyne, James, Clodomiro Ferreira, Maren Froemel, and Paolo Surico. 2018. "Monetary Policy, Corporate Finance and Investment." National Bureau of Economic Research.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2017. "Innocent Bystanders? Monetary Policy and Inequality." *Journal of Monetary Economics*, 88: 70–89.

Comin, Diego A, Ana Danieli, and Martí Mestieri. 2020. "Income-driven Labor Market Polarization." National Bureau of Economic Research Working Paper 27455.

Comin, Diego, Danial Lashkari, and Martí Mestieri. 2021. "Structural Change with Long-Run Income and Price Effects." *Econometrica*, 89(1): 311–374.

Crossley, Thomas F., and Hamish W. Low. 2011. "Is the Elasticity of Intertemporal Substituion Constant?" *Journal of the European Economic Association*, 9(1): 87–105.

Danieli, Ana. 2020. "Expenditure Elasticity and Inequality Over the Business Cycle." *Available at SSRN 3593499*.

Deaton, Angus. 1974. "A Reconsideration of the Empirical Implications of Additive Preferences." *The Economic Journal*, 84(334): 338–348.

Deaton, Angus. 1978. "Specification and Testing in Applied Demand Analysis." *The Economic Journal*, 88(351): 524–536.

Debortoli, Davide, and Jordi Galí. 2017. "Monetary Policy with Heterogeneous Agents: Insights from TANK Models."

De Nardi, Mariacristina, and Giulio Fella. 2017. "Saving and Wealth Inequality." *Review of Economic Dynamics*, 26: 280–300.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren. 2020. "Integrated Public Use Microdata Series, Current Population Survey: Version 7.0."

Foellmi, Reto, and Josef Zweimüller. 2008. "Structural Change, Engel's Consumption Cycles and Kaldor's Facts of Economic Growth." *Journal of monetary Economics*, 55(7): 1317–1328.

Gertler, Mark, and Peter Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics*, 7(1): 44–76.

Guerron-Quintana, Pablo, Atsushi Inoue, and Lutz Kilian. 2017. "Impulse response matching estimators for DSGE models." *Journal of Econometrics*, 196(1): 144–155.

Gurkaynak, Refet S, Brian Sack, and Eric 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(1): 39.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante. 2010. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006." *Review of Economic Dynamics*, 13(1): 15–51.

Holm, Martin Blomhoff, Pascal Paul, and Andreas Tischbirek. 2021. "The Transmission of Monetary Policy under the Microscope." *Journal of Political Economy*, 129(10): 2861–2904.

Jaimovich, Nir, Sergio Rebelo, and Arlene Wong. 2019. "Trading down and the business cycle." *Journal of Monetary Economics*, 102(C): 96–121.

Jaimovich, Nir, Sergio Rebelo, Arlene Wong, and Miao Ben Zhang. 2020. "Trading Up and the Skill Premium." *NBER Macroeconomics Annual*, 34(1): 285–316.

Jarociński, Marek, and Peter Karadi. 2020. "Deconstructing Monetary Policy Surprises—The Role of Information Shocks." *American Economic Journal: Macroeconomics*, 12(2): 1–43.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review*, 96(5): 1589–1610.

Johnson, Noah N. 2017. "Tradable and nontradable inflation indexes: replicating New Zealand's tradable indexes with BLS CPI data." *Monthly Labor Review*, 1–25.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review*, 95(1): 161–182.

Justiniano, Alejandro, and Giorgio E. Primiceri. 2008. "The Time-Varying Volatility of Macroeconomic Fluctuations." *American Economic Review*, 98(3): 604–641.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review*, 108(3): 697–743.

Li, Dake, Mikkel Plagborg-Møller, and Christian K. Wolf. 2022. "Local Projections vs. VARs: Lessons From Thousands of DGPs." National Bureau of Economic Research, Inc NBER Working Papers 30207.

Mankiw, N Gregory, and Ricardo Reis. 2007. "Sticky information in general equilibrium." *Journal of the European Economic Association*, 5(2-3): 603–613.

McKay, Alisdair, and Johannes F. Wieland. 2019. "Lumpy Durable Consumption Demand and the Limited Ammunition of Monetary Policy." National Bureau of Economic Research.

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

Mertens, Karel, and Morten O Ravn. 2011. "Understanding the aggregate effects of anticipated and unanticipated tax policy shocks." *Review of Economic dynamics*, 14(1): 27–54.

Mertens, Karel, and Morten O. Ravn. 2013. "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States." *American Economic Review*, 103(4): 1212–1247.

Miranda-Agrippino, Silvia, and Giovanni Ricco. 2021. "The Transmission of Monetary Policy Shocks." *American Economic Journal: Macroeconomics*, 13(3): 74–107.

Montiel Olea, José Luis, and Mikkel Plagborg-Møller. 2021. "Local Projection Inference Is Simpler and More Robust than You Think." *Econometrica*, 89(4): 1789–1823.

Nakamura, Emi, and Jón Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *The Quarterly Journal of Economics*, 133(3): 1283–1330.

Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703–708.

NIPA. 2021. "NIPA Handbook: Concepts and Methods of the U.S. National Income and Product Accounts."

Olea, José Luis Montiel, and Carolin Pflueger. 2013. "A Robust Test for Weak Instruments." *Journal of Business & Economic Statistics*, 31(3): 358–369.

Olivi, Alan, Vincent Sterk, and Dajana Xhani. 2023. "Optimal Monetary Policy during a Cost of Living Crisis." mimeo, UCL. 1.

Orchard, Jacob. 2022. "Cyclical Demand Shifts and Cost of Living Inequality." *Available at SSRN 4033572*.

Patterson, Christina. 2023. "The Matching Multiplier and the Amplification of Recessions." *American Economic Review*, 113(4): 982–1012.

Plagborg-Møller, Mikkel, and Christian K. Wolf. 2021. "Local Projections and VARs Estimate the Same Impulse Responses." *Econometrica*, 89(2): 955–980.

Ramey, Valerie A. 2011. "Identifying Government Spending Shocks: It's All in the Timing." *The Quarterly Journal of Economics*, 126(1): 1–50.

Ravn, Morten, and Vincent Sterk. 2017. "Job uncertainty and deep recessions." *Journal of Monetary Economics*, 90(C): 125–141.

Ravn, Morten O, and Vincent Sterk. 2021. "Macroeconomic Fluctuations with HANK & SAM: an Analytical Approach." *Journal of the European Economic Association*, 19(2): 1162–1202.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and frictions in US business cycles: A Bayesian DSGE approach." *American economic review*, 97(3): 586–606.

Sonnervig, Marcos. 2022. "Unequal Business Cycles."

Stiglitz, Joseph E. 1969. "Behavior towards risk with many commodities." *Econometrica: Journal of the Econometric Society*, 660–667.

Stock, James H., and Mark W. Watson. 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *The Economic Journal*, 128(610): 917–948.

Stock, James H., and Mark W. Watson. 2020. "Slack and Cyclically Sensitive Inflation." *Journal of Money, Credit and Banking*, 52(S2): 393–428.

Taylor, John B. 1993. "Discretion versus policy rules in practice." *Carnegie-Rochester Conference Series on Public Policy*, 39(1): 195–214.

Verona, Fabio, and Maik H Wolters. 2014. "Sticky information models in Dynare." *Computational Economics*, 43: 357–370.

Wachter, Jessica A., and Motohiro Yogo. 2010. "Why Do Household Portfolio Shares Rise in Wealth?" *The Review of Financial Studies*, 23(11): 3929–3965.

Online Appendix

A Measurement

This Appendix first outlines the approach for constructing new data series for essentials and non-essentials. As a first step, we classify goods and services into non-essentials and essentials, using CEX data. With this classification, we construct novel consumption and price series using PCE data. Next, using our final goods and services split across essentials and non-essentials and the input output tables, we classify industries according to the final goods and services that they ultimately sell to downstream. This industry classification is used to construct labour market series with CPS data. We also present additional macroeconomic data sources, summary statistics, alternative consumption categorisations, and the state-level analysis.

A.1 Classification procedure for consumption categories

This section outlines how we classify consumption categories into non-essentials and essentials, to build time series for consumption, prices, and labour market variables. Our first step uses consumption micro-data from the Consumer Expenditure Survey (CEX) to define what types of consumption goods are non-essential vs essential.

We classify consumption categories into into essentials and non-essentials by estimating Engel Curves closely following the approach used by Aguiar and Bils (2015). Aguiar and Bils (2015) use household microdata from three waves of the CEX and estimate the expenditure elasticities as β_j from:

$$\ln x_{hjt} - \ln \bar{x}_{hjt} = \alpha_{jt} + \beta_j \ln X_{ht} + \Gamma_j \mathbf{Z}_h + u_{hjt} \tag{4}$$

Where x_{hjt} is the expenditure by household h on goods of type j in year t, $\bar{x}_{hjt}$ is the equivalent average across households, X_{ht} is total household expenditure, instrumented by household income (dummies for category and log real after-tax income), α_{jt} are good fixed effects and $\mathbf{Z}_h$ are household characteristics (age range, earners and household size). For full details see the original paper, we replicate the identical empirical specification.

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), omitting the final two columns[28]. Essentials are defined as categories with an income/total

[28] Aguiar and Bils (2015) use two specifications, using either income to instrument total expenditure, or

expenditure elasticity of demand (β_j) less than one; non-essentials are defined as those with an elasticity greater than one.

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 (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.

We make two minor alterations to Aguiar and Bils (2015)'s approach. Firstly, we al-

lagged total expenditure to instrument current total expenditure. We use the former here. The reason is to hedge against the possible concern that the lagged spending instrument might bias downward the estimated elasticity for lumpy expenditure sectors, such a new cars. Whenever a household buys a car, they have higher total expenditure in that quarter, but the predicted expenditure from the instrument of the last quarter is lower, therefore associating a higher car expenditure with a lower total predicted expenditure, which biases the IED estimate towards a necessity. This attenuation in the elasticity estimate is not present with the income instrument, and in practice makes a substantial difference to the estimated IED for new cars.

ter slightly the set of product categories, introducing some narrower categories where the broader categories included goods that varied considerably in their elasticities. Specifically, we split “Applicances, phones, computers with associated services” into “Communications” and “Household appliances”, “All other transportation” into “Gas and vehicle maintenance” and “Public transport”, “Housing” into ”Rents” and “Owner-occupied housing consumption” and “Vehicle purchasing, leasing and insurace” into “New car purchases”, “Used car purchases” and “Other car spending (leasing, financing and insurance”. We also omit tobacco from the product categories, as the intertemporal substitutability of tobacco is likely more related to the addictive nature of the good than the income elasticity, so less related to our theoretical framework. Secondly, we estimate the Engel curves for 1995-1997 rather than 1994-1996, in order to use the more consistent goods categories reported in the CEX Interview FMLI files during these years. As Aguiar and Bils (2015) note, the expenditure elasticities do not vary considerably over time, and consistent with this using the slightly different sample period makes minimal difference to their original estimated elasticities.

Table A.2 shows the expenditure shares of non-essentials vs essentials by housing tenure type and by income group using the elasticities above, on the same CEX sample. For mortgagors the non-essential share is 63.9%, for owner-occupiers without a mortgage this is 60.6% and for renters it is 33.6%. Households in the lowest income tercile have a non-essential share of 44.3% and households in the top two income terciles have a non-essential share of 60.3%. We use this information to calibrate the structural model, as detailed in Table 1. Note that the expenditure shares here differ slightly from consumption shares reported from PCE data. There are two main reasons for this; i) expenditure shares here reflect nominal expenditure shares, rather than real consumption shares constructed from chained consumption series and ii) because of the differences in the underlying data.

A.2 Construction of Consumption and Price Indices

In this subsection, we show how we construct time series for consumption and price indices.

Using the estimated elasticities and classification into essential and non-essentials from the previous section, we match their counterparts in the *PCE by Type of Product* tables from the U.S. Bureau of Economic Analysis (BEA). The consumption categories included in the above do not cover the entire consumption bundle of households, but our approach is to maximise the coverage as much as possible. This mapping closely follows a similar exercise in Aguiar and Bils (2015). These omissions and adjustments largely follow Aguiar and Bils (2015) and include cases where either:

Table A.2: Non-essential expenditure shares: by tenure type and across income distribution

	Non-essential share	
By housing tenure type		
Mortgagor	63.9%	
Owner occupier (without mortgage)	60.6%	
Renter	33.6%	
	Non-essential share	
By income tercile		
First	44.3%	
Second	56.1%	Top 2/3: 60.3%
Third	63.3%	
	Non-essential share	
By income quintile		
First	43.1%	
Second	48.7%	
Third	55.8%	
Fourth	59.8%	
Fifth	64.9%	

Notes: Non-essential expenditure shares from CEX data (see text). Income terciles and quintiles are based on after tax income.

1. Expenditures not made entirely by private, US households for their own personal consumption; if they are made on behalf of households by non-profits, employers or insurers.
   - Includes: food on farms, food supplied to military, net expenditures abroad, expenditures relating to net foreign travel, final consumption expenditures of nonprofit institutions serving households, some categories of insurance.
2. The expenditure might reasonably not be considered consumption which generates personal utility, and is instead a form of saving or cost of saving or other expense.
   - Includes: financial services (bank/pension fund fees, investment service commissions), some categories of insurance.
3. We don't trust or unable to estimate reasonable Engel curve estimates using the CEX microdata, due to incomplete or inaccurate consumption reporting.
   - Includes: professional and other services (legal, accounting, union, professional associations, funerals), Foundations and grantmaking and giving services to households.

4. We classify children's clothing as essential and adults clothing as non-essential, using CEX data. In the PCE, there are three clothing categories; 'Women's and girls' clothing', 'Men and boys; clothing', and 'Children's and inflant's clothing'. We follow Aguiar and Bils (2015) in splitting the former two categories, attributing 22% to children's, essential clothing, and 78% to adults, non-essential clothing.

5. For health expenditures, we also follow Aguiar and Bils (2015) in only including the proportion of health expenditure made out of pocket by households, by adjusting down the health expenditure and net health insurance expenditures using National Health Expenditure Data from Centers for Medicare and Medicaid Services. This helps reduce the proportion of health expenditure which is contributed to by (for instance) government programmes and so not discretionary spending by households directly, but still included in PCE.

Following this process, we classify on average over the sample period 36% of expenditure reported in the PCE as essential, 44% as non-essential and the remaining 20% is left unclassified.

We then construct Fisher price and consumption quantity indices for essentials and non-essentials by aggregating the (nominal) expenditure and price subindices following the approach outlined NIPA (2021), Chapter 4. The quantity index aggregated from all the subindices i categorised as essentials (E) is given by:

$$Q_{t,E}^{F} = \sqrt{\frac{\sum_{i \in E} p_{i,t-1} q_{i,t}}{\sum_{i \in E} p_{i,t-1} q_{i,t-1}} \times \frac{\sum_{i \in E} p_{i,t} q_{i,t}}{\sum_{i \in E} p_{i,t} q_{i,t-1}}}$$

Where the (deflated) values within the summations are calculated using the nominal expenditure $e_{i,t}$ and price indices $p_{i,t}$ as appropriate, for instance:

$$p_{i,t-1} q_{i,t} = p_{i,t-1} * \frac{p_{i,t} q_{i,t}}{p_{i,t}} = p_{i,t-1} * \frac{e_{i,t}}{p_{i,t}}$$

And similarly for different combinations of lagged quantities and prices.

We construct the Fisher price indices for essentials as:

$$P_{t,E}^{F} = \sqrt{\frac{\sum_{i \in E} p_{i,t} q_{i,t-1}}{\sum_{i \in E} p_{i,t-1} q_{i,t-1}} \times \frac{\sum_{i \in E} p_{i,t} q_{i,t}}{\sum_{i \in E} p_{i,t-1} q_{i,t}}}$$

And the equivalent formulas for non-essentials. When we refer to consumption shares with the PCE data, we use chained consumption series also following the NIPA guidelines.

A.3 Mapping of final goods classification to industries for labour market variables

The next step in the process is to use the classification of the consumption goods to understand which industries are producing non-essentials vs essentials. This industry classification will allow us to classify workers into the sectors they work for and understand the labour market implications of non-essential consumption dynamics. The first step to do this is to classify final goods producing industries according to the goods and services they supply. However, we would also like to classify intermediate industries, in order to also account for upstream labour market implications of final good demand. To achieve this second step, we use the input-output matrix from the BEA to understand the downstream final goods that intermediate industries contribute to. The final step we take is to use this classification with CPS data to build time series of labour earnings, employment, and wages of worker who mainly produce essentials and non-essential goods and services.

Final goods producer classification. The first stage of this process is the final goods classification. We map consumption categories to all NAICS 2007 industries included in the input-output tables of the BEA. We manually classify all industry codes as either essential, non-essential or unclassified, based on whether the industry produces final consumption goods which fit into our classified consumption categories.

We take an unconservative approach to this final goods industry classification, in order to maximise the amount of employment we are able to categorise. If there is an industry which is primarily producing intermediate goods, but related to one consumption category, we still classify it according to that consumption category. This is because our second step using the input output approach we use will reassign an industry's sales of input goods to different sectors according to their eventual downstream use. For example, 'Photographic and Photocopying Equipment Manufacturing' (NAICS code 333316) would be an non-essential if purchased by households, but supplies a lot of intermediate inputs which are used in essential industries, so this is eventually classified as a essential industry. Sometimes we classify industries that produce a range of goods to the consumption category which they are *most* rather than entirely associated with. For instance, employees working for department stores may supply both essential and non-essential consumption goods, but we assume that the majority of goods supplied are within the non-essential consumption categories, and so classify these as non-essential.

Input-Output approach to classify intermediate industries. We would like to be classify industries which primarily produce intermediate goods based on the downstream final goods that they primarily supply. In order to do this, we use the input-output tables combined with the final good industries from the previous section. We take the *Use of commodities by industry* table from the BEA Input-Output Accounts Data for 2007 at the most detailed disaggregation of 405 industries. From there we exclude government, private households, secondary smelting and alloying of aluminum, scrap, used and secondhand goods, noncomparable imports, and rest of the world adjustment. This allows to have a square matrix of input-output linkages with 391 industries both as suppliers and buyers of intermediate inputs. We link each intermediate industry to the final products with the Leontief inverse, in order to assign each industry the essential or non-essential final products. For categories that we do not have downstream sales data, we use the final product classification from the CEX.

A simple production network model in the spirit of Acemoglu et al. (2012) can help to explain all the steps. We take an economy with N industries comprising intermediate and final products. Each industry i has total sales $X_i = p_i x_i$ which can be made to intermediate producers $p_i x_{i,j}$, consumers for personal consumption expenditures $C_i = p_i c_i$ or other agents for final good expenditures $Z_i = p_i z_i$ (these can be government, investment, inventories, or exports). Total quantity sold is:

$$x_i = \sum_{j=1}^{N} x_{i,j} + c_i + z_i$$

The production function of industry j uses intermediate inputs $x_{i,j}$ and other inputs l_j in order to produce x_j with a Cobb-Douglas production function:

$$x_j = A_j l_j^{\alpha_j} \prod_{i=1}^{N} x_{i,j}^{(1-\alpha_j)\omega_{i,j}}$$

The first order condition under perfect competition for each intermediate input is: $p_i x_{i,j} = p_j x_j (1-\alpha_j)\omega_{i,j}$. This allows a recursive structure on the industry sales by substituting it in:

$$X_i = \sum_{j=1}^{N} (1-\alpha_j)\omega_{i,j} X_j + C_i + Z_i$$

Which we can write in matrix form and invert it to find the Leontief inverse L. Notice that we use $\circ$ for the Hadamard product (the element-wise product).

$$
\begin{aligned}
X &= (((1-\alpha)\mathbf{1}'_N) \circ \Omega)X + C + Z \\
X &= (I_N - ((1-\alpha)\mathbf{1}'_N) \circ \Omega)^{-1}(C+Z) \\
X &= L(C+Z)
\end{aligned}
$$

We have a classification of final products as essential E, non-essential N, or unclassified U we can build three $N \times 1$ indicator vectors taking value one if the final product is of that category and zero otherwise: $\mathbb{1}_k$ for $k = \{E, N, U\}$. We can assign an industry to essential if this industry sells more to essential final goods than non-essential final goods and if the sum of these sales is higher than the sales to unclassified sectors. Mathematically, we assign industry i to essentials if:

$$
\begin{aligned}
\{L(C \circ \mathbb{1}_E)\}_i &> \{L(C \circ \mathbb{1}_N)\}_i \\
\{L(C \circ \mathbb{1}_E)\}_i + \{L(C \circ \mathbb{1}_N)\}_i &> \{L(C \circ \mathbb{1}_U)\}_i
\end{aligned}
$$

And similarly for non-essentials. We leave as unclassified each remaining industry. Intuitively this method allows to match intermediate industries to their most important final goods. As an example, we match *Grain farming* to essentials, and *Iron, gold, silver, and other metal ore mining* to non-essential, despite not being classified within final goods (as they are intermediates).

Given the intermediate input-output matrix cleaned with the steps above, $((1-\alpha)\mathbf{1}'_N) \circ \Omega$ is the IO matrix with each intermediate input sales $p_i x_{i,j}$ divided by the *Total industry output (basic value)* line: $p_j x_j$. The C vector we use to weight each sales to assign to the three categories is *Personal consumption expenditures* in the input-output data.

The outcome of this exercise is the classification in essentials and non-essentials of the intermediate and final industries, defined with NAICS 2007 codes.

Mapping between industry codes. Our objective is to create time series for labour market variables split by essentials and non-essentials, e.g. what are the labour earnings of workers who predominantly produce non-essentials. However, we must overcome one last intermediate step before merging the industry classification with the labour market data from the CPS: the datasets we use to classify industries and workers use different industry codes. To accommodate this, we have to map between two different industry codes; NAICS 2007 and census 1990. Table A.3 shows the steps we follow.

We primarily use the the cross-walk supplied by the Census Bureau for this. However,

Table A.3: Datasets and industry codes for labour market classification

CEX Table A.1	$\xrightarrow[\text{Classify industries}]{}$	Input-Output NAICS 2007	$\xrightarrow[\text{Adjust industry codes}]{}$	CPS Census 1990

Notes: This table shows the different dataset we use and the corresponding industry codes classification to classify labour market variables. Arrows show the direction of the mapping, from the initial final good classification to the final time series.

sometimes we use some discretion and make some assumptions to do map between the codes. First, we classify NAICS 2007 codes to categorise industries according to our essential/non-essential split from the CEX. Then we use these NAICs codes in our input-output adjustment process to classify intermediate industries. Once armed with intermediate industry classifications using the input-output approach, we then map the classification to census 1990 codes. This mapping between industry codes requires some approximations and adjustments:

1. Most importantly, for retail industry codes (census codes 580-691), many of the census codes are more disaggregated than the available NAICS codes. For those, we overwrite the intermediate industry classification from the input-output process, and instead we use the initial classification of the industry. This is because these industries primarily supply final goods which are more straightforward to classify directly than intermediate industries. We also directly classify private households as non-essential, as this is also a exclusively final goods industry.

2. A portion of NAICS codes have multiple NAICS codes in the industry data for one census code. An example of this is dairy product manufacturing (census code 101) which in the input output tables maps to four NAICS industry categories (Cheese manufacturing; Dry, condensed, and evaporated dairy product manufacturing; Fluid milk and butter manufacturing; Ice cream and frozen dessert manufacturing). For these cases, we apply the same classification for all NAIC codes that related to a particular census code, treat them as separate industries in the input-output table processing, and then average the final sales shares to different categories of industries (essential, non-essential and unclassified) across a census industry using the total sales of each NAICS industry as weights.

3. Some census codes are more detailed than the NAICS codes in the input-output tables. For example, there is a census code (402) for taxicab services, which corresponds to NAICS code 485300 but only the more aggregated NAICS code 485000 is available in the input-output tables. In these cases, we assign the sales shares of the more aggregated NAICS industry to the more disaggregated census industry. This assumes

that the disaggregated industry does not vary substantially in what it supplies goods to compared to the more aggregated industry.

4. Some census codes are only mapped to large NAICS categories in the crosswalk, often because they are non-specificed or miscellaneous industries. For example, the census code 472 (non-specified utilities) is part of NAICS code 22, although there are more direct mappings between the codes in NAICS 22 and the census codes. For those industries, we also take an weighted average of all sales shares of all relevant industries (here, for example, 221100, 221200 and 221300), again assuming that the average of the larger group will be representative of the industries in the census code. Where not possible, (in particular, for Manufacturing non-durable, allocated) we leave unclassified.

5. Finally, there a few remaining cases where the mapping is less straightforward, because industries are divided differently in the two industry classifications. For example, knitting mills (census code 132) corresponds to NAICS codes 31324 and 3151, but in the input-output tables only the larger categories 3132 and 315 are available. In the same spirit as the previous approaches, we select all NAICS codes at the more aggregated level that include relevant industries, and take a total sales-weighted average of the sales shares to essentials, non-essentials and apply this to the census industry. Again this assumes that the census industry's sales shares are represented reasonably by the more aggregated industry.

Full mappings between NAICS 2007 industries in the input-output tables and the 1990 census industry codes used are given in the replication files.

Final classification of industries into essentials and non-essentials. Using the classification from the Input-Output approach we classify all industries as either essential, non-essential or unclassified. The final industry classification is presented in Table A.4. This is the classification we use for labour market variables.

Table A.4: Industry classification

Essential
Coal mining; oil and gas extraction; meat products; dairy products; canned, frozen, and preserved fruits and vegetables; grain mill products; bakery products; sugar and confectionery products; misc. food preparations and kindred products; food industries, n.s; miscellaneous paper and pulp products; drugs; soaps and cosmetics; agricultural chemicals; industrial and miscellaneous chemicals; petroleum refining; miscellaneous petroleum and coal products; tires and inner tubes; farm machinery and equipment; construction and material handling machines; office and accounting machines; guided missiles, space vehicles, and parts; medical, dental, and optical instruments and supplies; photographic equipment and supplies; u.s. postal service; pipe lines, except natural gas; wired communications; telegraph and miscellaneous communications services; electric light and power; gas and steam supply systems; electric and gas, and other combinations; water supply and irrigation; sanitary services; utilities, n.s; professional and commercial equipment and supplies; drugs, chemicals, and allied products; groceries and related products; petroleum products; wholesale trade, n.s; grocery stores; dairy products stores; food stores, n.e.c; auto and home supply stores; gasoline service stations; drug stores; fuel dealers; retail florists; insurance; personnel supply services; automobile parking and carwashes; automotive repair and related services; beauty shops; barber shops; funeral service and crematories; miscellaneous personal services; offices and clinics of physicians; offices and clinics of dentists; offices and clinics of chiropractors; offices and clinics of optometrists; offices and clinics of health practitioners, n.e.c; hospitals; nursing and personal care facilities; health services, n.e.c; residential care facilities, without nursing; accounting, auditing, and bookkeeping services; management and public relations services
Non-essential
Metal mining; nonmetallic mining and quarrying, except fuels; all construction; beverage industries; knitting mills; dyeing and finishing textiles, except wool and knit goods; carpets and rugs; yarn, thread, and fabric mills; miscellaneous textile mill products; apparel and accessories, except knit; miscellaneous fabricated textile products; pulp, paper, and paperboard mills; paperboard containers and boxes; newspaper publishing and printing; printing, publishing, and allied industries, except newspapers; plastics, synthetics, and resins; paints, varnishes, and related products; other rubber products, and plastics footwear and belting; miscellaneous plastics products; leather tanning and finishing; footwear, except rubber and plastic; leather products, except footwear; logging; sawmills, planing mills, and millwork; wood buildings and mobile homes; miscellaneous wood products; furniture and fixtures; glass and glass products; cement, concrete, gypsum, and plaster products; structural clay products; pottery and related products; misc. nonmetallic mineral and stone products; blast furnaces, steelworks, rolling and finishing mills; iron and steel foundries; primary aluminum industries; other primary metal industries; cutlery, handtools, and general hardware; fabricated structural metal products; screw machine products; metal forgings and stampings; ordnance; miscellaneous fabricated metal products; metal industries, n.s; engines and turbines; metalworking machinery; computers and related equipment; machinery, except electrical, n.e.c; machinery, n.s; household appliances; radio, tv, and communication equipment; electrical machinery, equipment, and supplies, n.e.c; electrical machinery,

equipment, and supplies, n.s; motor vehicles and motor vehicle equipment; aircraft and parts; ship and boat building and repairing; railroad locomotives and equipment; cycles and miscellaneous transportation equipment; toys, amusement, and sporting goods; manufacturing industries, n.s; railroads; bus service and urban transit; taxicab service; warehousing and storage; water transportation; air transportation; services incidental to transportation; radio and television broadcasting and cable; motor vehicles and equipment; furniture and home furnishings; lumber and construction materials; metals and minerals, except petroleum; electrical goods; hardware, plumbing and heating supplies; machinery, equipment, and supplies; scrap and waste materials; miscellaneous wholesale, durable goods; paper and paper products; apparel, fabrics, and notions; farm-product raw materials; alcoholic beverages; farm supplies; miscellaneous wholesale, nondurable goods; lumber and building material retailing; hardware stores; retail nurseries and garden stores; mobile home dealers; department stores; variety stores; miscellaneous general merchandise stores; retail bakeries; motor vehicle dealers; miscellaneous vehicle dealers; apparel and accessory stores, except shoe; shoe stores; furniture and home furnishings stores; household appliance stores; radio, tv, and computer stores; music stores; eating and drinking places; liquor stores; sporting goods, bicycles, and hobby stores; book and stationery stores; jewelry stores; gift, novelty, and souvenir shops; sewing, needlework, and piece goods stores; catalog and mail order houses; vending machine operators; direct selling establishments; miscellaneous retail stores; retail trade, n.s; savings institutions, including credit unions; credit agencies, n.e.c; real estate, including real estate-insurance offices; advertising; services to dwellings and other buildings; computer and data processing services; detective and protective services; business services, n.e.c; automotive rental and leasing, without drivers; electrical repair shops; miscellaneous repair services; private households; hotels and motels; lodging places, except hotels and motels; laundry, cleaning, and garment services; shoe repair shops; dressmaking shops; theaters and motion pictures; bowling centers; miscellaneous entertainment and recreation services; elementary and secondary schools; colleges and universities; vocational schools; educational services, n.e.c; child day care services; family child care homes; museums, art galleries, and zoos; labor unions; religious organizations; membership organizations, n.e.c; engineering, architectural, and surveying services; miscellaneous professional and related services.

Unclassified

Tobacco manufactures; manufacturing, non-durable - allocated; scientific and controlling instruments; watches, clocks, and clockwork operated devices; miscellaneous manufacturing industries; trucking service; banking; security, commodity brokerage, and investment companies; legal services; libraries; job training and vocational rehabilitation services; social services, n.e.c; research, development, and testing services

Notes: Classification of 1990 Census industry codes into essential, non-essential and unclassified.

Classification of labour market variables. Our last step is to use microdata from the Current Population Survey (CPS) to construct labour market series across industries; e.g. employment in the sectors producing essential good and services or labour earnings in the sectors producing non-essential good and services.

We construct employment using the main sample, and weekly usual earnings from the CPS ORG sample. We omit all workers working in agriculture or for the government. We combine these two series to give overall labour earnings for each sector. We also calculate earnings distributions within each sector, based on weekly usual earnings, as described in the main text. Using this classification, over the sample period 62% of employment is classified as non-essential, 30% as essential and the remaining 8% is unclassified.

Rather than the binary classification of industries into essential and non-essential, an alternative approach would be to classify the *share* of an intermediate industry which supplies downstream to non-essentials. For instance if a worker is employed in an industry where 60% of downstream consumption is essential 30% is non-essential and the remainder unclassified, in our baseline series we classify this employee as one essential worker. In our shares series, the employee would be counted as 0.6 of a person in the essential total employment series and 0.3 of a person in the non-essential employment series. We verify that our baseline empirical results are robust to using this alternative approach (results available upon request).

A.4 Other macroeconomic data sources

In addition to the constructed non-essential and essential series for consumption, prices, employment and earnings, we also use additional aggregate macroeconomic time-series in our Proxy-SVAR and local projection estimation, the sources for which are detailed below.

In the Proxy SVAR:

- Industrial production (INDPRO), PCE price index (PCEPI) and end of month 1y Treasury yields (DGS1) - downloaded from St Louis Fed's FRED, specific variable names in brackets.
- Excess bond premium, from the Federal Reserve Board[29]
- Monetary policy surprise series - both taken from the replication files of Jarociński and Karadi (2020):
  - The Gertler and Karadi shocks we use are the FF4 surprises updated and provided by Jarocinski and Karadi, which go from 1990m2 to 2016m12. There is a missing value on 2001m9 which we fill as zero.
  - The Jarconski and Karadi shocks we use, mitigating the information effect, are the FF4 surprise if there is a negative correlation between the FF4 surprise and the SP500 surprise. These go from 1990m2 to 2016m12. There is a missing value on 2001m9 which we fill as zero.

[29] https://www.federalreserve.gov/econresdata/notes/feds-notes/2016/updating-the-recession-risk-and-the-excess-bond-premium-20161006.html

In the smooth local projection estimations, in addition to the Proxy SVAR, we add:

- Total employment - depending on the sample, this is aggregated from the CPS data described previously for employment and earnings IRFs, otherwise we use total private employment recorded by the Current Employment Statistics (Establishment Survey, CES), taken from FRED (variable name USPRIV).
- Overall earnings - to compare with our constructed earnings series, we use the BEA NIPA series Total Compensation of Employees (Received: Wage and Salary Disbursements)
- Per worker earnings - median earnings series constructed using CPS data described previously, for SLP-IV IRFs for earnings. Otherwise, to give a longer time-series, we use Average Weekly Earnings of Production and Nonsupervisory Employees, for Private employees from the CES, also taken from FRED (CES0500000030).
- For the price IRFs, we also use inflation expectations as an additional control. For these, we use University of Michigan Inflation expectations, also taken from FRED (MICH)

A.5 Descriptive statistics and additional charts

To complement Figure 2 in the main text, which shows the dynamics of non-essential and essential consumption and earnings after recessions, Figure A.1 shows the corresponding inflation and employment dynamics. Inflation in the non-essential decelerates more rapidly than in the essential sector, though this heterogeneity is more mild. Here, we focus on core inflation, to remove the more supply-driven dynamics of energy and food inflation[30]. Employment in the non-essential sector sharply contracts, to a peak of nearly 4% below trend in the second year of the recession, while essential earnings decline by only 1%.

Table A.5 gives descriptive statistics of the constructed essential and non-essential series described above and in the main text. Consumption, employment and median earnings of non-essentials are more volatile than that of essentials, and more positively covary with industrial production (used as a proxy for overall output). In contrast, prices of non-essentials are less volatile and more negatively correlated with output, a fact we ascribe to the volatility of food and energy prices and that their variation may primarily not be caused by demand

[30]For the rest of the paper where we analyse identified responses to exogenous monetary policy shocks, this is no longer necessary and we instead address the response of the complete price index.

Figure A.1: Response of Essentials and Non-essentials over the business cycle - Prices and Employment

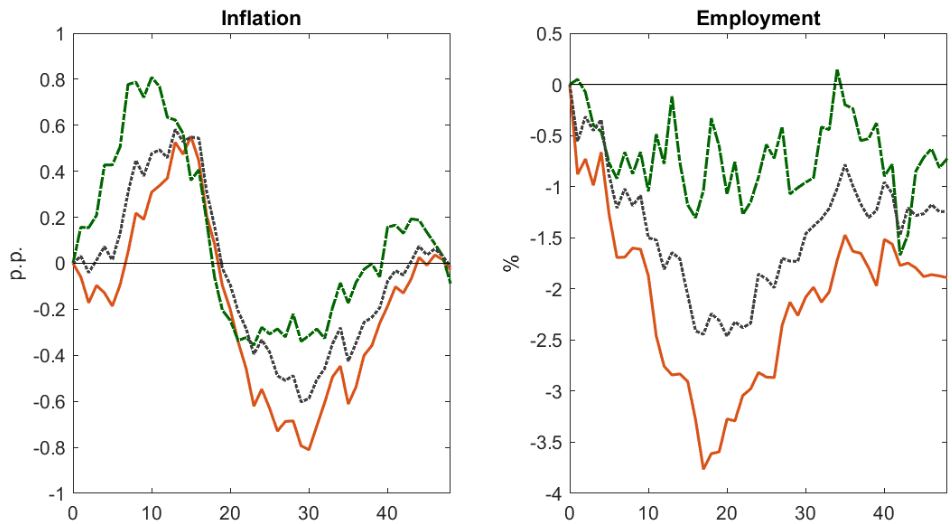

Response of essential and non-essential series, starting from the peak of the previous expansion, as defined by NBER. Includes all recession peaks since 1973 where non-essential and essential serie s for each variable are available for a full 48 months after the peak (peaks in 1973m11, 1981m7, 1990m7, 2001m3, 2007m12 and see sample definitions in text). For employment, this shows the cyclical component of the logged variable detrended using the HP filter ($\lambda = 14,440$). Inflation is y/y core inflation, also detrended using the HP filter. All series are normalised to 0 at the inital period by taking the peak observation from all periods.

shocks. Table A.6 shows average values of the series over the sample. Figure A.2 shows the underlying timeseries of Figure 2.

Table A.5: Descriptive statistics

	Consumption	Prices	Employment	Earnings
Correlation with Industrial Production				
Aggregate	0.68	-0.11	0.74	0.47
Essential	0.52	-0.02	0.36	0.17
Non-essential	0.73	-0.22	0.75	0.51
St. dev relative to Industrial Production				
Aggregate	0.15	0.10	0.16	0.24
Essential	0.14	0.19	0.13	0.22
Non-essential	0.21	0.09	0.21	0.32

Notes: Descriptive statistics for essentials and non-essentials series. All variables are year-on-year log differences. Panel 1 shows the standard deviation of the series, divided by the standard deviation for industrial production. Panel 2 shows the correlation between the series and industrial production. Monthly data used, the sample period ends in March 2020, and begins at the earliest available point for each series; January 1960 for consumption and prices, January 1977 for employment and January 1983 for earnings. Price and consumption are based on PCE data and employment and earnings are from CPS data, constructed as described in the text.

Table A.6: Average amount and share, Essentials vs Non-essentials

	Average annual amount			Non-essential % of overall
	Overall	Essential	Non-essential	
Consumption per cap. ($)	21,710	10,267	11,443	53%
Employment (mn)	93.4	30.6	62.9	67%
Median earnings	31,127	33,025	29,333	94%

Notes: The table shows the average annual amount of consumption, employment and median annual wages, in essentials and non-essentials, over the sample period. The final column shows the non-essential consumptiona and employment shares and the non-essential median wage as a % of overall median wages. Only the value of consumption and employment categorised into essentials and non-essentials is included in 'Overall', excluding uncategorised. Consumption is per capita chained PCE in 2012$, median wages are deflated to 2012$. Details of the calculations and data sources are included in text.

Figure A.2: Essentials and Non-essentials over time

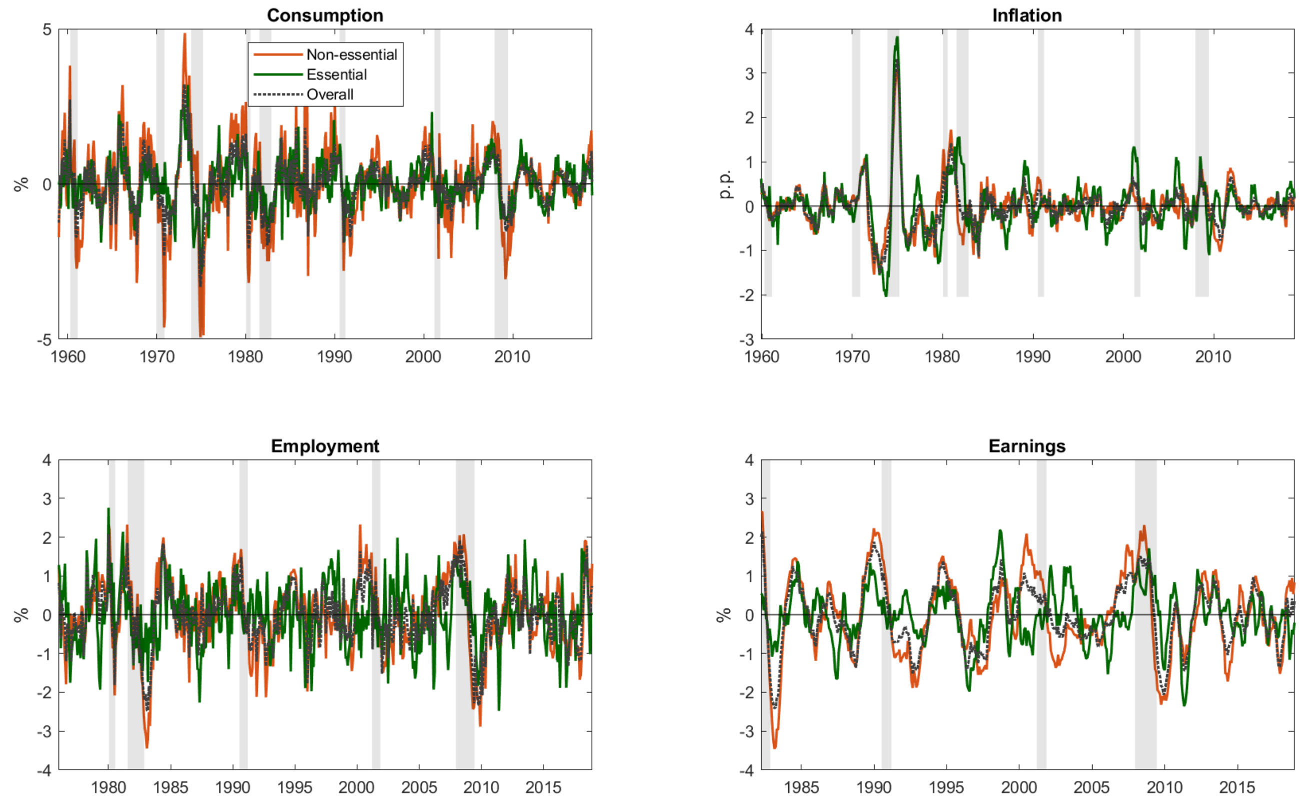

Underlying series of Figure 2. For consumption, employment and earnings, this shows the cyclical component of the logged variable detrended using the HP filter ($\lambda = 14,440$). For earnings, this refers to total earnings and the inital log series is centred 6-month rolling average, to reduce noise. Inflation is y/y core inflation, also detrended using the HP filter.

Core timeseries. Food and energy are essential categories, and may account for much of the variability in the essential price series, where the essential prices are (perhaps unexpectedly) more volatile than the non-essentials. We construct core essential and non-essential series, excluding the same categories as in the aggregate core series from the BEA. Comparison of the timeseries are in Figure A.3. These series are used in Figure A.1.

Figure A.3: Non-essential and essentials inflation - Headline vs Core

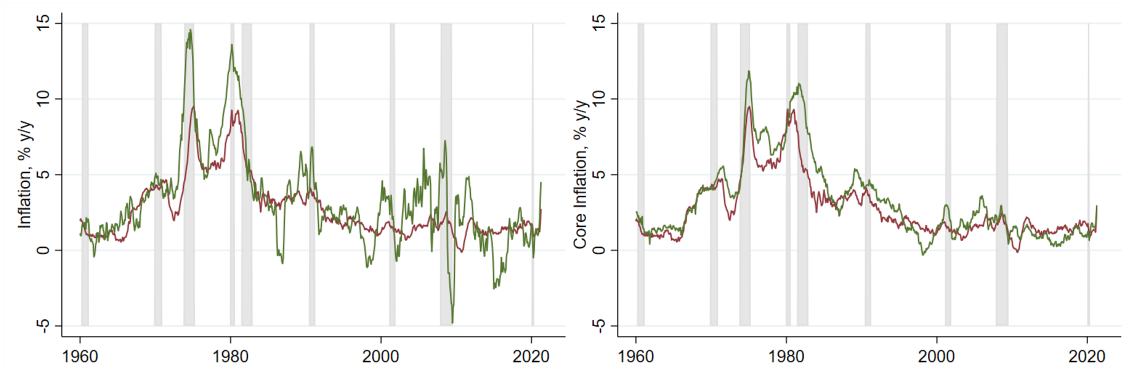

Notes: Non-essential and essential time-series inflation, LHS is headline, RHS is core (excluding food and energy). Underlying data sources are the PCE by Type of Product tables from the BEA, described in detail in the text. NBER recession dates shaded.

A.6 Durables versus non-durables

One major alternative characteristic that has been extensively discussed in the literature is the durablity of goods (Barsky, House and Kimball (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.

We discuss how these are constructed in detail below, but we follow the same broad approach and use the same data sources as for the overall non-essential and essential series.

Table A.7 shows 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 A.7: 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 A.4 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 cyclicality appears to be contributing relatively more to overall cyclicality. Figure A.5 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 cyclicality of durables. Both facts suggest that non-essentials' importance for business cycles is distinct from that of durables.

Figure A.4: Response of Durables and Non-durables over the business cycle

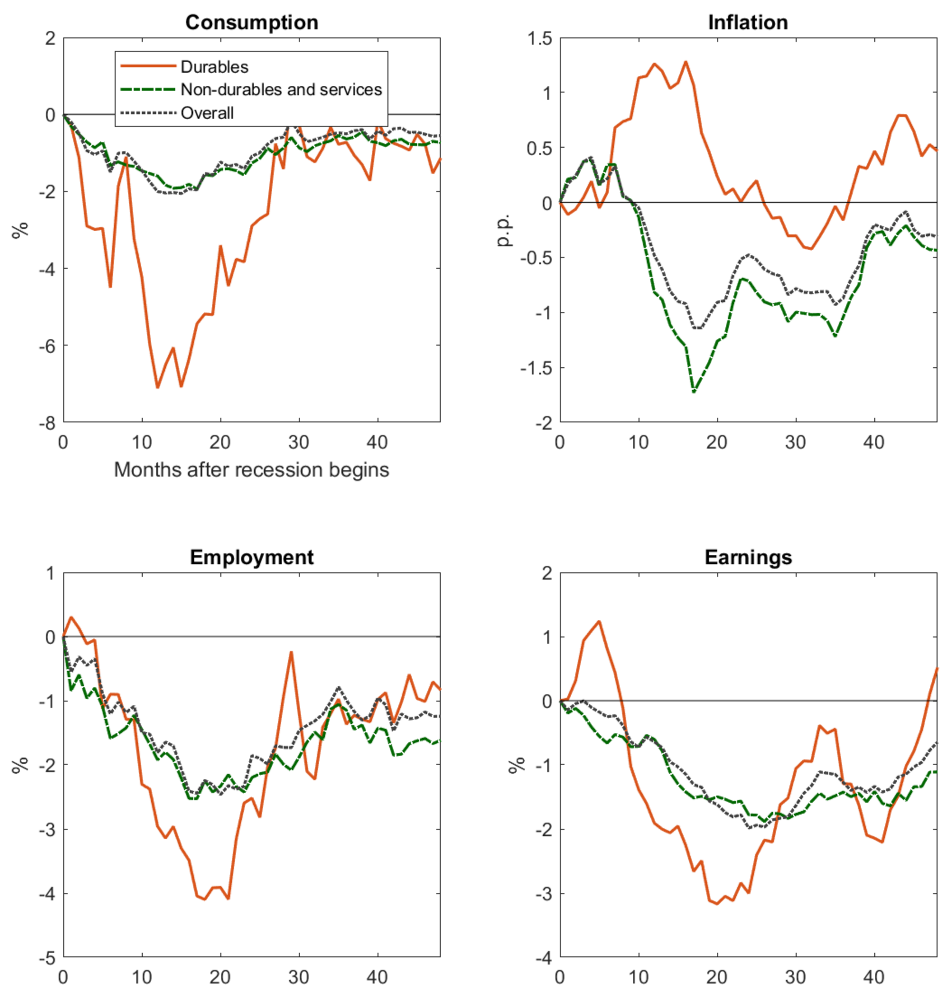

Equivalent of Figure 2 for of durables and non-durables. Series starting from the peak of the previous expansion, as defined by NBER. Includes all recession peaks since 1973 where non-essential and essential series for each variable are available for a full 48 months after the peak (peaks in 1973m11, 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 inital 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 inital period by taking the peak observation from all periods.

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 Figures A.6 and A.7. Focusing on the heterogeneity be-

Figure A.5: Earnings distribution - non-essentials and essentials, split into durables vs non-durables and services

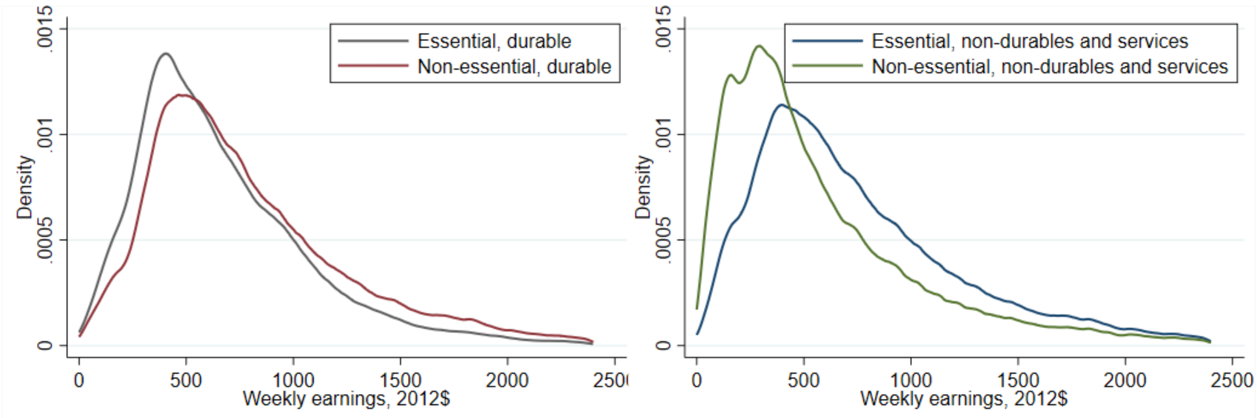

Notes: Kernel density of earnings 1982-2020, pooled, for non-essentials and essentials, split by durables vs non-durables and services.

tween non-durable essentials and non-durable non-essentials, we find 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 cyclicality 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.

Classification of durable/non-durable consumption and industries To understand the differences between durables and non-durables compared with non-essentials and essentials, we also construct series from disaggregated data for this alternative classification. 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

Figure A.6: IRFs of consumption - non-essentials and essentials, split into durables and non-durables

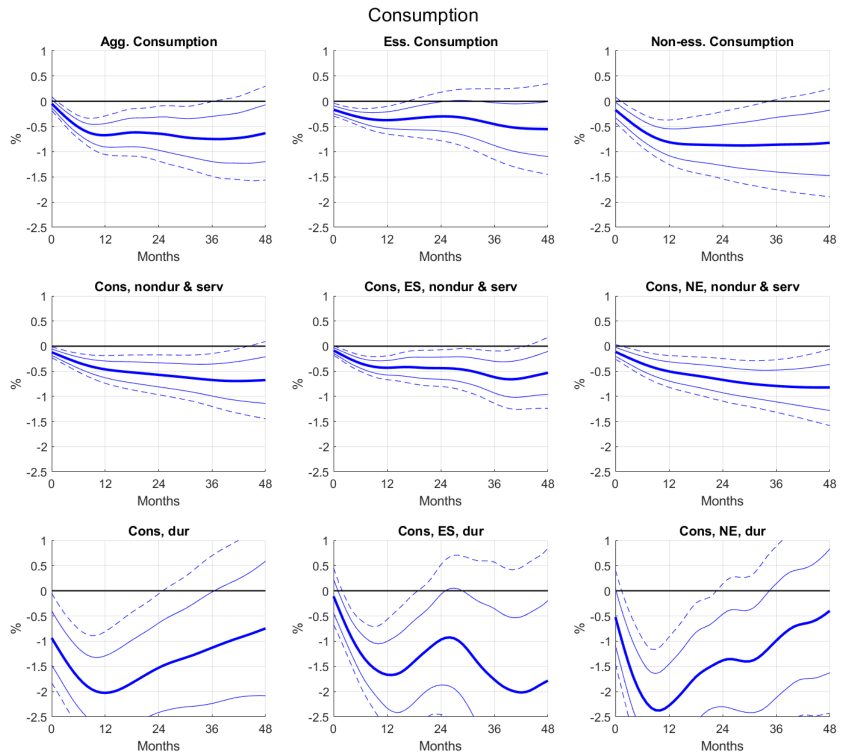

Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument. The sample periods and controls are as described in the main text. 68% and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

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.

Figure A.7: IRFs of earnings - non-essentials and essentials, split into durables and non-durables

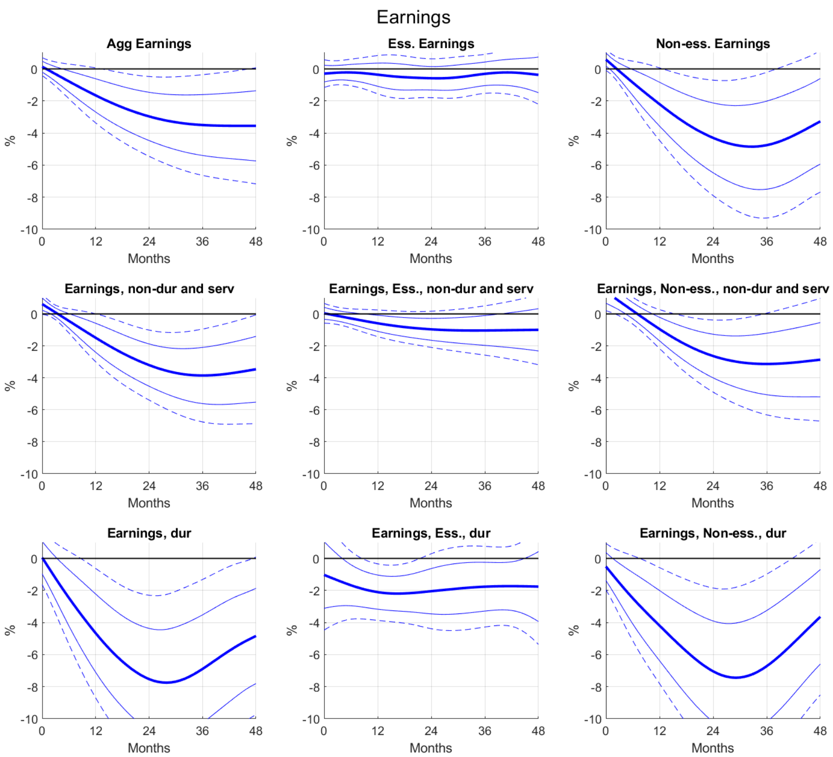

Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument. The sample periods and controls are as described in the main text. 68% and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

A.7 Other consumption categorisations

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 A.8 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.

Table A.8: Alternative categorisations

	Durables	Non-durables	Goods	Services	Tradeables	Non-tradables
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 tradaeables/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 catgeories 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).

To make this more concrete, Figure A.8 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, Rebelo and Wong (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

Figure A.8: Consumption IRFs of example consumption categories, goods versus services

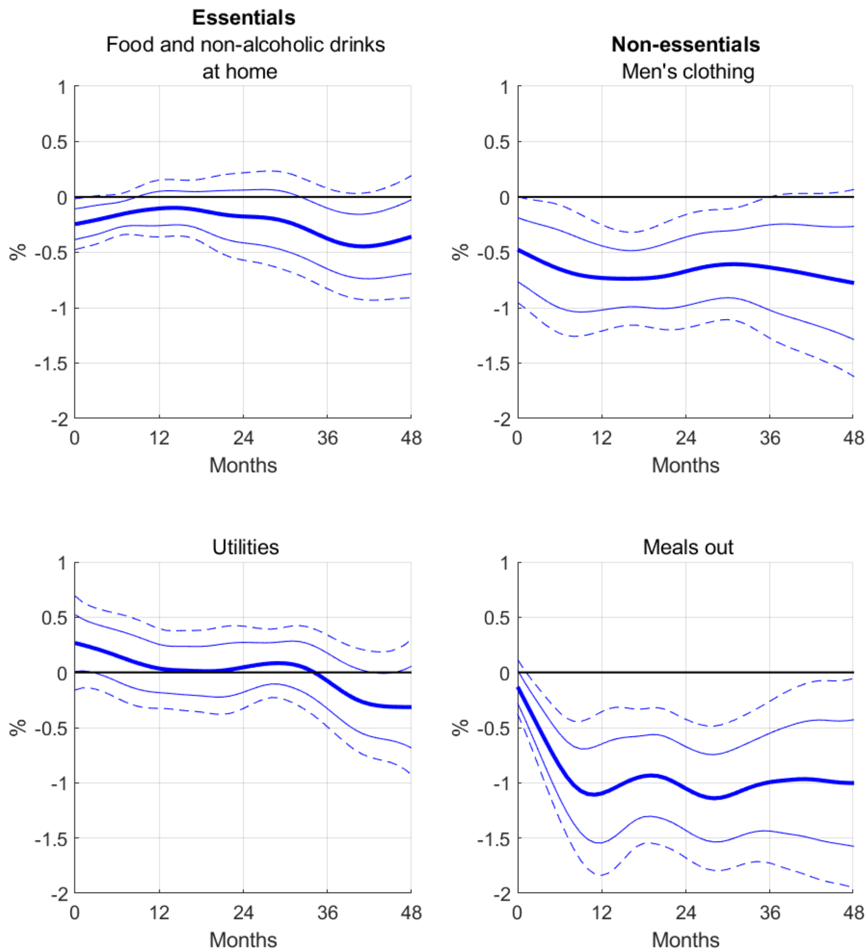

Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument. The sample period and controls are as described in the main text. 68% and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

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, Rebelo and Wong (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 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.

A.8 State-level analysis methodology

Figure 4 in the main text shows the the correlation between state-level employment changes during recessions and state-level non-essential consumption shares.

To construct this we used:

- Monthly BLS state-level employment data, derived from the CES.
  - We used the raw (non-seasonally adjusted) series, which starts in 1939, and seasonally adjusted it using the X-13ARIMA-SEATS approach. This gives seasonally adjusted state level employment series. These only start in 1973, due to limits on how long the series you can seasonally adjust can be using this procedure (but covers most of our sample period).
  - To identify state-specific recession dates by identifying the state-specific peak and trough of employment within 12 months before/after the NBER recession dates, excluding states where employment did not decline.
- State-level PCE series. The BEA provides these annually for 1997-present. The consumption categories available are somewhat more aggregated than those we are using for our main analysis, so the average non-essential shares do not exactly correspond. Non-essential shares are consumption shares from the BEA's state-level annual PCE series. We average these over the entire sample available for the series shown on the x-axis.

B SLP-IV implementation details

The point estimates for the IRFs for the SLP-IV approach have been estimated using the procedure suggested in Barnichon and Brownlees (2019):

1. We estimate a (standard) first stage by regressing the 1-year yield on the instrument and controls, and extract the predicted values of the endogenous variables $\hat{x}$
2. Use the predicted values in the SLP approach (following the notation in Barnichon and Brownlees (2019)):
   - $\hat{\mathcal{X}}_{\beta,t}$ is a d_t x K matrix where the (h,k)th element is $B_k(h)\hat{x}_t$, and this is stacked with the control matrices in the same way to produce the matrix $\hat{\chi}$.
3. Estimate the second stage SLP by generalised ridge regression: $\hat{\theta} = (\hat{\mathcal{X}}'\hat{\mathcal{X}} + \lambda\mathbf{P})^{-1}\hat{\mathcal{X}}'Y$

λ is selected using a five-fold cross-validation procedure, as suggested by Barnichon and Brownlees. We shrink towards a B-spline of order 2, which shrinks towards a line.

The SLP Newey-West standard errors Barnichon and Brownlees (2019) suggest are:

$$\begin{aligned}\widehat{V}(\hat{\theta}) =& T\left[\sum_{t=1}^{T-H_{\min}} \mathcal{X}_t'\mathcal{X}_t + \lambda\mathbf{P}\right]^{-1}\left[\hat{\Gamma}_0 + \sum_{l=1}^{L} w_l\left(\hat{\Gamma}_l + \hat{\Gamma}_l'\right)\right]\\ &\times\left[\sum_{t=1}^{T-H_{\min}} \mathcal{X}_t'\mathcal{X}_t + \lambda\mathbf{P}\right]^{-1}\end{aligned}$$

where $w_l = 1 - l/(1+L)$ and $\hat{\Gamma}_l = \frac{1}{T}\sum_{l+1}^{T-H_{\min}} \mathcal{X}_t'\hat{\mathcal{U}}_t\hat{\mathcal{U}}_{t-l}'\ \mathcal{X}_{t-l}$ where $\hat{\mathcal{U}}_t$ are the residuals from the second stage.

To contruct SLP-IV SEs, we use the generated regressor equivalent of this:

$$\begin{aligned}\widehat{V}(\hat{\theta}) =& T\left[\sum_{t=1}^{T-H_{\min}} \hat{\mathcal{X}}_t'\hat{\mathcal{X}}_t + \lambda\mathbf{P}\right]^{-1}\left[\hat{\Gamma}_0 + \sum_{l=1}^{L} w_l\left(\hat{\Gamma}_l + \hat{\Gamma}_l'\right)\right]\\ &\times\left[\sum_{t=1}^{T-H_{\min}} \hat{\mathcal{X}}_t'\hat{\mathcal{X}}_t + \lambda\mathbf{P}\right]^{-1}\end{aligned}$$

where $w_l = 1 - l/(1+L)$ and $\hat{\Gamma}_l = \frac{1}{T}\sum_{l+1}^{T-H_{\min}} \hat{\mathcal{X}}_t'\hat{\mathcal{U}}_t\hat{\mathcal{U}}_{t-l}'\ \hat{\mathcal{X}}_{t-l}$. Following Hansen (2021) $\hat{\mathcal{U}}_t = Y - \mathcal{X}\hat{\theta}$ are the residuals used, ie using the controls $\mathcal{X}$ constructed using the actual values of x rather than the first stage predicted values $\hat{x}$. If we set $\lambda = 0$, so no smoothing and penalising the results, this is the same as standard Newey-West standard errors for LP-IV. The autocorrelation lag used is the minimum between the Newey-West (1994) suggestion ($T^{1/4}$) and a linear increase with the estimation horizon. In an omitted robustness check, we also use lag-augmentation, with an extra lag of the controls and White standard errors, which set L=0, so no correction for auto-correlation.

B.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, 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.

B.2 Monetary policy surprises

Figure B.1 shows the monetary policy surprise series extracted from the proxy-SVAR. The Gertler-Karadi series is the main surprise series used in the empirical section, and Jarocinski-Karadi series is used in Appendix C.5.

C Additional empirical results and robustness

This section shows some additional empirical 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 accounting for the information effect, using alternative, more general business cycles and a longer sample, including COVID. We also show that our results hold using standard (not smoothed) local projection specification and our earnings responses at a quarterly frequency.

C.1 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 cyclicality of demand for non-essentials.

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

Figure B.1: Monetary policy surprise series

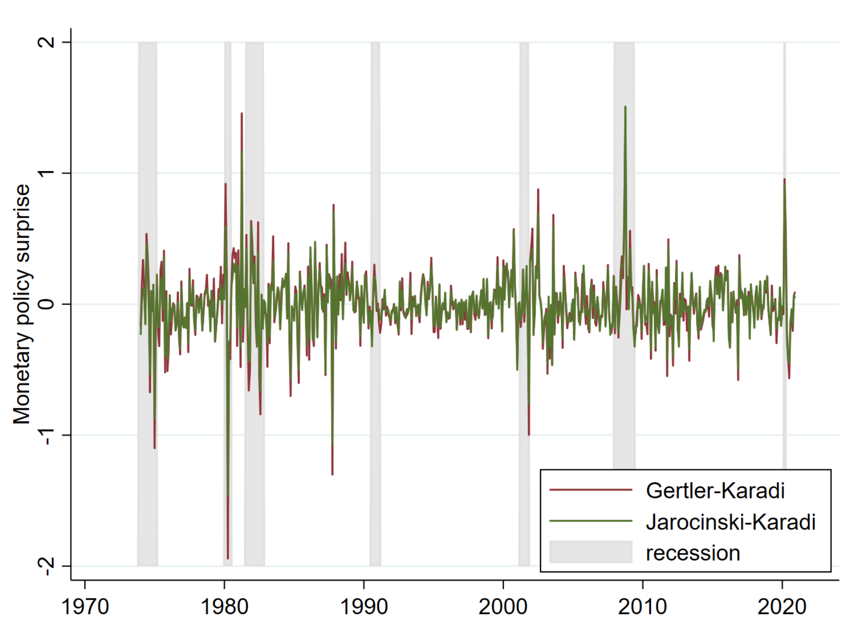

Notes: Monetary policy surprises, extracted from a proxy SVAR as described in Section 3.1. The Gertler-Karadi surprises are extracted from a proxy SVAR estimated using the (updated) monetary policy instrument proposed by Gertler and Karadi (2015), while the Jarocinski-Karadi surprises are from using the monetary policy instrument robust to the information effect proposed by Jarocinski and Karadi (2020).

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.2 Earnings distribution IRFs

Figure C.2 shows the IRFs of earnings percentiles show in the main text, Figure 6, with their confidence intervals.

C.3 IRFs for other macro aggregates

Figure C.3 shows the IRFs estimated using our SLP-IV specification for the other macroeconomic aggregate series used as controls and in the Proxy-SVAR. These are 1y yields, the excess bond premium, industrial production and the PCE price index. The results are broadly consistent with standard responses, for instance those given in Gertler and Karadi (2015) using their HFI instrument and SVAR. The shock is a 100bp exogenous rise in 1y yields, after which 1y yields fall and here fall significantly below their prior level by four years after the shock, rather than reverting back to their prior level. The excess bond premium rises about half the amount of 1y yields, but reverts to zero by 18 months after the shock. Industrial production falls by 2% by 15 months after the shock before recoving and becoming insignificantly different from zero by three years after the shock. Aggregate prices fall insignificantly.

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

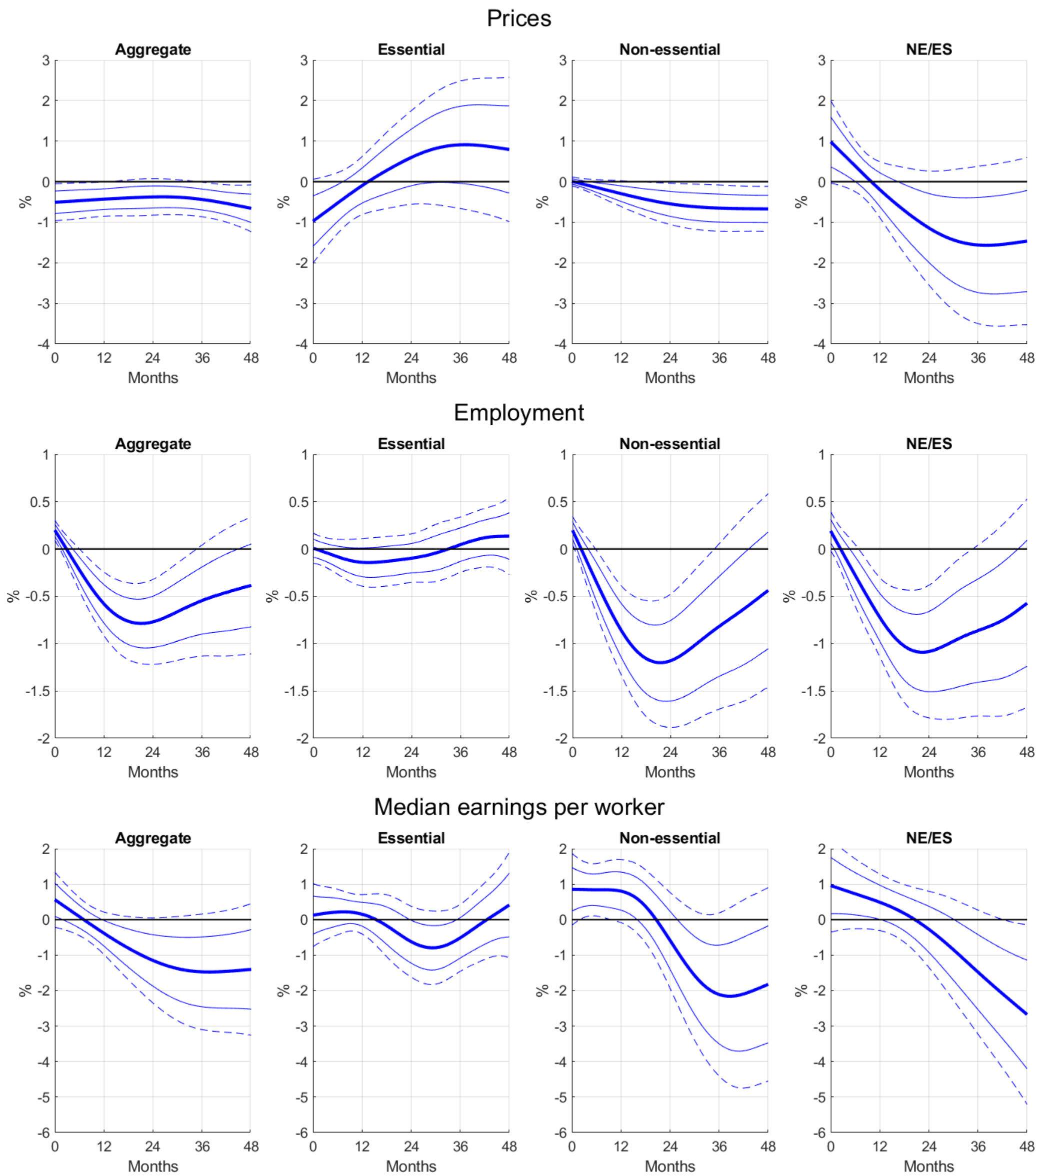

Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument. Sample periods and controls are specified in the main text and Appendix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

Figure C.2: IRFs to contractionary monetary policy shock - Earnings distribution

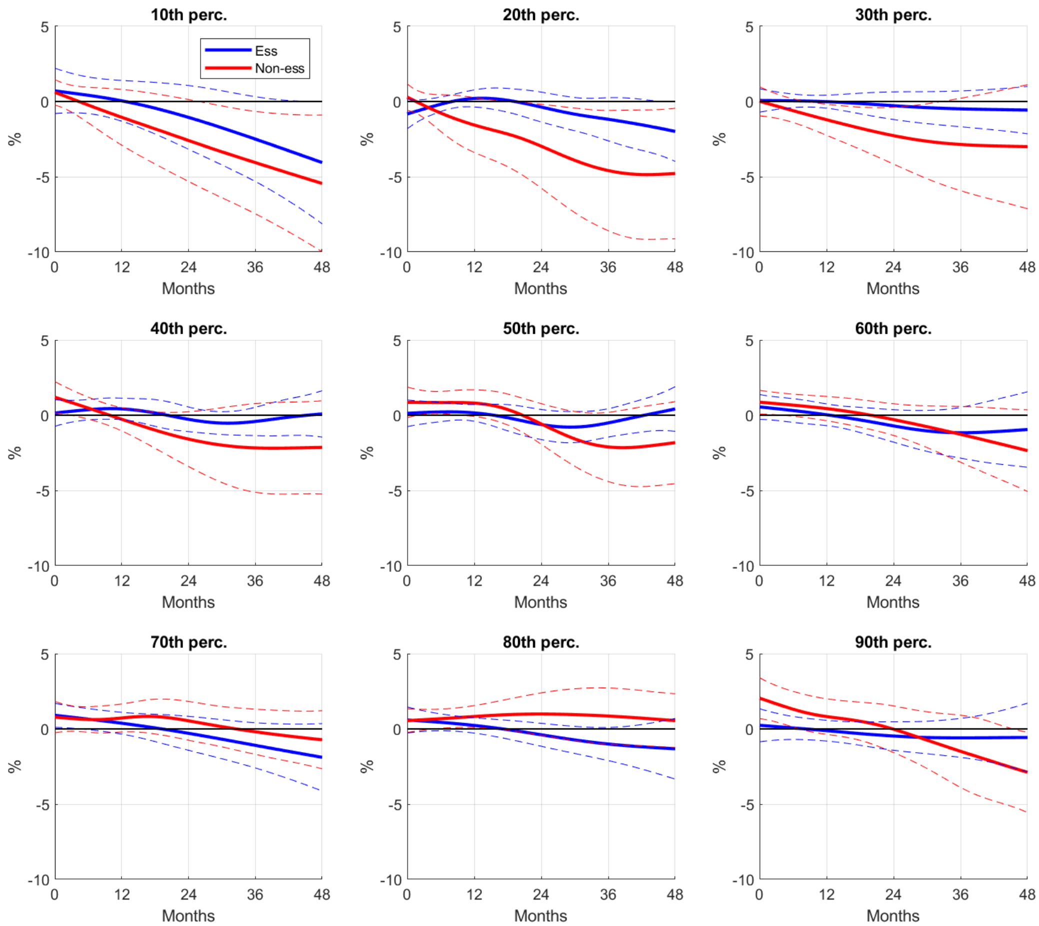

Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument. 90% confidence intervals displayed. Sample periods and controls are specified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

Figure C.3: IRFs to contractionary monetary policy shock - Macro aggregates

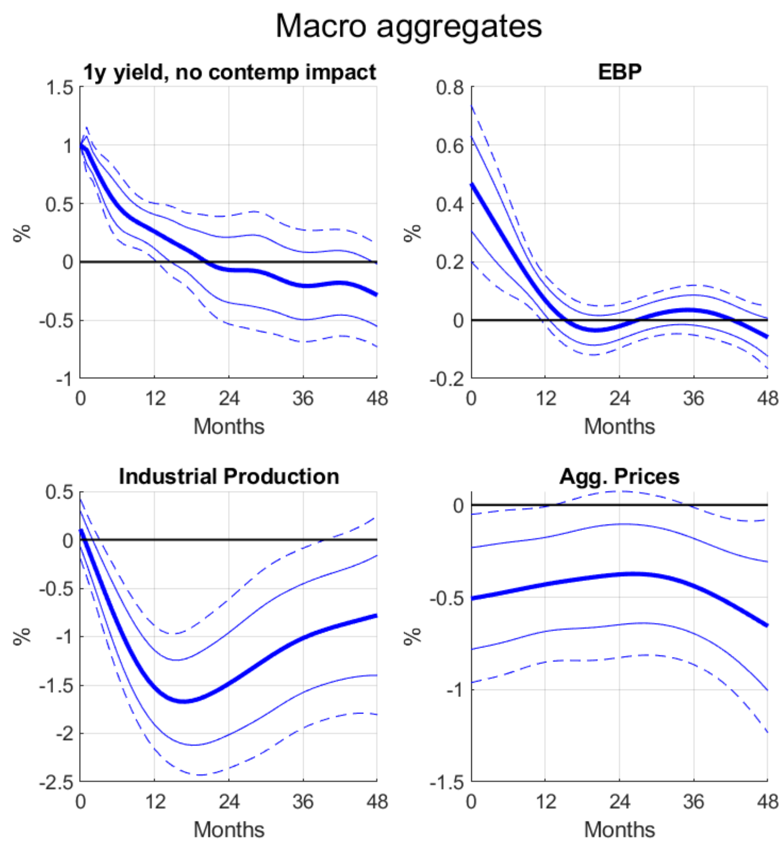

Notes: IRFs estimated by smooth local projections, response to 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument.

Figure C.4: IRFs to contractionary monetary policy shock - Comparison of total earnings series

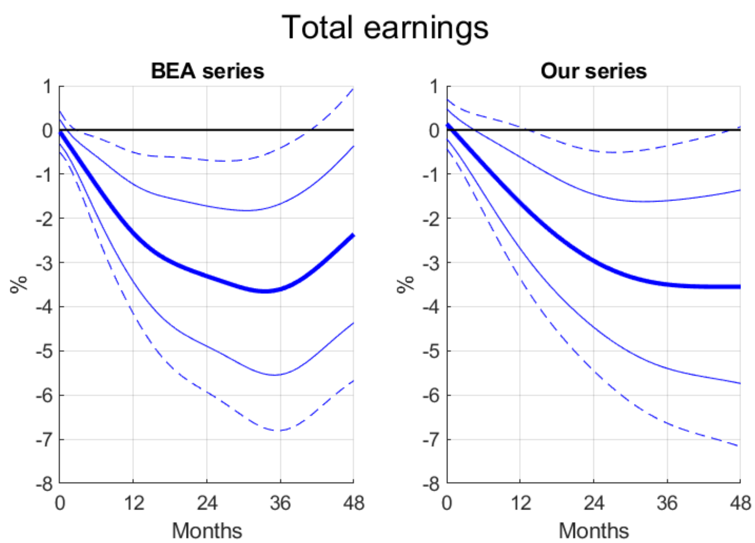

Notes: IRFs estimated by smooth local projections, response to 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument. The LHS series is the IRF of total compensation of employees (Received: Wage and Salary Disbursements) from the BEA NIPA data. The RHS series is the IRF our constructed equivalent series using CPS data.

C.4 Additional earnings distribution results

Figure C.5 shows the PDF and CDF of earnings distributions, plus the proportion of employment in non-essentials across the earnings deciles. The CDF demonstrates the the the CDF of essential earnings first order stochastically dominates the CDF of non-essential earnings.

Figure C.5: Non-essential and essential - Earnings distribution

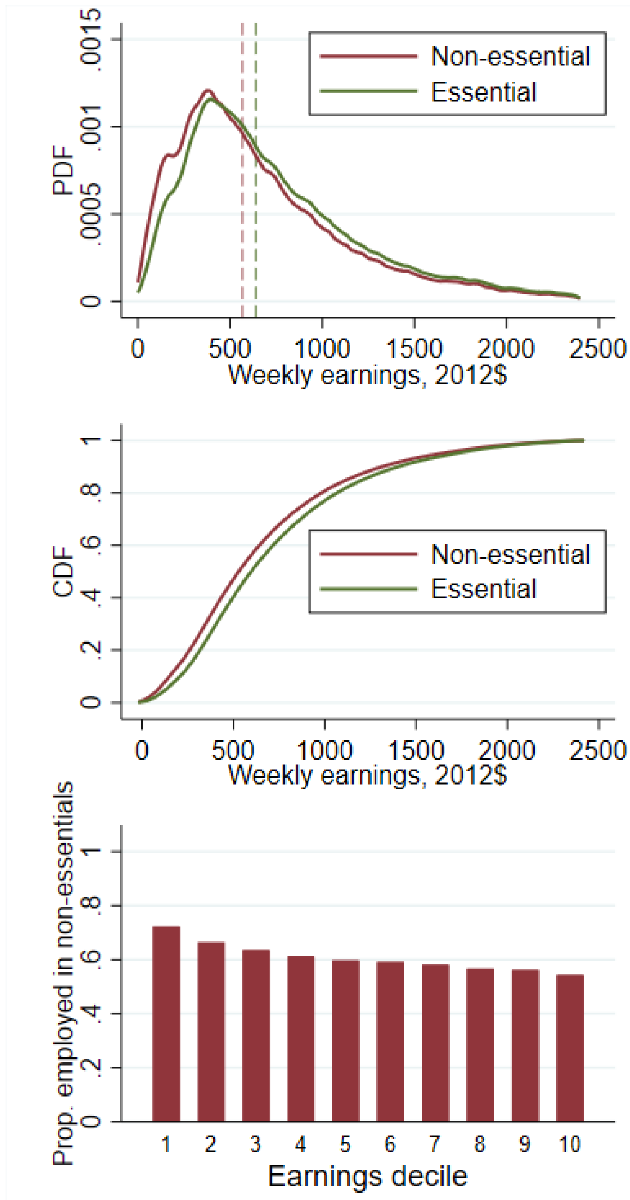

Notes: Earnings distributions within essential and non-essential industries. Underlying data is pooled Jan 1982 - December 2020, from the CPS, as described in the text. Panel 1 shows the kernel density plot along the median of each distribution, panel 2 shows the corresponding CDF, and panel 3 shows the percent of employees working in non-essential industries for each decile of the income distribution (deciles computed annually).

C.5 Robustness to information effect

In the main body, we use monetary policy shocks derived from Gertler and Karadi's high-frequency identified monetary policy instrument. Later literature, particularly Nakamura and Steinsson (2018), Jarociński and Karadi (2020) and Miranda-Agrippino and Ricco (2021) have explored how these shocks may be confounded by an 'information effect', if the monetary policy announcement also conveys information about the state of the economy that was privately held by the central bank.

To ensure our results are robust to this, we also estimate IRFs using monetary policy shocks derived from the instrument from Jarociński and Karadi (2020) - specifically, the shock to the Fed Funds futures (FF4) if there is a negative correlation between the FF4 surprise and the SP500 surprise. The rationale behind this instrument is that if there is privately held positive news about the economy that is part of the reason for a more contractionary monetary policy decision, then this may also result in a positive response of the stock market. Removing these instances leaves monetary policy surprises that are less likely to be contaminated by

Figure C.6: IRFs to contractionary monetary policy shock - Consumption and Prices

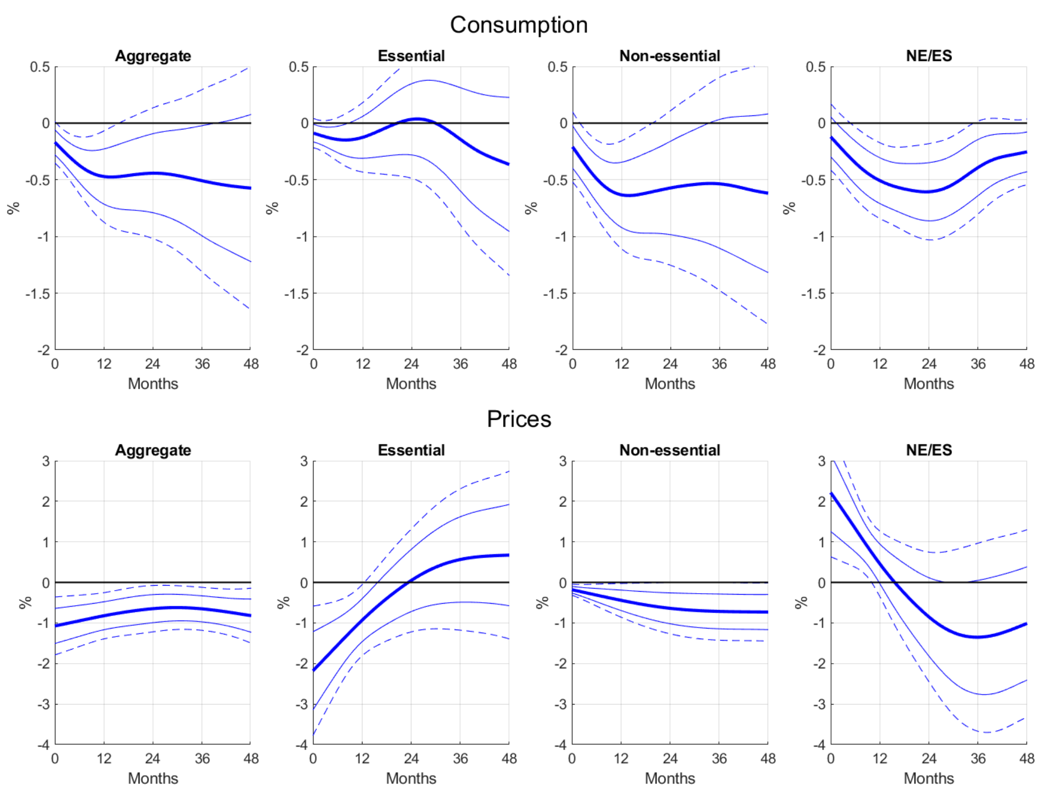

Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Jarocinski and Karadi (2020) high-frequency identified monetary policy instrument, robust to the information effect.

the information effect.

Figures C.6, C.7 and C.8 are the counterparts to the IRFs in the main text, using these alternative monetary policy surprise series. Sample and controls are as in the baseline regressions. The results are substantively similar; consumption, prices, employment and earnings (particularly at the lower end of the distribution) in non-essentials fall more than in essentials. However, the difference between essentials and non-essentials responses of prices and median earnings is no longer significant at the 90% level.

Figure C.7: IRFs to contractionary monetary policy shock - Employment and Earnings

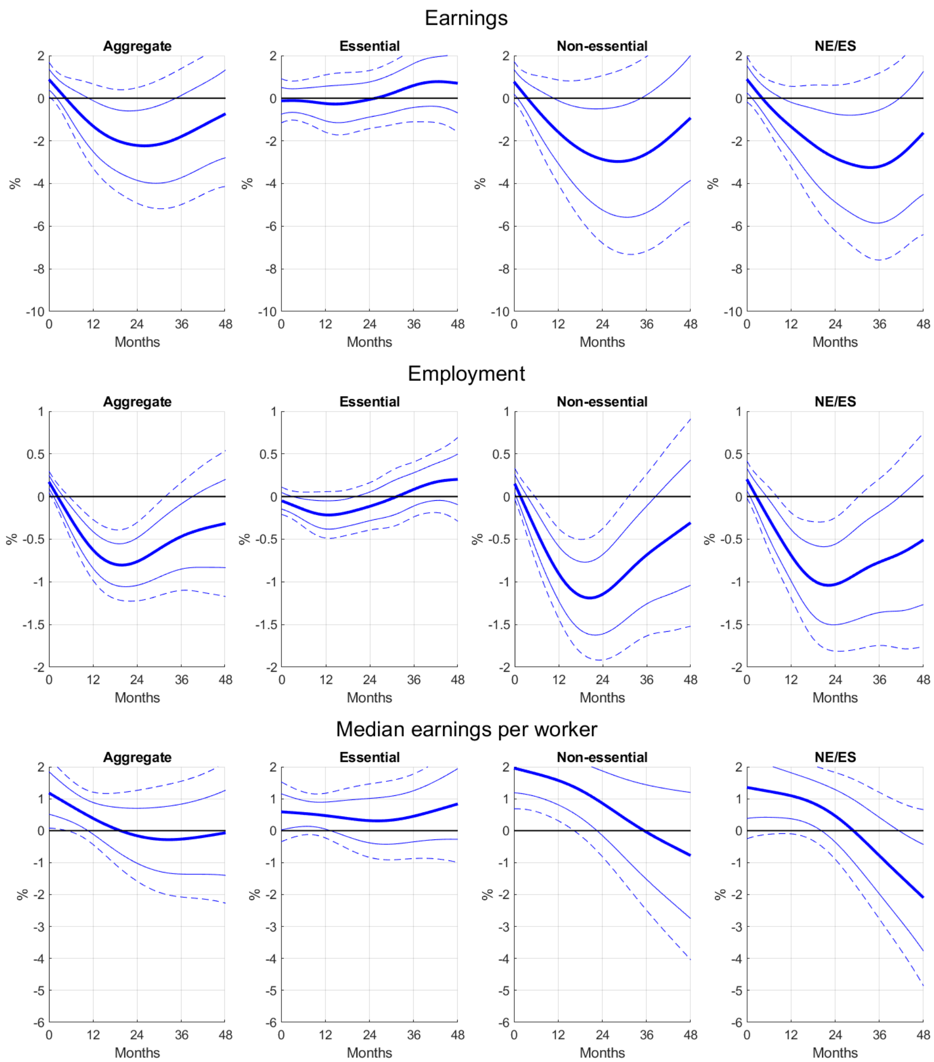

Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, i instrumented using monetary policy shocks derived from Jarocinski and Karadi (2020) high-frequency identified monetary policy instrument, robust to the information effect.

Figure C.8: IRFs to contractionary monetary policy shock - Earnings distribution

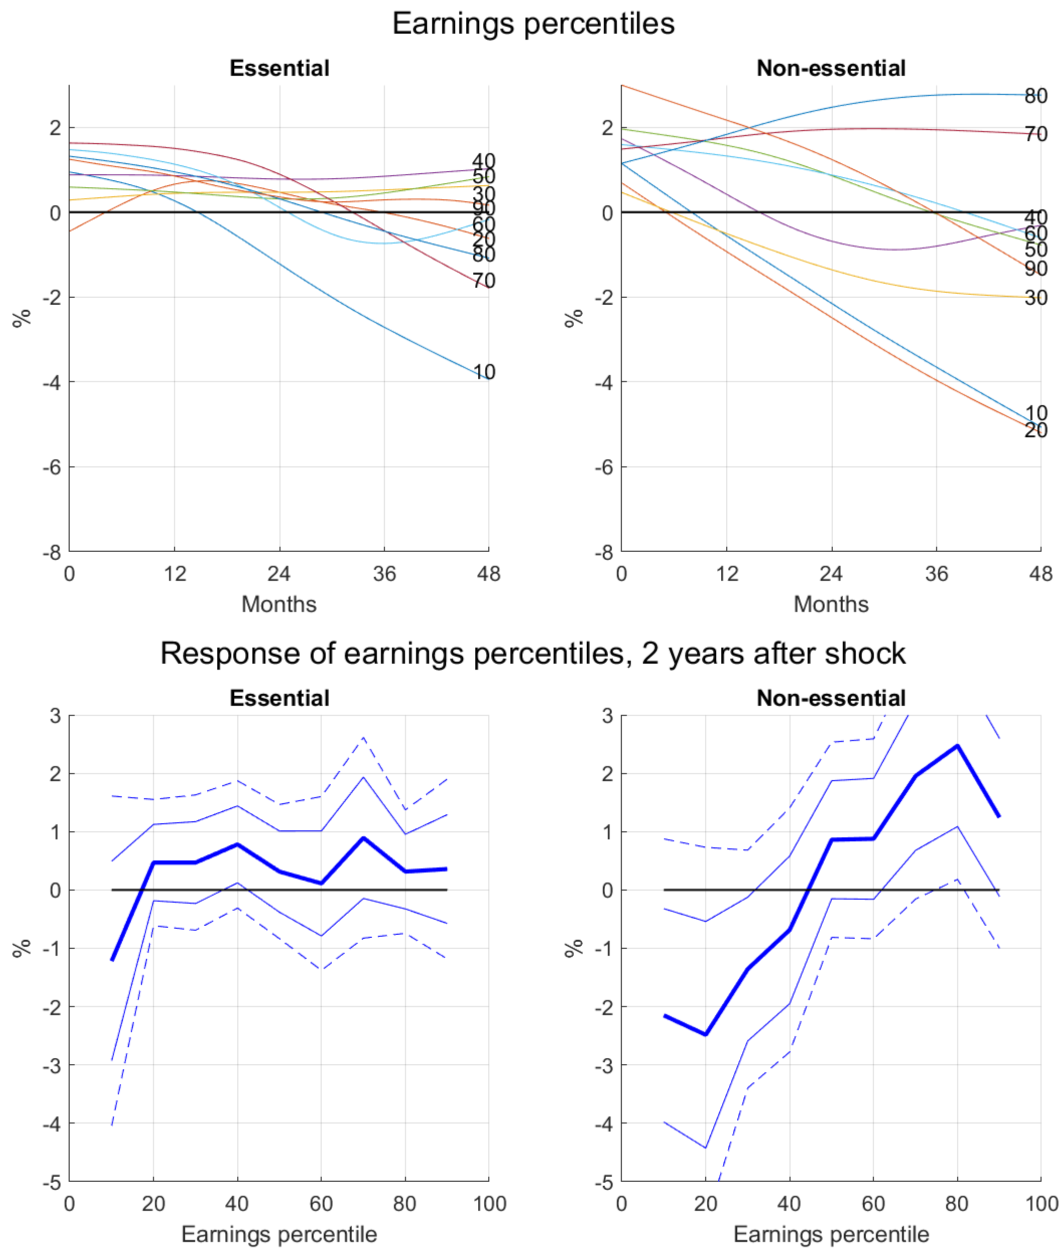

Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Jarocinski and Karadi (2020) high-frequency identified monetary policy instrument, robust to the information effect. Earnings IRFs estimated on data January 1982-December 2020, with 12 lags of 1 year yields, the excess bond premium, industrial production, consumption, prices, earnings, employment, and the LHS variable as controls. Earnings percentiles are from the CPS, and percentiles are calculated separately for the non-essential and essential earnings distributions. 90% confidence intervals displayed.

C.6 Shocks from Business Cycle Anatomy (Angeletos, Collard and Dellas (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, Collard and Dellas (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, Collard and Dellas (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.

Figure C.9 shows the results. 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.9: Response to business cycle shock from Angeletos, Collard and Dellas (2020)

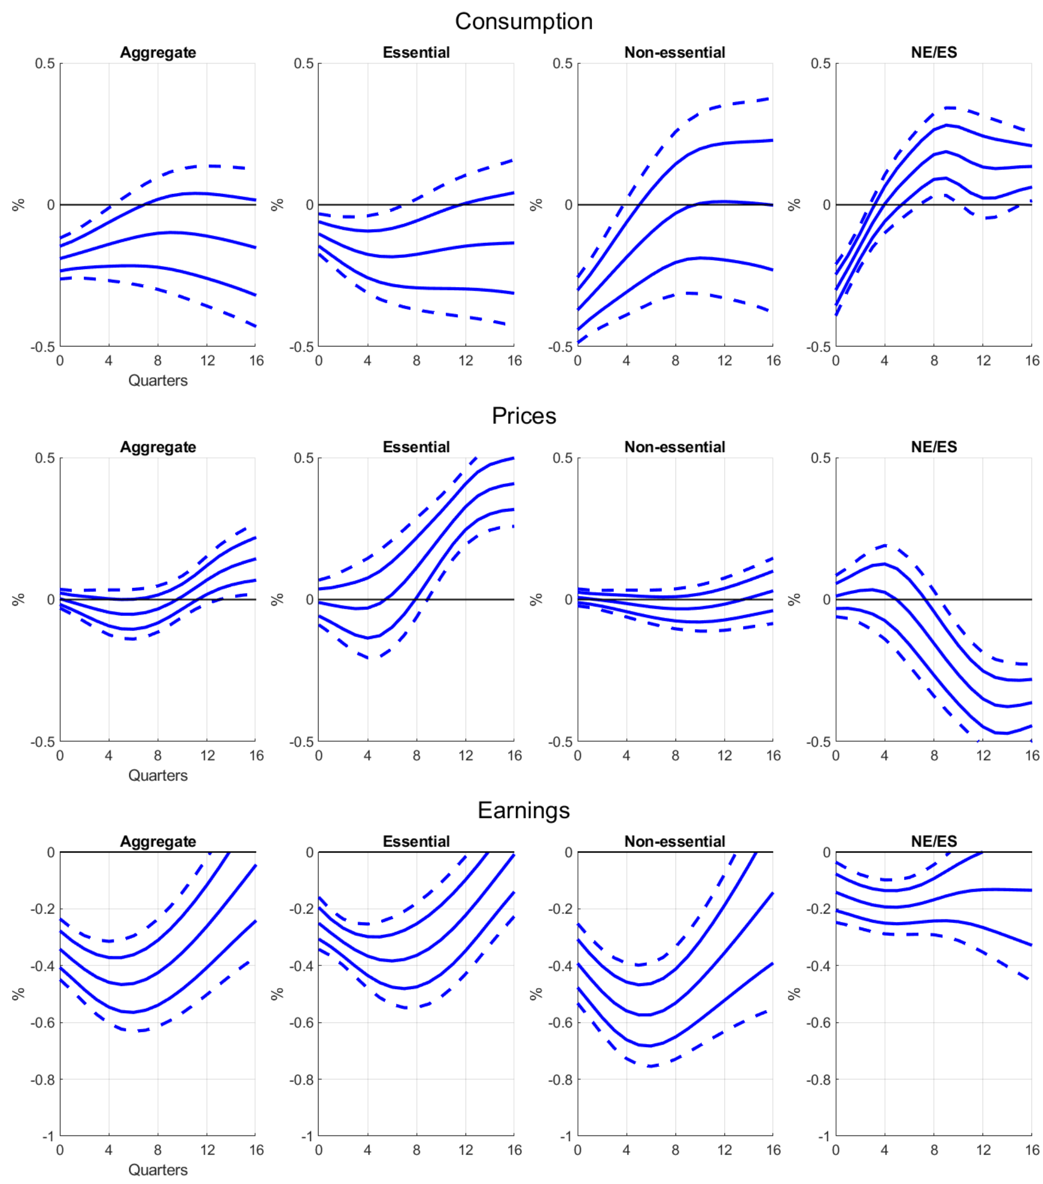

Notes: We use the shock that maximises the variation in unemployment, which is the main shock used by Angeletos, Collard and Dellas (2020). Results estimated using SLP, same specification as the main results, other than no instrument used and quarterly frequency and sample ending in 2017q4.

C.7 Adding COVID to the sample period

In our main sample, we end the estimation period in December 2019. This omits the effects of Covid-19, where non-essentials and essentials responded differently to the shock partly due to sector-specific reductions in activity not directly driven by the mechanism we propose here[31]. To check that our results are robust to adding the effects of the Covid-19 period, we estimate the IRFs for samples ending in December 2020 in Figures C.10 and C.11. The magnitude and degree of heterogeneity in responses is increased with this sample, but in our main results we prefer to focus on the more conservative set of results, excluding Covid, to ensure that only entirely voluntary deferral of non-essential consumption is considered.

[31] Although we envisage that a large reason for the differential shutdowns of different sectors were precisely because certain types of consumption are not intertemporally substitutable, consistent with our mechanism, and that our identification strategy of estimating the response to monetary policy shocks should alleviate this issue

Figure C.10: IRFs to contractionary monetary policy shock - Consumption and Prices

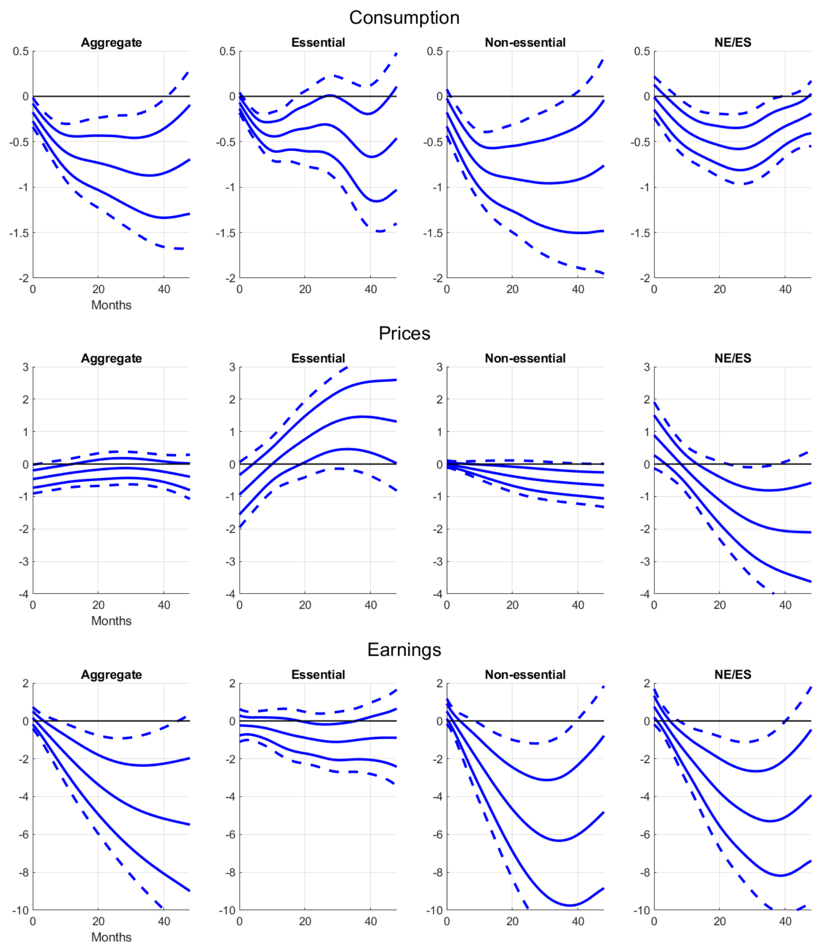

Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument, robust to the information effect. Sample period ends in December 2020, otherwise specification is as in the main text.

Figure C.11: IRFs to contractionary monetary policy shock - Earnings distribution

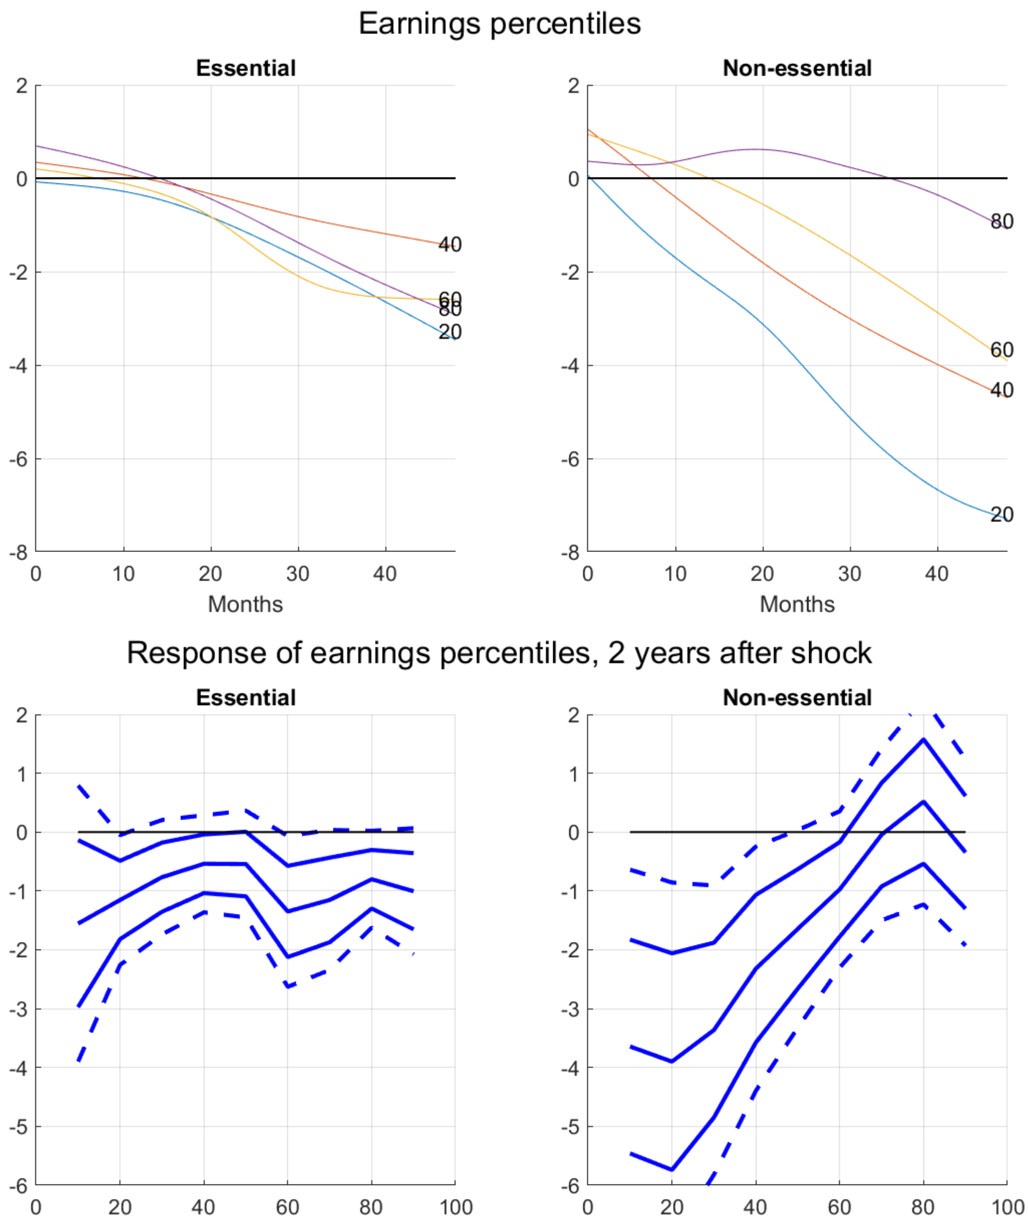

Notes: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument, robust to the information effect. Sample ends December 2020, otherwise the specification remains in the main body of the text. 68 and 90% confidence intervals displayed.

C.8 IRFs with (unsmoothed) local projections

To show that our results are robust to using standard local projections, rather than smoothed local projections, Figures C.12 shows our main results for consumption, prices and earnings are similar for standard LP, but the introduction of smoothing allows us to more clearly see the key results. C.13 shows the IRFs for selected percentiles of the earnings distribution; due to the noise in the earnings series, it is harder to see clear patterns from the LP results.

Figure C.12: IRFs to contractionary monetary policy shock - Consumption, Prices and Earnings

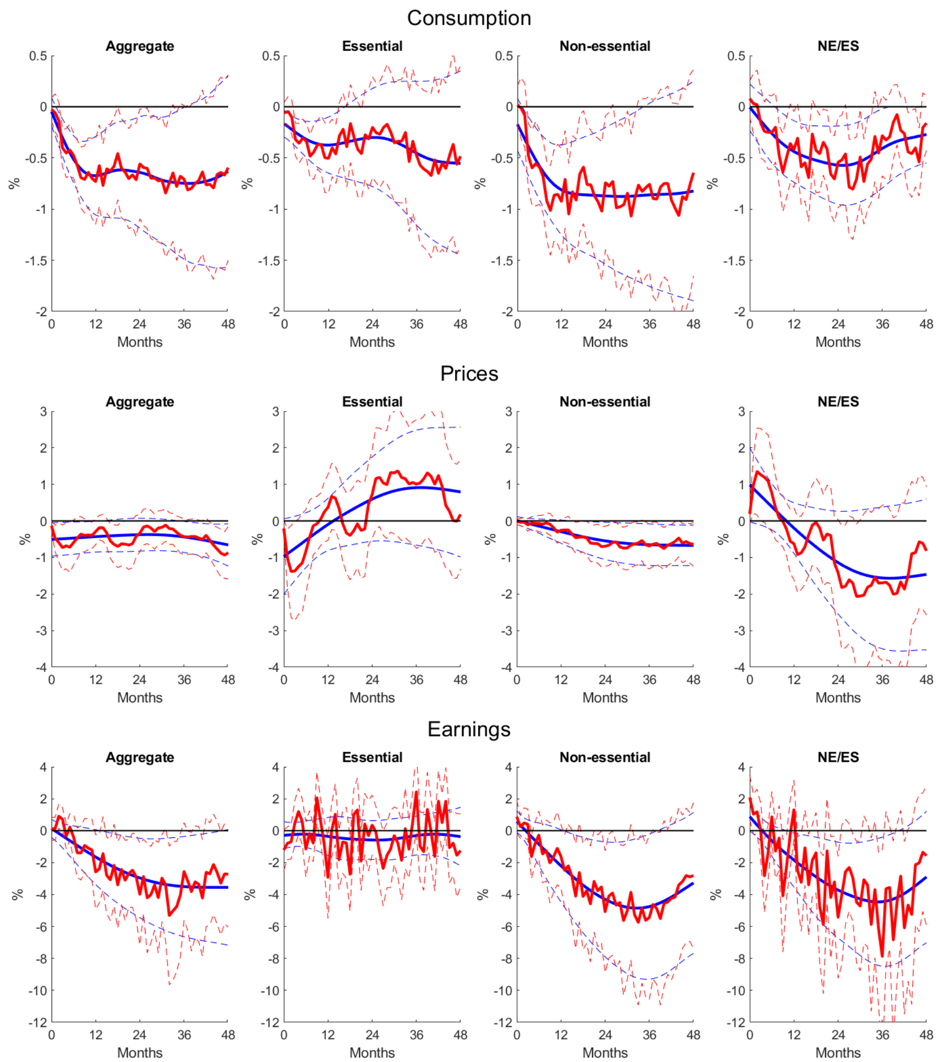

Notes: IRFs estimated by smooth local projections (blue) and standard local projections (red), response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument, robust to the information effect. Samples and specifications as described in the main text.

Figure C.13: IRFs to contractionary monetary policy shock - Earnings distribution

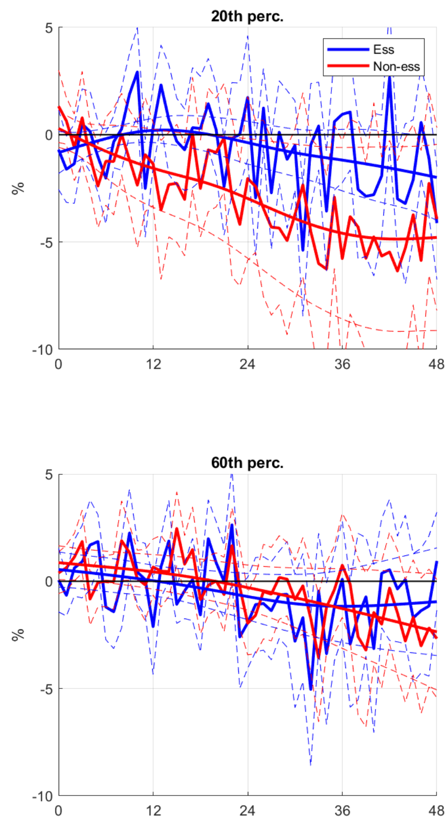

Notes: IRFs estimated by smooth local projections (smooth IRFs) and standard local projections (non-smooth IRFs), response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument, robust to the information effect. Sample and specification as in main text. 90% confidence intervals.

C.9 Quarterly earnings IRFs

The CPS ORG sample is formally designed to be representative only at the quarterly frequency, but in our main results we use monthly frequency. To verify our results still hold at the lower frequency, Figure C.14 shows our main results for earnings using quarterly frequency data. As the quarterly frequency removes some of the higher frequency variation useful for identifying responses, the results are less significant but qualitatively similar to the baseline results.

Figure C.14: IRFs to contractionary monetary policy shock - Earnings at quarterly frequency

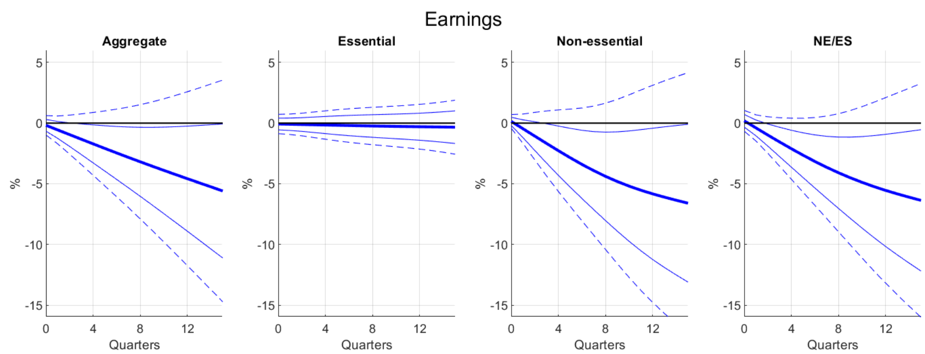

Notes: IRFs estimated by smooth local projections (smooth IRFs) and standard local projections (non-smooth IRFs), response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument, robust to the information effect. Sample and specification as in main text, quarterly frequency data used. 90% and 68% confidence intervals.

D Model derivations

In this appendix, we provide more detailed derivations for the theoretical model. The objective is to highlight the solution method, the steady state computation, and the full log-linear equilibrium.

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 $a_{H,t}$:

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

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: $C^E_{i,t,0}$ and for the future if they don't update $C^E_{i,t+j,j}$ 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.

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

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:

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

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

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

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:

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

Use the recursive structure to write the FOC as a traditional Euler equation:

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

Use a similar method to massage the FOC for j periods ahead:

$$(C^E_{H,t+j,j})^{-\frac{1}{\gamma^E}} = \mathbb{E}_t(C^E_{H,t+j,0})^{-\frac{1}{\gamma^E}}$$

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

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

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. Start with the time zero problem:

$$(C^N_{H,t,0})^{-\frac{1}{\gamma^N}}\varphi = p^N_t\beta\lambda\sum_{m=0}^{\infty}\beta^m(1-\lambda)^m\mathbb{E}_tV'(a_{H,t+m+1})\prod_{k=0}^{m}\tilde{R}_{t+k+1}$$
$$(C^N_{H,t,0})^{-\frac{1}{\gamma^N}}\varphi = p^N_t(C^E_{H,t,0})^{-\frac{1}{\gamma^E}}$$

The FOC for j periods ahead:

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

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,[32] 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.

$$\begin{aligned}
(C^E_{H,t,0})^{-\frac{1}{\gamma^E}} &= \beta \mathbb{E}_t \left((C^E_{H,t+1,0})^{-\frac{1}{\gamma^E}} \frac{R_t}{\pi^E_{t+1}} \right) \\
\varphi (C^N_{H,t,0})^{-\frac{1}{\gamma^N}} &= p^N_t (C^E_{H,t,0})^{-\frac{1}{\gamma^E}} \\
(C^E_{H,t+j,j})^{-\frac{1}{\gamma^E}} &= \mathbb{E}_t \left((C^E_{H,t+j,0})^{-\frac{1}{\gamma^E}} \right) \\
(C^N_{H,t+j,j})^{-\frac{1}{\gamma^N}} &= \mathbb{E}_t \left((C^N_{H,t+j,0})^{-\frac{1}{\gamma^N}} \right)
\end{aligned}$$

The Ricardian agents budget constraint, which drops out in the equilibrium definition due to Walras law:

$$C^E_{H,t} + p^N_t C^N_{H,t} + b_{H,t} = w_{H,t} N_{H,t} + \Pi^r_{H,t} + t_{H,t} + R_{t-1} b_{H,t-1} \frac{R_{t-1}}{\pi^E_t}$$

Consumption aggregation across attentive and non-attentive consumers:

$$\begin{aligned}
C^E_{H,t} &= \lambda \sum_{j=0}^{\infty} (1-\lambda)^j C^E_{H,t-j,j} \\
C^N_{H,t} &= \lambda \sum_{j=0}^{\infty} (1-\lambda)^j C^N_{H,t-j,j}
\end{aligned}$$

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

$$\begin{aligned}
C^E_{H,t} &= \lambda \sum_{j=0}^{\infty} (1-\lambda)^j \left[\mathbb{E}_{t-j} \left((C^E_{H,t,0})^{-\frac{1}{\gamma^E}} \right) \right]^{-\gamma^E} \\
C^N_{H,t} &= \lambda \sum_{j=0}^{\infty} (1-\lambda)^j \left[\mathbb{E}_{t-j} \left((C^N_{H,t,0})^{-\frac{1}{\gamma^N}} \right) \right]^{-\gamma^N}
\end{aligned}$$

[32] Notice 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: $-\frac{\partial\left(\frac{C^N_{H,t,0}}{C^E_{H,t,0}}\right)}{\partial p^N_t} \frac{p^N_t}{\left(\frac{C^N_{H,t,0}}{C^E_{H,t,0}}\right)} = \gamma^E - \frac{\gamma^N-\gamma^E}{\gamma^N} \frac{\partial C^N_{H,t,0}}{\partial p^N_t} \frac{p^N_t}{C^N_{H,t,0}} = \gamma^N - \frac{\gamma^N-\gamma^E}{\gamma^E} \frac{\partial C^E_{H,t,0}}{\partial p^N_t} \frac{p^N_t}{C^E_{H,t,0}}$. As $\frac{\partial C^N_{H,t,0}}{\partial p^N_t} \leq 0$ and $\frac{\partial C^E_{H,t,0}}{\partial p^N_t} \geq 0$ we can see that the elasticity of substitution across the two goods is bounded between γ^E and γ^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 satisfied[33].

First, we show the budget constraint in terms of wealth:

$$
\begin{aligned}
P_t^E C_{L,t}^E + P_t^N C_{L,t}^N &\leq W_{L,t} N_{L,t} + \Pi_{L,t} + T_{L,t} \\
C_{L,t}^E + p_t^N C_{L,t}^N &\leq a_{L,t} = w_{L,t} N_{L,t} + \Pi_{L,t}^r + t_{L,t}
\end{aligned}
$$

Their maximisation problem, for the periods in which they cannot update:

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

To find the solution, take the FOC for the two goods and equate the marginal utilities:

$$
\beta^j (1-\lambda)^j \left(C_{L,t+j,j}^E\right)^{-\frac{1}{\gamma^E}} = \beta^j (1-\lambda)^j \left(C_{L,t+j,j}^N\right)^{-\frac{1}{\gamma^N}} \varphi \frac{1}{\mathbb{E}_t(p_{t+j}^N)}
$$

Use this condition to arrive to the three equilibrium conditions as for the Ricaridan agents, minus the Euler equation:

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

We can still aggregate goods consumption across attentive and non-attentive consumers. By

[33] This 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.

assuming risk sharing across consumers, agents can follow through with their plans ex-post.

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

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 consumption to compute the labour supply choice. We end up with two standard intra-temporal equilibrium conditions:

$$\xi \frac{N^{\chi}_{L,t}}{(C^E_{L,t})^{-\frac{1}{\gamma^E}}} = w_{L,t}$$
$$\xi \frac{N^{\chi}_{H,t}}{(C^H_{L,t})^{-\frac{1}{\gamma^E}}} = w_{H,t}$$

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.

$$Y^i_t = \left(\int_0^1 (y^i_{k,t})^{\frac{\varepsilon-1}{\varepsilon}} dk \right)^{\frac{\varepsilon}{\varepsilon-1}} \quad i = \{E, N\}$$

This leads to a standard demand that the final good producers have for different varieties of a given good category (essential and non-essentials):

$$y^i_{k,t} = Y^i_t \left(\frac{P^i_{k,t}}{P^i_t} \right)^{-\varepsilon} \quad i = \{E, N\}$$

Calvo retailers. Retailers of a given type of good, say essential[34], buy a wholesale good of the same type at a wholesale price $P^{E,w}_t$ and use it to produce the retail variety $y^E_{k,t}$ 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.

[34] The non-essential retailers one is fully symmetric.

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:

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

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 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 household[35], and rearrange to arrive to three equations representing the non-linear New Keynesian Phillips Curve.

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

For the non-essential goods we can solve the same problem and arrive to the same equations:

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

[35]As 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.

$$\frac{K_t^{N,f}}{F_t^{N,f}} = \left(\frac{1-\theta(\pi_t^N)^{\varepsilon^N-1}}{1-\theta}\right)^{\frac{1}{1-\varepsilon^N}}$$

D.3.1 Wholesalers

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

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

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 optimisation problem of the essential wholesalers is:

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

For the non-essential wholesalers:

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

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 marker 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:

$$1 = \mu_H + \mu_L$$

The market clearing conditions for the two goods markets:

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

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

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

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. It is imporant to note that 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.

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

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:

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

The only role of fiscal policy 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.

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

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 policy[36]:

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

D.6 Equilibrium

The competitive equilibrium consists of 29 endogenous allocations $\{C_t$, C_t^E, C_t^N, $C_{H,t}^E$, $C_{H,t}^N$, $C_{L,t}^E$, $C_{L,t}^N$, $C_{H,t,0}^E$, $C_{H,t,0}^N$, $C_{L,t,0}^E$, $C_{L,t,0}^N$, $N_{H,t}$, $N_{L,t}$, $N_{H,t}^E$, $N_{H,t}^N$, $N_{L,t}^E$, $N_{L,t}^N$, $b_{H,t}$, $\Pi_{H,t}^r$, $\Pi_{L,t}^r$, $\Pi_t^{r,N}$, $\Pi_t^{r,E}$, Y_t, Y_t^E, Y_t^N, $K_{E,t}^f$, $F_{E,t}^f$, $K_{N,t}^f$, $F_{N,t}^f\}$, 13 prices $\{w_{H,t}$, $w_{L,t}$, π_t, π_t^E, π_t^N, $\pi_{t,Lasp}$, $\pi_{t,Paasche}$, p_t^N, P_t^E, P_t^N, R_t, $\mathcal{S}_t^E$, $\mathcal{S}_t^N\}$, and 3 exogenous processes $\{A_t^E$, A_t^N, $\varepsilon_t^{mp}\}$, with P_0^E normalised to one; such that households, final good producers, retailers, and wholesalers optimise, the central bank follows a Taylor rule, the treasury follows the tax rule, profits are disbursed according to the profit rule, and markets clear. To avoid repetition, we re-write the full set of equations only in the linearised equilibrium.

[36] Recall capital letter Π are profits, small letter π are inflation rates.

D.7 Steady state computation

We define a steady state variable simply without the time subscript. We solve for a zero-inflation steady state ($\pi^E = \pi^N = 1$). We set the transfers to the Calvo retailers at $\tau^E = 1/\varepsilon^E$ and $\tau^N = 1/\varepsilon^N$ to ensure no steady state markups ($\mathcal{S}^E = \mathcal{S}^N = 1$) and zero steady state profits. We normalise the steady state price level for the essential good at 1 ($P^E = 1$) and solve for the steady state relative price p^N.

Wages. We solve for wages from the wholesalers problem. As long as $\alpha^E \neq \alpha^N$, the formula is:

$$w_L = (p^N)^{\frac{1-\alpha^E}{\alpha^N-\alpha^E}} \left[\left(A^E (1-\alpha^E)^{1-\alpha^E} (\alpha^E)^{\alpha^E} \right)^{-(1-\alpha^N)} \left(A^N (1-\alpha^N)^{1-\alpha^N} (\alpha^N)^{\alpha^N} \right)^{(1-\alpha^E)} \right]^{\frac{1}{\alpha^N-\alpha^E}} \tag{5}$$

$$w_H = (p^N)^{\frac{-\alpha^E}{\alpha^N-\alpha^E}} \left[\left(A^E (1-\alpha^E)^{1-\alpha^E} (\alpha^E)^{\alpha^E} \right)^{\alpha^N} \left(A^N (1-\alpha^N)^{1-\alpha^N} (\alpha^N)^{\alpha^N} \right)^{-\alpha^E} \right]^{\frac{1}{\alpha^N-\alpha^E}} \tag{6}$$

Consumption. To solve for consumption, first note that in steady state attentive and inattentive consumers all have the same consumption level. Next, plug the labour supply choice, the intra-temporal choice between essential and non-essential goods in the budget constraint and use the zero profit and zero transfer in steady state. This leads for each household $k = \{H, L\}$ to a one non-linear equation in the consumption of essentials:

$$C_k^E + \varphi^{\gamma^N} \left(p^N\right)^{1-\gamma^N} (C_k^E)^{\frac{\gamma^N}{\gamma^E}} = w_k^{1+\frac{1}{\chi}} \xi^{-\frac{1}{\chi}} (C_k^E)^{-\frac{1}{\chi\gamma^E}} \quad k = \{H, L\} \tag{7}$$

With non-homotheticity, this equation cannot be solved analytically, but can be solved easily numerically.

Algorithm to find the steady state. For a given set of structural parameters, we compute the steady state with the following algorithm. Vary p^N such that we compute:

1. w_H, and w_L analytically with (5) and (6).
2. C_H^E and C_L^E numerically with (7).
3. C_H^N, C_L^N, N_H, N_L from the household/union problem.
4. Y^E and Y^N from the goods market clearing conditions.
5. N_H^E and N_L^N from firms' labour demand functions.

6. The difference between $N_H^E + N_H^N$ and $\mu_H N_H$.

Iterate on p^N until the difference it is zero. Alternatively, the last step can be substituted with the difference between $N_L^E + N_L^N$ and $\mu_L N_L$ by Walras law (one market clearing condition can be ignored).

In each estimation draw, we target the steady state consumption shares of Ricardian and hand-to-mouth agents of non-essentials: $\bar{C}_H^N \equiv \frac{p^N C_H^N}{p^N C_H^N + C_H^E}$ and $\bar{C}_L^N \equiv \frac{p^N C_L^N}{p^N C_L^N + C_L^E}$. To do so, we vary the relative preference parameter for non-essentials φ and the relative productivity of the two sectors $a^E \equiv A^E/A^N$. φ affects the average consumption share. a^E affects the relative wage, and, therefore, the relative consumption shares, thanks to the non-homotheticity in the utility function.

D.8 Log-linear equilibrium

We solve the log-linearised model. Steps are standard, we log-linearise each variable, except for profits, which we linearise as they are zero in steady state. Log-linearised and linearised variables are hatted. The only feature to note is that all CPI inflation indices simplify to the same steady states weighted average of inflation:

$$\hat{\pi}_t = \hat{\pi}_{t,Lasp} = \hat{\pi}_{t,Paasche} = \frac{C^E}{C^E + p^N C^N}\hat{\pi}_t^E + \frac{p^N C^N}{C^E + p^N C^N}\hat{\pi}_t^N$$

Equilibrium. The competitive equilibrium consists of 25 endogenous allocations $\{\hat{C}_t$, $\hat{C}_t^E$, $\hat{C}_t^N$, $\hat{C}_{H,t}^E$, $\hat{C}_{H,t}^N$, $\hat{C}_{L,t}^E$, $\hat{C}_{L,t}^N$, $\hat{C}_{H,t,0}^E$, $\hat{C}_{H,t,0}^N$, $\hat{C}_{L,t,0}^E$, $\hat{C}_{L,t,0}^N$, $\hat{N}_{H,t}$, $\hat{N}_{L,t}$, $\hat{N}_{H,t}^E$, $\hat{N}_{H,t}^N$, $\hat{N}_{L,t}^E$, $\hat{N}_{L,t}^N$, $\hat{\Pi}_{L,t}^r$, $\hat{\Pi}_t^{r,N}$, $\hat{\Pi}_t^{r,E}$, $\hat{Y}_t$, $\hat{Y}_t^E$, $\hat{Y}_t^N$, $\hat{Earn}_t^E$ $\hat{Earn}_t^N\}$, 9 prices $\{\hat{w}_{H,t}$, $\hat{w}_{L,t}$, $\hat{\pi}_t$, $\hat{\pi}_t^E$, $\hat{\pi}_t^N$, $\hat{p}_t^N$, $\hat{R}_t$, $\hat{\mathcal{S}}_t^E$, $\hat{\mathcal{S}}_t^N\}$, and 3 exogenous processes $\{\hat{A}_t^E$, $\hat{A}_t^N$, $\varepsilon_t^{mp}\}$; such that households, final good producers, retailers, and wholesalers optimise, the central bank follows a Taylor rule, the treasury follows the tax rule, profits are disbursed according to the profit rule, and markets clear. The equilibrium is characterised by the following equations:

$$-\frac{1}{\gamma^E}\hat{C}_{H,t,0}^E + \frac{1}{\gamma^N}\hat{C}_{H,t,0}^N = -\hat{p}_t^N$$

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

$$\hat{C}_{H,t}^E = \lambda\sum_{j=0}^{\infty}(1-\lambda)^j\mathbb{E}_{t-j}\left(\hat{C}_{H,t,0}^E\right)$$

$$\hat{C}_{H,t}^N = \lambda\sum_{j=0}^{\infty}(1-\lambda)^j\mathbb{E}_{t-j}\left(\hat{C}_{H,t,0}^N\right)$$

$$
\begin{aligned}
-\frac{1}{\gamma^E}\hat{C}^E_{L,t,0} + \frac{1}{\gamma^N}\hat{C}^N_{L,t,0} &= -\hat{p}^N_t \\
C^E_L\hat{C}^E_{L,t} + p^N C^N_L(\hat{p}^N_t + \hat{C}^N_{L,t}) &= w_L N_L(\hat{w}_{L,t} + \hat{N}_{L,t}) + \frac{\hat{\Pi}^r_{L,t}}{\mu_L} \\
\hat{C}^E_{L,t} &= \lambda\sum_{j=0}^{\infty}(1-\lambda)^j\mathbb{E}_{t-j}\left(\hat{C}^E_{L,t,0}\right) \\
\hat{C}^N_{L,t} &= \lambda\sum_{j=0}^{\infty}(1-\lambda)^j\mathbb{E}_{t-j}\left(\hat{C}^N_{L,t,0}\right) \\
\chi\hat{N}_{H,t} + \frac{1}{\gamma^E}\hat{C}^E_{H,t} &= \hat{w}_{H,t} \\
\chi\hat{N}_{L,t} + \frac{1}{\gamma^E}\hat{C}^E_{L,t} &= \hat{w}_{L,t} \\
\hat{\pi}^N_t &= \beta\mathbb{E}_t(\hat{\pi}^N_{t+1}) + \kappa^N\hat{\mathcal{S}}^N_t \\
\hat{\pi}^E_t &= \beta\mathbb{E}_t(\hat{\pi}^E_{t+1}) + \kappa^E\hat{\mathcal{S}}^E_t \\
\pi^N_t &= \pi^E_t + p^N_t - p^N_{t-1} \\
\hat{Y}^N_t &= \hat{A}^N_t + \alpha^N\hat{N}^N_{L,t} + (1-\alpha^N)\hat{N}^N_{H,t} \\
\hat{\mathcal{S}}^N_t + \hat{Y}^N_t - \hat{N}^N_{H,t} &= \hat{w}_{H,t} - \hat{p}^N_t \\
\hat{\mathcal{S}}^N_t + \hat{Y}^N_t - \hat{N}^N_{L,t} &= \hat{w}_{L,t} - \hat{p}^N_t \\
\hat{Y}^E_t &= \hat{A}^E_t + \alpha^E\hat{N}^E_{L,t} + (1-\alpha^E)\hat{N}^E_{H,t} \\
\hat{\mathcal{S}}^E_t + \hat{Y}^E_t - \hat{N}^E_{H,t} &= \hat{w}_{H,t} \\
\hat{\mathcal{S}}^E_t + \hat{Y}^E_t - \hat{N}^E_{L,t} &= \hat{w}_{L,t} \\
N^E_H\hat{N}^E_{H,t} + N^N_H\hat{N}^N_{H,t} &= \mu_H N_H\hat{N}_{H,t} \\
N^E_L\hat{N}^E_{L,t} + N^N_L\hat{N}^N_{L,t} &= \mu_L N_L\hat{N}_{L,t} \\
\hat{\pi}_t &= \frac{C^E}{C^E + p^N C^N}\hat{\pi}^E_t + \frac{p^N C^N}{C^E + p^N C^N}\hat{\pi}^N_t \\
Y\hat{Y}_t &= Y^E\hat{Y}^E_t + p^N Y^N\hat{Y}^N_t \\
\hat{R}_t &= \rho_R\hat{R}_{t-1} + (1-\rho_R)\left(\phi_\pi\left(\mathbb{E}_t(\hat{\pi}_{t+1})\right) + \phi_Y\hat{Y}_t\right) + \varepsilon^{mp}_t \\
\hat{\Pi}^r_{L,t} &= \phi^E_{\Pi,L}\hat{\Pi}^{r,E}_t + \phi^N_{\Pi,L}\hat{\Pi}^{r,N}_t \\
\hat{\Pi}^{r,E}_t &= -Y^E\hat{\mathcal{S}}^E_t \\
\hat{\Pi}^{r,N}_t &= -Y^N p^N\hat{\mathcal{S}}^N_t \\
C^E\hat{C}^E_t &= \mu_H C^E_H\hat{C}^E_{H,t} + \mu_L C^E_L\hat{C}^E_{L,t} \\
C^N\hat{C}^N_t &= \mu_H C^N_H\hat{C}^N_{H,t} + \mu_L C^N_L\hat{C}^N_{L,t} \\
\hat{Y}^E_t &= \hat{C}^E_t
\end{aligned}
$$

$$
\begin{aligned}
\hat{Y}_t^N &= \hat{C}_t^N \\
\hat{Earn}_t^E &= \frac{w_H N_H^E}{w_H N_H^E + w_L N_L^E}(\hat{w}_{H,t} + \hat{N}_{H,t}^E) + \frac{w_L N_L^E}{w_H N_H^E + w_L N_L^E}(\hat{w}_{L,t} + \hat{N}_{L,t}^E) \\
\hat{Earn}_t^N &= \frac{w_H N_H^N}{w_H N_H^N + w_L N_L^N}(\hat{w}_{H,t} + \hat{N}_{H,t}^N) + \frac{w_L N_L^N}{w_H N_H^N + w_L N_L^N}(\hat{w}_{L,t} + \hat{N}_{L,t}^N)
\end{aligned}
$$

Notice that the equilibrium conditions include four equations with an infinite sum of past expectations (the mapping from each inattentive consumer consumption to the family wide one). To solve the model with a state space representation, we adopt a method proposed by Verona and Wolters (2014) for sticky expectations models. We solve for a truncated set of past expectations. The key insight is that, if we care only about IRFs, as we do here (our estimation uses IRF matching), we can truncate the expectations at the horizon of the IRFs and have no loss in precision (say in period 16). $\mathbb{E}_{t-j}(\hat{C}_{H,t,0}^E)$ will be zero for each $j > 16$, that is before the shock happens.

E Model estimation and counterfactual

In this appendix, we present the estimation procedure, the full set of estimated IRFs, and the details of our counterfactual exercise.

E.1 Estimation

We estimate the model with a limited-information Bayesian approach, that is, with a impulse response matching with a maximum a posteriori (MAP) estimation procedure. We follow the estimation procedure of Mertens and Ravn (2011), with the weighting matrix choice of Guerron-Quintana, Inoue and Kilian (2017), extended to a MAP setting. Given our model, we estimate a vector of parameters Θ_2 (the parameters in Panel A of Table 1) conditional on a vector of calibrated parameters Θ_1 (the parameters in Panel B of Table 1). The quasi-likelihood:

$$F(\hat{\Lambda}_d|\Theta_2, \Theta_1) = \left(\frac{1}{2\pi}\right)^{\frac{T}{2}} |\Sigma_d| \exp\left[-\frac{1}{2}\left(\hat{\Lambda}_d - \Lambda(\Theta_2|\Theta_1)\right)' \Sigma_d^{-1} \left(\hat{\Lambda}_d - \Lambda(\Theta_2|\Theta_1)\right)\right]$$

This maps the difference in the estimated IRFs with smooth local projections $\hat{\Lambda}_d$ to the model based IRFs $\Lambda(\Theta_2|\Theta_1)$. We stack the IRFs in a vector of dimension T, in the baseline setting equal to 112 (16 quarters times 7 variables). As weighting matrix, we follow Guerron-Quintana, Inoue and Kilian (2017) and use a diagonal matrix with the squared standard errors from the smooth local projection estimates for each IRF element. We denote $p(\Theta_2)$ the prior distribution over the estimated parameters. We follow the common procedure of imposing bounds in the prior draws, but none bind at the estimated values. The quasi-posterior:

$$F(\Theta_2|\hat{\Lambda}_d, \Theta_1) \propto F(\hat{\Lambda}_d|\Theta_2, \Theta_1)p(\Theta_2)$$

Maximum a posterior estimation maximises the posterior over estimated parameters. The practical benefit, over frequentist impulse response matching matching, is that it allows to incorporate priors over parameters.

$$\hat{\Theta}_2 = \arg\max_{\Theta_2} F(\Theta_2|\hat{\Lambda}_d, \Theta_1)$$

We compute the standard errors of $\hat{\Theta}_2$ with the delta method. The formula for the asymptotic covariance matrix, from Mertens and Ravn (2011):

$$\begin{aligned}\Sigma_{\Theta_2} &= \Lambda_{\Theta_2}\frac{\partial\Lambda(\Theta_2|\Theta_1)'}{\partial\Theta_2}\Sigma_d^{-1}\Sigma_S\Sigma_d^{-1}\frac{\partial\Lambda(\Theta_2|\Theta_1)}{\partial\Theta_2}\Lambda_{\Theta_2}\\ \Lambda_{\Theta_2} &\equiv \left[\frac{\partial\Lambda(\Theta_2|\Theta_1)'}{\partial\Theta_2}\Sigma_d^{-1}\frac{\partial\Lambda(\Theta_2|\Theta_1)}{\partial\Theta_2}\right]^{-1}\\ \Sigma_S &\equiv \Sigma_d + \Sigma_m\end{aligned}$$

Where we use Σ_d in the last line, following Guerron-Quintana, Inoue and Kilian (2017). Notice that we use the model based IRFs, not the IRFs estimated on data simulated from the model as Mertens and Ravn (2013) do, so that $\Sigma_m = 0$ and the overall expression for the parameters covariance matrix simplifies.

Estimated IRFs. Figure E.1 shows the empirical IRFs in blue and the estimated IRFs in red for the whole set of matched variables. We estimate the end-quarter impulse response for these variables, as described in the main text. For consumption, earnings and prices we match the estimated IRFs for non-essentials and essentials using our SLP empirical approach. For 1y yields, we estimate the impulse response from the proxy-SVAR.

E.2 Counterfactual

In Table 2, we compared the cumulative response of aggregate consumption in counterfactual exercises. In this appendix, we complete this analysis by showing the dis-aggregated consumption responses by different goods.

Table E.1 shows the cumulative IRFs of non-essential and essential consumption between the non-homothetic and homothetic representative agent counterfactuals. As seen in Table 2, aggregate consumption responds equally in both cases. However, unlike for aggregates, non-homotheticity does change sectoral outcomes. Notice that our irrelevance result of Appendix Section F demonstrates the irrelevance of sectoral heterogeneities for aggregates in the representative agent setting (albeit for a simpler model than that used in the counterfactuals). This table demonstrates numerically the same result applies for the representative agent model used in the counterfactuals.

Figure E.1: IRFs to contractionary monetary policy shock - Matched variables from model

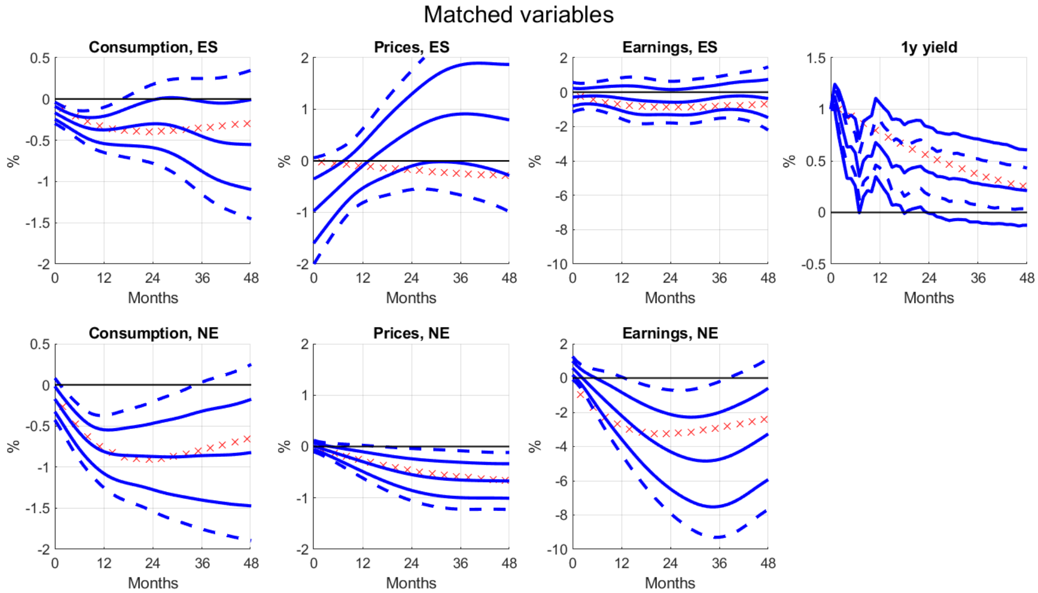

Notes: Consumption, prices, earnings: IRFs estimated by smooth local projections, response to a 100bp increase in 1y yields, instrumented using monetary policy shocks derived from Gertler and Karadi (2015) high-frequency identified monetary policy instrument. Sample periods and controls are specified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

Table E.1: Counterfactuals of Essentials and Non-essentials in the Representative Agent Model

	Representative Agent		
	C	C^E	C^N
Homothetic	1.00	1.00	1.00
Non-Homothetic	1.00	0.34	1.51

Notes: Each cell display the ratio of the cumulative IRF of counterfactual experiment over the cumulative IRF of representative agent with homothetic preferences with the estimated model parameters. The first columns shows aggregate consumption, the second essential consumption, and the third non-essential consumption. In the homothetic case, we set the IES equal to the estimated average IES in the baseline model.

F When non-homotheticity matters for aggregates, a proof

In this appendix, we present the a proof on when non-homotheticity does not amplify business cycles. We show that the non-homothetic RANK has the same response to monetary policy of aggregate variables then a homothetic RANK with the IES equal to the IES of the non-homothetic RANK. This implies that non-homotheticity does not matter per-se for amplification, but it matters only when interacts with other features, as labour market heterogeneity, finanacial constraints, price stickiness, heterogeneous capital intensities, etc. We formalize this idea with Proposition 1 and Corollary 1.

Proposition 1 *Consider a simplified version of the model of Sections 4 and D. Take an attentive representative agent version with non-homothetic utility (2) and a simplified Taylor rule of the form $R_t = \phi_\pi \mathbb{E}_t(\pi_{t+1}) + \varepsilon_t^{mp}$. The impact of the monetary policy shock on total consumption is characterised by the average intertemporal elasticity of substitution and on CPI inflation by the average intertemporal elasticity of substitution and the slope of the Phillips curves:*

$$\frac{\partial \hat{C}_t}{\partial \varepsilon_t^{mp}} = -(\underbrace{\gamma^E(1-\bar{C}^N) + \gamma^N \bar{C}^N}_{Average\ IES})$$
$$\frac{\partial \hat{\pi}_t}{\partial \varepsilon_t^{mp}} = -\underbrace{\kappa}_{Slope\ of\ NKPC}(1 + \underbrace{\gamma^E(1-\bar{C}^N) + \gamma^N \bar{C}^N}_{Average\ IES})$$

Corollary 1 *Consider a simplified version of the model of Sections 4 and D. Take an attentive representative agent version with homothetic utility*

$$U(C_t^E, C_t^N, N_t) = \frac{(C_t^E)^{1-\frac{1}{\gamma}}}{1-\frac{1}{\gamma}} + \varphi \frac{(C_t^N)^{1-\frac{1}{\gamma}}}{1-\frac{1}{\gamma}} - \xi \frac{N_t^{1+\chi}}{1+\chi}$$

such that the intertemporal elasticity of substitution γ is equal to $\gamma^E(1-\bar{C}^N) + \gamma^N \bar{C}^N$ of the model presented in Proposition 1, and a simplified Taylor rule of the form $R_t = \phi_\pi \mathbb{E}_t(\pi_{t+1}) + \varepsilon_t^{mp}$. The impact of the monetary policy shock on total consumption is characterised by the intertemporal elasticity of substitution and on CPI inflation by the intertemporal elasticity of substitution and the slope of the Phillips curves:

$$\frac{\partial \hat{C}_t}{\partial \varepsilon_t^{mp}} = -\underbrace{\gamma}_{IES}$$

$$\frac{\partial \hat{\pi}_t}{\partial \varepsilon_t^{mp}} = - \underbrace{\kappa}_{Slope\ of\ NKPC} (1 + \underbrace{\gamma}_{IES})$$

We now move to prove both statements. The intuition of the result is that relative prices are a state variable but they do not respond to an aggregate shock in the representative agent model. In addition, the two New-Keynesian Phillips curves have the same expressions for the map from aggregate consumption to overall inflation.

Proof of Proposition 1.

We solve analytically the model which features non-homothetic preferences with a representative agent who is attentive. Operationally, we set $\alpha^N = \alpha^E = 0$ as we have one agent only. We set $\lambda = 1$. We as have only one agent, we have $C_{H,t} = C_t$ and similarly for sector specific variables and employment variables. We can rewrite the first set of equilibrium conditions:

$$\begin{aligned}
\hat{p}_t^N &= \frac{1}{\gamma^E}\hat{C}_t^E - \frac{1}{\gamma^N}\hat{C}_t^N \\
\hat{N}_t + \frac{1}{\gamma^E}\hat{C}_t^E &= \hat{w}_t \\
\hat{Y}_t^N &= \hat{N}_t^N \\
\hat{\mathcal{S}}_t^N &= \hat{w}_t - \hat{p}_t^N \\
\hat{Y}_t^E &= \hat{N}_t^E \\
\hat{\mathcal{S}}_t^E &= \hat{w}_t \\
\hat{Y}_t^N &= \hat{C}_t^N \\
\hat{Y}_t^E &= \hat{C}_t^E \\
\hat{C}_t &= (1-\bar{C}^N)\hat{C}_t^E + \bar{C}^N\hat{C}_t^N \\
\hat{N}_t &= (1-\bar{C}^N)\hat{N}_t^E + \bar{C}^N\hat{N}_t^N
\end{aligned}$$

We can solve this systems to express $\hat{\mathcal{S}}_t^N$ and $\hat{\mathcal{S}}_t^N$ as function of $\hat{C}_t$ and $\hat{p}_t^N$:

$$\begin{bmatrix} \hat{\mathcal{S}}_t^E \\ \hat{\mathcal{S}}_t^N \end{bmatrix} = \begin{bmatrix} \frac{1+\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} & \frac{\gamma^N\bar{C}^N}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} \\ \frac{1+\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} & \frac{\gamma^E(1-\bar{C}^N)}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} \end{bmatrix} \begin{bmatrix} \hat{C}_t \\ \hat{p}_t^N \end{bmatrix}$$

Compactly:

$$\begin{bmatrix} \hat{S}_t^E \\ \hat{S}_t^N \end{bmatrix} = \begin{bmatrix} a_C^{SE} & a_p^{SE} \\ a_C^{SN} & a_p^{SN} \end{bmatrix} \begin{bmatrix} \hat{C}_t \\ \hat{p}_t^N \end{bmatrix}$$

Next, we map goods specific consumption and inflation to their aggregate counterparts. First, express consumption of essentials as a function of overall consumption and relative prices with the overall consumption definition and the intra-termprial consumption good choice.

$$\hat{C}_t = (1-\bar{C}^N)\hat{C}_t^E + \bar{C}^N\hat{C}_t^N$$
$$\hat{C}_t = (\gamma^E(1-\bar{C}^N) + \gamma^N\bar{C}^N)\frac{1}{\gamma^E}\hat{C}_t^E - \gamma^N\bar{C}^N\hat{p}_t^N$$

We can express inflation in essential and non-essential as function of overall inflation and relative prices with the mapping between relative prices and the inflation rates:

$$\hat{\pi}_t = (1-\bar{C}^N)\hat{\pi}_t^E + \bar{C}^N\hat{\pi}_t^N$$
$$\hat{\pi}_t^N = \hat{\pi}_t + (1-\bar{C}^N)(\hat{p}_t^N - \hat{p}_{t-1}^N)$$

and symmetrically:

$$\hat{\pi}_t^E = \hat{\pi}_t - \bar{C}^N(\hat{p}_t^N - \hat{p}_{t-1}^N)$$

We can now turn to the inter-temporal part of the model. The equations are:

$$\hat{\pi}_t^E = \beta\mathbb{E}_t(\hat{\pi}_{t+1}^E) + \kappa\hat{\mathcal{S}}_t^E$$
$$\hat{\pi}_t^N = \beta\mathbb{E}_t(\hat{\pi}_{t+1}^N) + \kappa\hat{\mathcal{S}}_t^N$$
$$\frac{1}{\gamma^E}\mathbb{E}_t\left(\hat{C}_{t+1}^E\right) = \frac{1}{\gamma^E}\hat{C}_t^E - \mathbb{E}_t(\hat{\pi}_{t+1}^E) + \hat{R}_t$$

Substitute-in the mappings from inflation in essential and non-essentials and essential consumption to overall consumption, inflation, and relative prices.

$$\hat{\pi}_t - \bar{C}^N(\hat{p}_t^N - \hat{p}_{t-1}^N) = \beta\mathbb{E}_t(\hat{\pi}_{t+1}) - \beta\bar{C}^N(\mathbb{E}_t(\hat{p}_{t+1}^N) - \hat{p}_t^N) + \kappa\hat{\mathcal{S}}_t^E$$
$$\hat{\pi}_t + (1-\bar{C}^N)(\hat{p}_t^N - \hat{p}_{t-1}^N) = \beta\mathbb{E}_t(\hat{\pi}_{t+1}) + \beta(1-\bar{C}^N)(\mathbb{E}_t(\hat{p}_{t+1}^N) - \hat{p}_t^N) + \kappa\hat{\mathcal{S}}_t^N$$
$$\frac{1}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N}\mathbb{E}_t\left(\hat{C}_{t+1}\right) + \frac{\bar{C}^N(1-\bar{C}^N)(\gamma^N-\gamma^E)}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N}\mathbb{E}_t\left(\hat{p}_{t+1}^N\right) =$$
$$= \frac{1}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N}\hat{C}_t + \frac{\bar{C}^N(1-\bar{C}^N)(\gamma^N-\gamma^E)}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N}\hat{p}_t^N - \mathbb{E}_t(\hat{\pi}_{t+1}) + \hat{R}_t$$

We can substitute in a simplified Taylor rule: $\hat{R}_t = \phi_\pi\mathbb{E}(\pi_{t+1}) + \varepsilon_t^{mp}$ and the expressions that map responses of consumption and relative prices to marginal costs and write the system in matrix form. In the final system the only parameter or convolutions that matter are: γ^E,

γ^N, β, κ, $\bar{C}^N$.

$$\begin{bmatrix} 0 & \beta & -\beta\bar{C}^N \\ 0 & \beta & \beta(1-\bar{C}^N) \\ \frac{1}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} & \phi_\pi - 1 & \frac{\bar{C}^N(1-\bar{C}^N)(\gamma^N-\gamma^E)}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} \end{bmatrix} \begin{bmatrix} \mathbb{E}_t(\hat{C}_{t+1}) \\ \mathbb{E}_t(\hat{\pi}_{t+1}) \\ \mathbb{E}_t(\hat{p}^N_{t+1}) \end{bmatrix} +$$

$$\begin{bmatrix} \kappa\frac{1+\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} & -1 & \bar{C}^N(\beta+1)+\kappa\frac{\bar{C}^N\gamma^N}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} \\ \kappa\frac{1+\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} & -1 & -(1-\bar{C}^N)(\beta+1)-\kappa\frac{(1-\bar{C}^N)\gamma^E}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} \\ -\frac{1}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} & 0 & -\frac{\bar{C}^N(1-\bar{C}^N)(\gamma^N-\gamma^E)}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} \end{bmatrix} \begin{bmatrix} \hat{C}_t \\ \hat{\pi}_t \\ \hat{p}^N_t \end{bmatrix} +$$

$$\begin{bmatrix} 0 & 0 & -\bar{C}^N \\ 0 & 0 & (1-\bar{C}^N) \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{C}_{t-1} \\ \hat{\pi}_{t-1} \\ \hat{p}^N_{t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix} \varepsilon^{mp}_t = 0$$

$$A\mathbb{E}(X_{t+1}) + BX_t + CX_{t-1} + H\varepsilon^{mp}_t = 0$$

We solve this system in the case of iid monetary policy shock. We solve it with the undetermined coefficient method. The solution depends on the monetary policy shock and on the state variable, the relative price in the previous period $\hat{p}^N_{t-1}$:

$$\begin{bmatrix} \hat{C}_t \\ \hat{\pi}_t \\ \hat{p}^N_t \end{bmatrix} = \begin{bmatrix} e_1\hat{p}^N_{t-1} + d_1\varepsilon^{mp}_t \\ e_2\hat{p}^N_{t-1} + d_2\varepsilon^{mp}_t \\ e_3\hat{p}^N_{t-1} + d_3\varepsilon^{mp}_t \end{bmatrix} = \begin{bmatrix} e_1 & d_1 \\ e_2 & d_2 \\ e_3 & d_3 \end{bmatrix} \begin{bmatrix} \hat{p}^N_{t-1} \\ \varepsilon^{mp}_t \end{bmatrix}$$

The system with the solution plugged in becomes:

$$A\mathbb{E}(X_{t+1}) + BX_t + CX_{t-1} + H\varepsilon^{mp}_t = 0$$

$$A \begin{bmatrix} e_1(e_3\hat{p}^N_{t-1} + d_3\varepsilon^{mp}_t) \\ e_2(e_3\hat{p}^N_{t-1} + d_3\varepsilon^{mp}_t) \\ e_3(e_3\hat{p}^N_{t-1} + d_3\varepsilon^{mp}_t) \end{bmatrix} + B \begin{bmatrix} e_1\hat{p}^N_{t-1} + d_1\varepsilon^{mp}_t \\ e_2\hat{p}^N_{t-1} + d_2\varepsilon^{mp}_t \\ e_3\hat{p}^N_{t-1} + d_3\varepsilon^{mp}_t \end{bmatrix} + C \begin{bmatrix} 0 \\ 0 \\ \hat{p}^N_{t-1} \end{bmatrix} + H\varepsilon^{mp}_t = 0$$

This creates two sets of systems of equations to solve for, from the coefficients associated with the state variable and with the monetary policy shock:

$$Ae_3 \begin{bmatrix} e_1 \\ e_2 \\ e_3 \end{bmatrix} + B \begin{bmatrix} e_1 \\ e_2 \\ e_3 \end{bmatrix} + C \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = 0$$

$$Ad_3 \begin{bmatrix} e_1 \\ e_2 \\ e_3 \end{bmatrix} + B \begin{bmatrix} d_1 \\ d_2 \\ d_3 \end{bmatrix} + H = 0$$

This would be daunting to solve analytically if monetary policy affected the relative price d_3. However, we show that the solution has $d_3 = 0$ by guessing it and verifying it. The uniqueness of the solution is guaranteed by the Taylor principle $\phi_\pi > 1$. The key idea is that the responses of consumption and inflation to the monetary policy shock depend on the average IES only $\gamma^E(1-\bar{C}^N) + \gamma^N \bar{C}^N$ and not on its elements separately. Moreover, the two NKPC display the same terms for inflation and consumption. If this was not the case, say due to labour market heterogeneity or price stickiness heterogeneity, the proof would not go through, showing that non-homotheticity matters only in conjunction with other relevant heterogeneity for aggregate fluctuation.

Guess $d_3 = 0$, then:

$$A0 \begin{bmatrix} e_1 \\ e_2 \\ e_3 \end{bmatrix} + B \begin{bmatrix} d_1 \\ d_2 \\ 0 \end{bmatrix} + H = 0$$

$$B \begin{bmatrix} d_1 \\ d_2 \\ 0 \end{bmatrix} + H = 0$$

$$\begin{bmatrix} \kappa \frac{1+\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} d_1 - d_2 = 0 \\ \kappa \frac{1+\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} d_1 - d_2 = 0 \\ -\frac{1}{\gamma^E(1-\bar{C}^N)+\gamma^N\bar{C}^N} d_1 - 1 = 0 \end{bmatrix}$$

$$d_1 = -(\gamma^E(1-\bar{C}^N) + \gamma^N\bar{C}^N)$$

$$d_2 = -\kappa(1 + \gamma^E(1-\bar{C}^N) + \gamma^N\bar{C}^N)$$

That is consumption responds by the average IES to a monetary policy shock and inflation responds by the Phillips curve slope times by one plus the average IES. This concludes the proof that in a non-homothetic RANK, only the average IES matters for aggregate fluctuations. This concludes the proof. ■

We now move to the corollary: the non-homothetic RANK responses of aggregate variables to monetary policy are the same to a homothetic-RANK with the same average IES.

Proof of Corollary 1. This is immediate, substitute $\gamma = \gamma^E(1-\bar{C}^N) + \gamma^N\bar{C}^N$ for γ^E

and γ^N. The system becomes:

$$
\begin{bmatrix} 0 & \beta & -\beta\bar{C}^N \\ 0 & \beta & \beta(1-\bar{C}^N) \\ \frac{1}{\gamma} & \phi_\pi - 1 & 0 \end{bmatrix} \begin{bmatrix} \mathbb{E}_t(\hat{C}_{t+1}) \\ \mathbb{E}_t(\hat{\pi}_{t+1}) \\ \mathbb{E}_t(\hat{p}^N_{t+1}) \end{bmatrix} +
$$
$$
\begin{bmatrix} \kappa\frac{1+\gamma}{\gamma} & -1 & \bar{C}^N(\beta+1+\kappa) \\ \kappa\frac{1+\gamma}{\gamma} & -1 & -(1-\bar{C}^N)(\beta+1+\kappa) \\ -\frac{1}{\gamma} & 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{C}_t \\ \hat{\pi}_t \\ \hat{p}^N_t \end{bmatrix} +
$$
$$
\begin{bmatrix} 0 & 0 & -\bar{C}^N \\ 0 & 0 & (1-\bar{C}^N) \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{C}_{t-1} \\ \hat{\pi}_{t-1} \\ \hat{p}^N_{t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix} \varepsilon^{mp}_t = 0
$$

The proof goes through in the same way, with the solution to a monetary policy shock being:

$$
\begin{aligned}
d_1 &= -\gamma \\
d_2 &= -\kappa(1+\gamma) \\
d_3 &= 0
\end{aligned}
$$

This concludes the proof. ■

Notice that the same result would go through also with more complicated models, as long as non-homotheticity does not interact directly with other heterogeneity. It would go through with inattentiveness, sticky wages, or persistent monetary policy. We showed this numerically in Table E.1.