

NBER WORKING PAPER SERIES

ARE INFLATIONARY SHOCKS REGRESSIVE? A FEASIBLE SET APPROACH

Felipe N. Del Canto
John R. Grigsby
Eric Qian
Conor Walsh

Working Paper 31124
<http://www.nber.org/papers/w31124>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2023

Grigsby acknowledges financial support from Princeton University. Walsh acknowledges financial support from Columbia Business School and the National Science Foundation. The authors have no further sources of funding nor material or relevant financial relationships. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Felipe N. Del Canto, John R. Grigsby, Eric Qian, and Conor Walsh. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Are Inflationary Shocks Regressive? A Feasible Set Approach
Felipe N. Del Canto, John R. Grigsby, Eric Qian, and Conor Walsh
NBER Working Paper No. 31124
April 2023
JEL No. E2,E30,E50

ABSTRACT

We develop a framework to measure the welfare impact of inflationary shocks throughout the distribution. The first-order impact of a shock is summarized by the induced movements in agents' feasible sets: their budget constraint and borrowing constraints. To measure this impact, we combine estimated impulse response functions with micro-data on household consumption bundles, asset holdings and labor income for different US households. We find that inflationary oil shocks are regressive, but monetary expansions are progressive, and there is substantial heterogeneity throughout the life cycle. In both cases, the dominant channel is the effect of the shock on asset prices, not movements in goods prices or labor income.

Felipe N. Del Canto
Columbia Business School
350 Kravis Hall
665 W 130th St
New York, NY 10027
felipe.delcanto@icloud.com

John R. Grigsby
Department of Economics and
School of Public and International Affairs
Princeton University
Princeton, NJ 08544
and NBER
jg6005@princeton.edu

Eric Qian
Princeton University
ericqian@princeton.edu

Conor Walsh
Columbia Business School
350 Kravis Hall
665 W 130th St
New York, NY 10027
caw2226@columbia.edu

1. INTRODUCTION

The recent inflationary episode has renewed interest in understanding the distributional incidence of inflationary macroeconomic shocks. Whether inflation is regressive may appear a simple question, but confronting it requires overcoming two challenges. First, the distributional consequences of inflation may depend on the source of the inflationary shock: supply shocks, such as oil price movements, may have a different effect than aggregate demand shocks, such as monetary expansions. Second, inflation affects all parts of the budget constraint: consumption prices, asset prices, transfer income and labor income. Inflationary shocks might have a regressive effect if poor households disproportionately consume goods that are responsive to aggregate inflation. On the other hand, inflation might be progressive if it erodes the real value of nominal debt, which the poor disproportionately owe, or if wages rise more at the bottom of the distribution than the top.

This paper studies the first-order impact of inflationary shocks on heterogeneous households. We develop a new empirical framework which allows for the incidence of inflation to differ based on the source of the shock and accounts for movements in all pieces of the budget constraint. The framework allows households to have different preferences over consumption goods, assets, and labor supply, and for these preferences to evolve as they age, permitting rich heterogeneity in consumption and asset holdings both cross-sectionally and over the life cycle. We also consider additional constraints on the household, such as borrowing or net worth constraints, and hence term our approach a *feasible set* approach.

We show that the first-order impact of a macroeconomic shock on a household's well-being is summarized by the shock's effect on 1) the price of the goods the household purchases, 2) the wage income the household earns, 3) the dividend stream on the assets owned by the household, 4) the prices of assets that the household trades and 5) transfer income from the government. Crucially, the envelope theorem implies that all consumption, labor supply, and asset holding choices are measured at what the household would have chosen absent the shock. While households may wish to substitute away from higher priced goods in response to the shock, such substitutions are not welfare-relevant to a first-order. This holds without needing to specify the general equilibrium structure of the economy that leads to these price responses and is robust to allowing for unemployment, durable consumption goods and assets in the utility function. Our methodology applies for generic stationary macroeconomic shocks, so long as the shock does not directly shift household preferences, as is the case for a wide set of macroeconomic shocks (e.g. monetary shocks, oil price shocks, fiscal policy shocks, TFP shocks, exchange rate shocks, etc.).

Our methodology requires two measurable inputs. First, we require empirical impulse response functions (IRFs), which may be estimated using standard time series techniques. Second, we aggregate these IRFs into welfare movements for different household types using cross-sectional data on consumption patterns, labor income and asset holdings of different household groups, the likes of which are readily available from household surveys.

We apply the framework to study two inflationary shocks which appear important in recent periods: oil supply and monetary shocks. Using “internal instrument” Structural Vector Autoregression (SVAR) techniques, we estimate impulse response functions of disaggregated CPI price indices, asset price and dividend indices and labor income series to the oil supply news shocks of Känzig (2021) and monetary shocks from Gertler and Karadi (2015). We then combine these IRFs with US survey data on consumption, labor income, asset holdings and accumulation patterns over the life cycle for three education groups.

Our main result is that the different sources of inflation carry radically different distributional consequences. Oil supply contractions appear regressive, while nominal interest rate cuts are progressive. After a one standard deviation oil price increase, those with less than a high school education must be paid around 800 dollars (around 2% of annual consumption) in 2019 to be able to afford their pre-shock level of utility. Meanwhile, middle-aged college educated households *gain* the equivalent of 833 dollars (1.1% of annual consumption) from the oil price increase. In contrast, a decrease in nominal rates of 25 basis points – which generates a similar response of aggregate inflation as our oil price shock – has little effect on low-education households but middle-aged high-education households lose around \$4,000 dollars (5.5% of annual consumption). Thus the answer to the question “is inflation regressive?” depends crucially on the source of the inflationary shock.

The difference between oil supply and monetary shocks is primarily explained by the different effects the two shocks have on asset prices. Consistent with Känzig (2021), we estimate that oil supply contractions lead to substantial declines in equity prices, but limited impact on the prices of other assets such as bonds or housing. This primarily benefits those who would have accumulated equities absent the shock, specifically middle-aged households with a college education, as they can now acquire equities more cheaply. This force causes oil price shocks to be highly regressive, even though dividends payouts fall in response to the shock. Monetary expansions have the opposite effect on asset prices: rate cuts raise the price of equities, housing and bonds. This hurts those who would be accumulating such assets, who are primarily middle-aged households, especially those with a college education. The response of assets pushes for inflation driven by monetary policy shocks to be somewhat progressive, as argued by Doepke and Schneider (2006) (but for different reasons).

We estimate that oil price increases push up unemployment and reduce the weekly earnings of the employed. While these effects are quite small, they are strongest amongst low-education households. Expansionary monetary policy has the opposite effect: it reduces the unemployment rate and increases weekly earnings, especially for the lowest-education households.

On the consumption side, both oil supply and monetary shocks induce disproportionate inflation in motor fuel and fuel and utilities. This force pushes towards regressive inflation, since low-education households spend a larger share of their income on these goods. However, its effects are similar for both monetary and oil price shocks.

Overall, our paper makes three contributions. The first is conceptual: the source of inflationary shocks matters for inflation’s distributional consequences. The second is methodological: we

demonstrate how one can measure the distributional welfare consequences of generic macroeconomic shocks by combining cross-sectional data on household consumption, labor income and asset holdings with impulse response functions estimated using standard time series techniques. The final is empirical: expansionary monetary policy is progressive, while oil supply contractions are regressive.

Ours is not the first paper seeking to assess the distributional impact of inflation. Several study the impact of inflation shocks on individual pieces of the budget constraint. Doepke and Schneider (2006) considers whether rich or poor households lose from aggregate inflation by examining heterogeneity in households' net nominal positions: whether the household is a net creditor or debtor. They argue that the losers from inflation are rich, old households who are large nominal creditors, while young, middle-class households with fixed rate mortgages are the main winners. Recently, Orchard (2022) studies cyclical variation in inflation rates by income level, finding that low-income households' experience higher consumption price inflation during recessions than do high-income households. Fang, Liu, and Roussanov (2022) show that stock returns are negatively correlated with core inflation, meaning holding stocks offers little scope to hedge against inflation risk. Our framework clarifies that the effect on asset holdings and consumption prices are both just one margin by which inflation can be redistributive; we argue that one must consider all sides of the budget constraint to fully assess the welfare impact of inflation. What's more, these conclusions may depend on the underlying drive of business cycles and inflation.

In addition, Bartscher, Kuhn, Schularick, and Wachtel (2021) finds that accommodative monetary policy increases employment more for black households than for white households, but widens wealth inequality by increasing the prices of assets held by white households. Broer, Kramer, and Mitman (2022) estimates the effect of European monetary shocks throughout the permanent income distribution using German administrative data and finds a stronger positive response of labor income to monetary expansions for low-income households. Coglianesi, Olsson, and Patterson (2022) studies the sudden tightening of monetary policy in Sweden in 2010-11 and finds that unemployment increases were concentrated among lower-wage workers with more rigid wages. Amberg, Jansson, Klein, and Rogantini Picco (2021) finds a U-shaped relationship between monetary policy and income gains. Lee, Macaluso, and Schwartzman (2022) study the effect of monetary shocks on real income volatility of black and white households. McKay and Wolf (2023c) argue that consumption responses to monetary easing are relatively homogeneous across households. We contribute to this literature by combining the effect of monetary shocks on all sides of the budget constraint into one welfare calculation, which has a few benefits over simply considering consumption responses. First, households may derive utility from things other than consumption, such as leisure or asset holdings. Second, the consumption, labor income and asset channels interact in interesting ways. Third, it can be difficult to directly measure the response of *lifetime* consumption to shocks, either because accounting for non-homothetic utility functions is difficult or because short-run responses of asset prices could affect long-run consumption several years down the road. Addressing these issues typically requires explicit functional form assumptions on utility, which

our framework does not require.

In a related paper, Cardoso, Ferreira, Leiva, Nuño, Ortiz, Rodrigo, and Vazquez (2022) use survey and administrative data from Spain to assess the distributional consequences of inflation. They argue that inflation has three effects on wealth inequality: a redistribution from net nominal lenders to net nominal borrowers (the “Fisher Channel”), an impact on labor income, and a heterogeneous impact on consumption prices. Furthermore, they argue that the Fisher Channel and labor income channels are an order of magnitude larger than the impact on consumption prices. We complement this paper in three ways. First, we study money metric welfare movements, rather than wealth inequality. Second, we combine cross-sectional and time series data to estimate the response to identified shocks, which allows us to assess whether different shocks carry different impacts. Finally, we allow greater flexibility in labor income and consumption bundles.

While we focus on the effect of short-run inflationary shocks, a recent literature studies differences in long-run consumption price inflation rates across households. Jaravel (2019), Kaplan and Schulhofer-Wohl (2017) and Argente and Lee (2021) use administrative pricing data to show that inflation rates have been lower for high-income households than low-income households on average.¹

This large reduced form literature is complemented by a set of papers which fully specify a structural model and use it to study the distributional effects of shocks. Auclert (2019) studies the role of redistribution for the aggregate effects of monetary policy in a heterogeneous agent New Keynesian (HANK) model.² He finds that those who gain from monetary policy are those with high marginal propensities to consume (MPCs), primarily due to the Fisher Channel and labor income channel. Yang (2022) studies optimal monetary policy rules in a HANK model when monetary shocks affect all sides of the budget constraint for different households differently. Glover, Heathcote, Krueger, and Ríos-Rull (2020) studies how the Great Recession redistributed income across generations in a general equilibrium stochastic overlapping-generations (OLG) model in which aggregate shocks affect asset valuations and income across the distribution. Gagliardone and Gertler (2023) study a New Keynesian model augmented to account for oil supply, and find a large role for oil supply shocks in accounting for the recent severe inflationary episode. Erosa and Ventura (2002) model the role of deficit-financing inflation as a regressive income tax on the poor due to their higher propensity to use cash in transactions. This paper contributes reduced form evidence for the redistributive effect of monetary policy and oil supply shocks as well as the important role that asset price responses play in this redistribution.

Our paper is also related to a host of recent papers seeking to measure the welfare effects of price movements. Baqaee and Burstein (forthcoming) considers how welfare responds to changes in budget sets or technologies with taste shocks and non-homothetic preferences in a representative agent economy. Baqaee and Burstein (2022) extends this to a heterogeneous

¹Jaravel (2021) surveys this literature and the associated policy implications.

²Zamourani (2021) measures the unhedged interest rate exposure channel from Auclert (2019) in the Eurozone.

agent economy. Davila and Schaab (2022) considers how to make welfare assessments in heterogeneous agent economies and shows that one can decompose welfare effects of policies into four components: aggregate efficiency, risk-sharing, intertemporal-sharing, and redistribution. Baqaee, Burstein, and Koike-Mori (2022) and Jaravel and Lashkari (2022) both provide methods to estimate the long-run changes in consumer money-metric welfare over time, allowing for arbitrary non-homothetic preferences. Oberfield (2023) shows that inequality in measured inflation need not reflect inequality in growth of living standards in a model with learning-by-doing and non-homothetic preferences. These papers are useful for characterizing long-run changes in societal welfare. Our focus is instead on studying the distributional money-metric welfare impact of short-run identified macroeconomic shocks.³

Methodologically, our paper is related to those that rely on the envelope theorem to study distributional effects of price movements. The closest is Fagereng, Gomez, Gouin-Bonenfant, Holm, Moll, and Natvik (2022), who study whether long run changes in asset prices have redistributed resources across the income distribution. Like us, they rely on the envelope theorem to point out that asset price increases increase the money-metric welfare of households that would *sell* assets, but are not welfare-relevant for those that would hold assets. They use high-quality administrative data in Norway to argue that rising asset valuations redistributed welfare from the young towards the old and from the poor towards the wealthy. Kim and Vogel (2020) likewise use the envelope theorem in a static model with no asset accumulation to argue that the welfare effect of expanded trade is summarized by the effect that trade has on the price that households' pay for consumption and their labor income. We contribute to these papers by combining cross-sectional data with the well-identified estimation of impulse response functions to specific stationary macroeconomic shocks in order to assess the distributional effects of inflationary shocks.

The rest of the paper is structured as follows. Section 2 outlines the framework which organizes our empirical exercise. Section 3 describes the data we use for our estimation. Section 4 details our approach to estimating impulse response functions. Section 5 presents estimated impulse response functions to oil supply and monetary shocks. Section 6 applies the framework to study the distributional effects of specific shocks. It first provides an overview of the key cross-sectional moments of consumption, asset accumulation and wage income which underlie our results, then estimates the money-metric welfare effects of oil supply new shocks and monetary policy announcements. Section 7 presents robustness checks and extensions to our framework. Section 8 concludes.

2. FRAMEWORK

This section offers an organizing framework to guide our empirical analysis. We consider agents who differ in their preferences over consumption bundles, labor supply and asset holdings, and who face different prices for their labor. The framework shows how to aggreg-

³Aguiar, Amador, and Arellano (2021) follows a budget set approach to study Pareto-improving fiscal policies when the risk-free rate on government bonds is less than the growth rate.

ate empirical impulse response functions using cross-sectional data to form estimates of the first-order welfare impact of inflationary shocks. We first present a benchmark deterministic version of the framework to build intuition, before introducing a stochastic environment and empirically-relevant extensions.

2.1 A Benchmark Deterministic Economy

Setting. Time is discrete and indexed by t . There are J consumption goods in the economy, indexed by $j \in \{1, \dots, J\}$, with good j having price p_{jt} in period t .

There are $K + 1$ long-lived assets, indexed by $k \in \{0, 1, \dots, K\}$, available for trading in each period. Asset k pays a nominal dividend $\{D_{kt}\}_{t=0}^{\infty}$ and may be traded at a price Q_{kt} in period t . We assume that asset $k = 0$ is a one period nominal bond which pays a return $R_t = 1/Q_{0t}$. Let $R_{t \rightarrow T} \equiv \prod_{\tau=t}^T R_{\tau}$ represent the cumulative return between periods t and T of buying a sequence of such bonds. Asset $k = 1$, which we term “money,” serves as the numeraire in this economy, pays a zero dividend forever and is completely durable.

The economy is populated by a finite set of G different household types with overlapping generations. Let a denote the age of a household at some reference time $t = 0$, which we call the household’s “initial age.” A household type is determined by a combination of their initial age a and their group g . Within an initial age a and group g , households are identical, and earn labor income W_t^{ag} . They die at group- and age-dependent rates, and we denote the cumulative survival rate of a cohort of initial age a by time t as δ_t^{ag} . Note that this nests both the canonical infinitely lived household with $\delta_t^{ag} = 1$, constant death rates, and finitely-lived overlapping generations structures with realistic death probabilities.

Let N_{kt}^{ag} denote the amount of asset k held by group g of initial age a at time t , where a negative value for N_{kt}^{ag} represents borrowing. We let Δ represent the first difference operator so that $\Delta X_t^{ag} \equiv X_t^{ag} - X_{t-1}^{ag}$. Assets are subject to convex adjustment costs $\chi_k(\Delta N_{kt})$, which gives one motive for holding assets of different types which is not purely the deterministic return. The one-period bond is not subject to adjustment costs.⁴

Let T_t^{ag} denote government transfers (or taxes, if negative) to households of group g and initial age a in period t .

Households have time-separable preferences with subjective discount factor $\beta \in (0, 1)$. The household has preferences from consumption, labor and asset holdings. We assume that each household type derives utility from an implicitly-defined aggregator of goods

$$(1) \quad C^{ag}(\{c_{jt}^{ag}\}_{j=1}^J, C_t^{ag}) = 1$$

where c_{jt}^{ag} is the consumption of good j chosen in period t by household g that is of initial age a . We assume that $C^{ag}(\cdot)$ is increasing in all but the last argument (where it is decreasing), and

⁴These costs can potentially be group specific (e.g. the poor are excluded from trading in stocks).

continuously differentiable.⁵ Let P_t^{ag} be the ideal price index over consumption for group g of initial age a at time t , defined as $P_t^{ag} \equiv E(\{p_{jt}\}_j, C_t^{ag})/C_t^{ag}$, where $E(\cdot)$ is the household's expenditure function.

Households' preferences may be summarized by the differentiable utility function $U(C_t^{ag}, \{N_{kt}^{ag}\}_{k=1}^K, L_t^{ag})$, where L_t^{ag} is the labor supplied by households of initial age a at time t . We assume that $U(\cdot)$ is weakly increasing and concave in its first two arguments, and weakly decreasing and convex in labor. Note we assume that bonds do not enter the utility function, but money or other assets might.⁶

The representative type g household of initial age a takes as given its initial stock of asset holdings $\{N_{k0}\}_k$ and the path of prices, wages, dividends and transfers, which we collect into a vector $\mathbf{s} \equiv \{\{p_{jt}\}_j, W_t, \{D_{kt}, Q_{kt}\}_k, T_t\}_t$. It solves the following present value utility maximization problem

$$(2) \quad V^{ag}(\mathbf{s}, \{N_{k0}\}_{k=1}^K) = \max_{\{\{c_{jt}^{ag}\}_j, L_t^{ag}, \{N_{kt}^{ag}\}_k\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \delta_t^{ag} U(C_t^{ag}, \{N_{kt}^{ag}\}_{k=1}^K, L_t^{ag}),$$

subject to period-by-period budget constraints for all t ,

$$(3) \quad \sum_j p_{jt} c_{jt}^{ag} = \sum_k \left[N_{kt-1}^{ag} D_{kt} - Q_{kt} (\Delta N_{kt}^{ag}) - \chi_k (\Delta N_{kt}^{ag}) \right] + W_t^{ag} L_t^{ag} + T_t^{ag},$$

the consumption aggregator (1), and a series of no-Ponzi conditions

$$(4) \quad \lim_{T \rightarrow \infty} R_{0 \rightarrow T}^{-1} N_{kT}^{ag} Q_{kT} \geq 0.$$

Welfare Response to Shocks. Consider some arbitrary perturbation dz which induces a change of the household's state vector, but leaves unchanged their preferences $U(\cdot)$, consumption aggregator \mathcal{C} , discount rates and survival rates. For example, dz could arise from a monetary policy shock, a shock to government transfers, an oil price shock, a technology shock which shifts wages, or some combination of all of these at different horizons. Each of these shocks may induce changes in the price of labor, consumption goods or assets, or changes in dividend streams. With a slight abuse of language, we refer to dz as a "shock" to be in keeping with the stochastic version of the model described in section 2.2 below. To save on notation, we drop the explicit dependence on g below, but all analysis should be viewed as group-specific.

Denote by dV^a/dz the first order impact of the perturbation dz on the welfare of a household of age a defined in equation (2). Following Fagereng et al. (2022), we define the *money metric welfare gain* from the shock dz as its welfare gain scaled by the households marginal utility of

⁵For example, a non-homothetic constant elasticity of substitution (CES) aggregator has $\mathcal{C}(\{c_{jt}^{ag}\}_{j=1}^J, C_t^{ag}) \equiv \sum_j (c_{jt}^{ag} (C_t^{ag})^{\epsilon_j})^{\frac{\sigma-1}{\sigma}} = 1$. This is equivalent to homothetic CES if $\epsilon_j = -1$ for all j .

⁶Allowing assets to directly impact utility is a common tool in monetary and financial economics to capture the liquidity values of assets (Sidrauski, 1967; Van den Heuvel, 2008). Indeed, one may consider bequest motives as being a form of assets directly affecting utility, as individuals receive utility from gifting assets to their descendants.

dollar of consumption at time zero, λ_0 :

$$\text{Money Metric Welfare Gain} \equiv \frac{dV^a}{dz} \Big/ \lambda_0$$

It may be interpreted as the household's willingness to pay to receive the shock measured in time 0 dollars, or an equivalent variation welfare measure. It is a particularly useful measure for considering the incidence of shocks across the distribution, since it collects welfare effects in a common unit—dollars—that is readily interpretable even when households have very different utility functions. While direct interpersonal comparisons of utility are generally not possible, this measure allows policymakers to assess distributional consequences in the same units in which inter-household transfers are made. We now state the key proposition of the deterministic framework.

Proposition 1. *Consider a household who chooses a sequence $\{\{c_{jt}^a\}_j, L_t^a, \{N_{kt}^a\}_k\}_t$. To a first order, the money metric welfare gain for the household in response to a small perturbation dz is*

$$\begin{aligned} \left(\frac{dV^a}{dz} \right) \Big/ \lambda_0 = & \sum_{t=0}^{\infty} R_{0 \rightarrow t}^{-1} \left(\underbrace{\sum_j -p_{jt} c_{jt}^a \cdot \frac{d \ln p_{jt}}{dz}}_{\text{Consumption Price Changes}} + \underbrace{W_t^a L_t^a \frac{d \ln W_t^a}{dz}}_{\text{Wage Income Changes}} \right. \\ & \left. + \underbrace{\sum_k D_{k,t} N_{kt-1}^a \cdot \frac{d \ln D_{kt}}{dz}}_{\text{Asset Income changes}} - \underbrace{\sum_k Q_{kt} \Delta N_{kt-1}^a \cdot \frac{d \ln Q_{kt}}{dz}}_{\text{Asset Price changes}} + \underbrace{T_t^a \frac{d \ln T_t^a}{dz}}_{\text{Gov't Transfer Changes}} \right) \end{aligned}$$

Proposition 1 states that the money metric welfare gain in response to an arbitrary shock dz is equal to the discounted sum of five terms. First, the shock may induce changes in the price of the household's consumption bundle. The first order effect of the shock simply weights the percentage change in the price of each good induced by the shock by total spending on each good, and ignores substitution effects. For instance, an increase in the price of food will have a larger effect on households for which food occupies a large share of consumption bundle. We term this the "consumption channel" of the perturbation.

Second, the shock may induce changes in labor income for the household if the wages they face move. For instance, an increase in the price of oil may increase the marginal revenue productivity of households employed in the oil extraction sector, resulting in an increase in their wage. We term this the "labor income channel" of the perturbation.

Third, the shock may change the household's asset income if it changes either the dividend stream paid out by their planned asset holdings or affects the prices at which they trade their assets. Crucially, as shown in Fagereng et al. (2022), one need only consider changes in the prices of assets for households that *would have changed their asset holdings absent the shock*.⁷ A rise in the price of the S&P500 at a certain time is mainly relevant for those at a point in their

⁷This may not be true if the household is subject to borrowing constraints. We consider the quantitative importance of borrowing constraints in Section 7.

lifecycle where they are accumulating stocks (in which case it is welfare negative), or for those selling down their holdings (in which case it is welfare positive). We dub the effect of the perturbation on asset prices and dividends as its “portfolio channel.”

Finally, the shock may shift the present value of taxes owed or transfers paid to the household. We term this the perturbation’s “transfer channel.”

A proof is provided in Appendix A. Intuitively, a perturbation in the path of prices, wages, dividends or transfers faced by the household may induce substitutions away from high-priced goods and time periods. However, the envelope theorem guarantees that these substitutions are not welfare-relevant to a first order. Thus, one need not account for the substitution patterns to measure the first-order welfare effects of price movements: it suffices to simply consider the present discounted value of movements in the household’s budget constraint. If households are on their Euler equation, and do not receive utility from bonds, they discount future movements in prices by the risk free rate.

Proposition 1 forms the foundation of our empirical strategy. It provides a method to appropriately aggregate arbitrary movements in prices into a welfare metric given (i) measured perturbations in households’ feasible sets and (ii) an estimate of the choices a household would make absent said perturbation. This is true without specifying the general equilibrium structure of the economy or the nature of the utility function. Indeed, it allows features other than an aggregate consumption good, such as leisure or asset holdings, to enter utility and thus holds even if consumption is not a sufficient statistic for welfare. Households do not care directly about the nature of the supply side of the economy which determines *why* prices move; they only care *that* prices move. One could employ this proposition to study any episode of price movements: for instance, one could feed in sequences of realized price movements after a particular recession. The two sets of price movements considered below—those associated with oil supply cuts and monetary expansions—appear particularly interesting ones to study, as oil supply and monetary policy are often considered two of the primary drivers of inflation in U.S. history. Next, we adapt the framework to a stochastic environment, and show that the welfare formula is unchanged, except that empirical impulse response functions replace perturbations.

2.2 Macroeconomic Shocks and Welfare

We now move to a stochastic setting to relate changes in household welfare to commonly studied fundamental macroeconomic shocks. We assume that dividends, asset prices, goods prices and wage income are stochastic, and take the form

$$(5) \quad D_{kt} = \bar{D}_{kt} \exp\left(v_{kt}^D\right)^\sigma, \quad Q_{kt} = \bar{Q}_{kt} \exp\left(v_{kt}^Q\right)^\sigma, \quad p_{jt} = \bar{p}_{jt} \exp\left(v_{jt}^p\right)^\sigma, \quad W_t^a = \bar{W}_t^a \exp\left(v_t^{W^a}\right)^\sigma,$$

where $\sigma > 0$ is a parameter that scales the variance of the stochastic processes. These variables depend on a deterministic time component, denoted with a bar (e.g. \bar{D}_{kt}), and a stationary shock process (e.g. v_{kt}^D). We assume that the shock processes are functions of current and

lagged values of a structural shock vector ϵ_t , such that

$$v_{kt}^D = \theta_k^D(L)\epsilon_t, \quad v_{kt}^Q = \theta_k^Q(L)\epsilon_t, \quad v_{jt}^P = \theta_j^P(L)\epsilon_t, \quad v_t^W = \theta^W(L)\epsilon_t,$$

where $\theta(L)$ is a lag operator matrix of finite dimension, and the elements of ϵ_t are mutually uncorrelated. We collect each of these into a vector \mathbf{v}_t . $\theta(L)$ is general in the sense that it nests both persistent and transitory shocks. We further assume that the structural shocks \mathbf{v}_t have no direct effect on household utility functions.⁸

Following Stock and Watson (2018), we define the vector of structural impulse response of each of the variables affecting the household's budget constraint (consumption and asset prices, wages, dividend streams and transfers) at time $t + h$ to the n^{th} entry of the structural shock vector ϵ at time t as

$$\Psi_{n,t+h} \equiv \mathbb{E}_t[\mathbf{v}_{t+h} | \epsilon_t^n = 1] - \mathbb{E}_t[\mathbf{v}_{t+h} | \epsilon_t^n = 0]$$

We are now ready to present the key result that guides our empirical analysis.

Proposition 2. *In the limit as $\sigma \rightarrow 0$, the first-order change in money-metric welfare from an impulse to an element i of the fundamental shock vector at $t = 0$ is*

$$(6) \quad (dV^a) / \lambda_0 = \sum_t R_{0 \rightarrow t}^{-1} \left(\underbrace{-\sum_j p_{j,t} c_{jt}^a \Psi_{n,t}^{p,j}}_{\text{Consumption}} + \underbrace{W_t^a L_t^a \Psi_{n,t+h}^{W^a}}_{\text{Labor Income}} + \underbrace{T_t^a \Psi_{n,t}^T}_{\text{Transfer Income}} \right) + \underbrace{\sum_k \left[N_{kt-1}^a D_{kt} \Psi_{n,t}^{D,k} - Q_{kt} \Delta N_{kt}^a \Psi_{n,t}^{Q,k} \right]}_{\text{Portfolio Channel}}$$

where $\Psi_{n,t}^{p,j}$ is the impulse response of the log price of good j , $\Psi_{n,t}^{W^a}$ is the impulse response of age a 's log wages, $\Psi_{n,t}^T$ is the response of log transfer income, $\Psi_{n,t}^{D,k}$ and $\Psi_{n,t}^{Q,k}$ are the impulse responses of logged dividends and prices, respectively, for asset k , all to the n^{th} entry of the structural shock vector ϵ .

This result shows how one can aggregate estimated impulse responses Ψ of goods prices, asset prices and income to a shock using cross-sectional information on the individuals spending and lifecycle asset holdings to form an estimate of the welfare effects of shocks. Crucially, these choices and prices are evaluated in steady state, and their movements are discounted by the risk-free rate. A full proof is provided in Appendix A and the intuition is identical to that of Proposition 1. Note again the existence of (i) the consumption channel, (ii) the labor income channel, (iii) the portfolio channel and (iv) the transfer channel.

Discussion Proposition 2 guides our empirical analysis of inflationary shocks as it permits the econometrician to aggregate impulse response functions of very different objects – the price of food and the S&P500, for instance – into a common unit. Our strategy is thus to estimate empirical impulse response functions with respect to identified shocks, then use cross-sectional data to aggregate to first-order welfare effects. Pursuing this strategy overcomes the two

⁸We therefore rule out preference shocks, such as discount rate shocks.

primary challenges in studying whether inflationary shocks are regressive: one can include movements of all aspects of the budget constraint and allow the impulse response functions to differ depending on the source of the shock.

The measure we derive is an approximation in an important sense. The formula is valid if the noise facing the agent is "small", such that their asset holdings primarily reflect the joint influence of differential return structures and adjustment costs, and consumption smoothing across different states of the world is not a central motive for portfolio choice. Changes in risk exposure from macroeconomic shocks are interesting, but beyond the scope of this paper.⁹ Given this, one should interpret our approach as measuring whether a particular inflationary shock makes the household's no-shock choices more or less affordable.

One major benefit of this approach is that it does not necessitate specifying the production side of the economy, nor solving for the general equilibrium of a heterogeneous agent economy. This permits us to incorporate more heterogeneity than is usually tractable in a structural model. Indeed, our framework does not impose any restrictions on the general equilibrium relationship between household choices and price movements. Rather, we seek to *estimate* the general equilibrium effects that shocks exert on prices in the economy. The key assumption underlying this approach is that the household only cares about general equilibrium relationships insofar as they affect the prices the household faces. For example, households do not care whether oil price shocks increase food prices because they increase marginal costs of production or because they induce a monetary policy response: they simply care that food prices increased. This assumption is standard in most macroeconomic models.¹⁰

There are two senses in which this reasoning may be incomplete. The first is on the response of asset prices. We do not impose that assets are in fixed net supply, so that accumulation by one group of households necessarily corresponds to decumulation by another. As a result, the aggregate welfare effects of asset price fluctuations need not equal zero. In practice, this could be because some other agents – such as governments or foreign investors – take the other side of asset trades, or because asset supply is not perfectly inelastic. We do not seek to measure the incidence of these shocks on these other groups: our exercise should be understood as trying to assess the distributional effects of inflationary shocks within the U.S. household sector.

Second, we assume the shock has no effect on households' risk preferences and that households' intertemporal decision making is driven by a rational expectations Euler Equation. If the shock also influences household risk or time preferences, or if it generated behavioral shifts in expectations, our exercise would measure the household's first order money-metric welfare change evaluated under their pre-shock utility function and expectations. We feel this is a reasonable statistic to compute when trying to assess the first-order distributional effects of inflationary shocks.

The framework is limited in a few additional ways. The small noise approximation we con-

⁹We additionally abstract from additional constraints on household decisions in our baseline formulation. We explore robustness to incorporating borrowing and short-selling constraints in Section 7.

¹⁰McKay and Wolf (2023b) make a similar assumption to show that policy shocks inform policy counterfactuals.

sider essentially removes the role of asset choice in insuring the consumer against aggregate shocks. The only reason for holding different asset classes is to trade off differential return structures, the utility flow from assets and the costs of changing the holdings. Any importance the consumer places on hedging the risk of inflation shocks (say, via holding gold or other durable assets) is effectively ignored. Our approach thus estimates whether the household's initial hedging choices become more expensive relative to their income after the shock.

In addition, evaluating the welfare formula in Proposition 2 requires knowledge of the asset accumulation profile in a deterministic setting with no shocks. To approximate this, we use average asset accumulation profiles over the lifetime, and explore the sensitivity of this to including different cross-sections of our data. However, even in a "small shock" world, the simple average of life-cycle data observed in repeated cross-sections may not exactly correspond to the solution in the deterministic setting. In addition, for our result to apply, the value function must be differentiable. It is therefore not applicable in settings with non-convex adjustment costs and kinks in the value function.

Both of these limitations should be noted when interpreting our analysis below. Our focus is on understanding who the first-order winners and losers are from inflationary macroeconomic shocks. While most papers in this space either study wealth inequality or one aspect of the budget constraint in isolation, Proposition 2 provides one method to assess the welfare effects by considering the present value of movements in the budget set. One key implication of our framework is that one must consider movements in prices relative to movements in income. The fact that inflation erodes the nominal value of debt, for instance, only has a welfare effect if income moves commensurately with inflation. Likewise, prices moving more for the goods consumed by the poor only has a disproportionate welfare effect if income for the poor does not keep up with their specific inflation rate. This core insight is not subject to the limitations outlined above. Thus, our approach to measuring the distributional impact of inflation follows the logic of Proposition 2, as we estimate the components of equation (6).

Before discussing our approach to estimating the elements of equation (6), we discuss a number of model extensions that are useful to bring the analysis closer to the data.

2.3 Empirical Extensions

Unemployment. We suppose that only a fraction $(1 - u_t^a)$ of the labor supplied to the market is actually employed by the market, while a fraction u_t^a of labor supply is rationed through unemployment. This is taken to represent a probability of unemployment within household.¹¹ We then suppose that earnings are given by

$$W_t^a = \bar{W}_t^a \exp(v_t^W)^\sigma (1 - u_t^a) L_t^a \quad u_t^a = \bar{u}_t^a \exp(v_t^u)^\sigma$$

Earnings have some deterministic component that is age-dependent and varies over the life-cycle, as well as a shock component that varies the wage depending on the state of the struc-

¹¹Implicitly, we assume perfect insurance within households with many individuals or within groups.

tural shock vector. They are scaled by the fraction of time spent employed, which also carries a shock component which loads on the structural shock vector ϵ_t . Now we interpret L_t^a as incorporating all margins of labor supply chosen by the household: that is, L_t^a incorporates both hours worked and labor force participation. In contrast, u_t^a and W_t^a are taken as given by the household: it does not choose its unemployment rate nor its wage.

A similar envelope argument shows that the welfare response is now

$$(7) \quad (dV^a)/\lambda_0 = \sum_t R_{0 \rightarrow t}^{-1} \left(- \sum_j p_{j,t} c_{jt}^a \Psi_{n,t}^{p,j} + T_t^a \Psi_{n,t}^T \right) + \underbrace{\sum_k \left[N_{kt-1}^a D_{kt} \Psi_{n,t}^{D,k} - Q_{kt} \Delta N_{kt}^a \Psi_{n,t}^{Q,k} \right] + W_t^a L_t^a (1 - u_t^a) \left(\Psi_{n,t}^w + \frac{u_t^a}{1 - u_t^a} \Psi_{n,t}^u \right)}_{\text{Hourly Wage and Unemployment Impulses}}$$

where again each term is understood to be evaluated at the deterministic value, and $\Psi_{n,t}^u$ is the impulse response of log unemployment rates. Comparing with equation (6), the only change is that the labor income channel now has two terms: the response of wages and the response of the unemployment rate. Again via the envelope theorem, endogenous changes in labor supply L_t^a in response to the impulse may be ignored.¹²

Durable Goods – The baseline model presumes all consumption goods fully depreciate between periods. In reality, durable goods act as an important input both to households' consumption bundles and their asset portfolios.

To account for this dual role of durable goods, we assume that the utility-relevant consumption of a durable good j is given by $c_{jt}^a \equiv \varrho_{jt}^a d_{jt}^a$, where d_{jt}^a is household of age a 's stock of the durable good at the beginning of period t and $\varrho_{jt}^a \in [0, 1]$ is the intensity with which the household uses the durable to produce its consumption-relevant good. For instance, if the good j is "vehicles", d_{jt}^a would be the quantity of vehicles owned while ϱ_{jt}^a would be related to the number of miles driven. We assume the household freely chooses the intensity of use ϱ .

Durable goods depreciate with use. In particular, we assume that a fraction $\delta(\varrho)$ of the stock of a durable depreciates between two periods if it is used with intensity ϱ . Under this assumption, one can write the law of motion for the durable as

$$(8) \quad d_{jt+1}^a = (1 - \delta(\varrho_{jt}^a)) d_{jt}^a + I_{jt}^a$$

where I_{jt} is the gross real investment in the durable, which may be negative if the household sells its durable and may be subject to convex adjustment costs $\chi_j(\Delta d_{jt}^a, \varrho_{jt}^a)$. Our treatment of durable goods thus mirrors the usual treatment of capital utilization often considered in the investment literature (Greenwood, Hercowitz, and Huffman, 1988; Burnside, Eichenbaum, and Rebelo, 1995).

To account for durable goods, suppose without loss of generality that consumption goods

¹²Empirically, we further split wages and unemployment rates for the household's primary and secondary earner, estimating a separate impulse response for each, assuming a unitary household.

$j \in \{1, 2, \dots, \hat{j}\}$ are non-durable, while goods $j \in \{\hat{j} + 1, \dots, J\}$ are durable and may be carried across periods. The household's period t budget constraint becomes

$$\underbrace{\sum_{j=1}^{\hat{j}} p_{jt} c_{jt}^a}_{\text{Non-Durable Consumption}} + \underbrace{\sum_{j=\hat{j}+1}^J \left(Q_{jt} I_{jt}^a + \chi_j(I_{jt}^a) \right)}_{\text{Net Durable Purchases}} = W_t^a L_t^a + T_t^a + \underbrace{\sum_k [N_{k,t-1}^a D_{k,t} - Q_{k,t} \Delta N_{k,t}^a - \chi_k(\Delta N_{k,t}^a)]}_{\text{Financial Assets}}$$

where Q_{jt} is the price of durable good j 's purchases. Substituting in for I_{jt} with the law of motion (8) and re-arranging, we have

$$(9) \quad \underbrace{\sum_{j=1}^{\hat{j}} p_{jt} c_{jt}^a}_{\text{Non-Durable Consumption}} + \underbrace{\sum_{j=\hat{j}+1}^J Q_{jt} \delta_j(q_{jt}^a) d_{jt}^a}_{\text{Durable Consumption}} = W_t^a L_t^a + T_t^a + \underbrace{\sum_k [N_{k,t-1}^a D_{k,t} - Q_{k,t} \Delta N_{k,t}^a - \chi_k(\Delta N_{k,t}^a)]}_{\text{Financial Assets}} - \underbrace{\sum_{j=\hat{j}+1}^J \left(Q_{jt} \Delta d_{jt}^a + \chi_j(\Delta d_{jt}^a, q_{jt}^a) \right)}_{\text{Cost of Increasing Durables Stock}}$$

This expression clarifies the dual role of durable goods as an asset and a consumption good. On the expenditure side, consumption of durable goods behaves identically to consumption of non-durables, only with a price proportional to the depreciation and foregone sale price of the durable. Indeed, multiplying and dividing the durable consumption expression by q_{jt}^a , one can see the price of utility-relevant durable consumption is $p_{jt}^a \equiv Q_{jt} \delta_j(q_{jt}^a) / q_{jt}^a$: the dollar value lost to depreciation as a result of usage. On the income side, durables behave identically to a financial asset with zero dividend. Proposition 2 therefore directly applies to the case with durable goods, so long as one can appropriately measure the depreciation rate of the durables in question. We discuss our approach to doing so in Section 3 below.

This formula is particularly useful for clarifying the welfare impacts of inflation through house price changes. Housing is both a durable good that delivers utility and a store of wealth. A commonly encountered view is that for homeowners, rental inflation is irrelevant and rises in house prices are only positive for welfare. This result shows in fact that house price inflation does negatively impact homeowners on the consumption side of the budget constraint, reflecting the increased cost of depreciation from use. This is similar to the "implicit rent" of owning a home. Counterbalancing this consumption channel, a house price increase raises welfare for those planning to decumulate housing through the portfolio channel, as they may sell at a higher price.¹³ The opposite is true for those who accumulate housing. Thus the welfare effect of an unexpected house price increase is more subtle than a clear benefit for homeowners.

¹³Many households also have a mortgage. Adjustments to mortgage interest payments are included as a negative dividend (D_{kt}) for mortgages.

3. DATA

This section describes the data used to estimate the distributional impact of inflationary oil supply and monetary shocks, accounting for movements in all parts of the budget constraint. It describes our data on household consumption, income and portfolios, as well as time series information on prices, dividends and wages. It additionally details the shocks and controls used in our estimation. Appendix B provides further information.

Throughout, household groups g are defined by the educational attainment of the household head. We distinguish households by their educational attainment for three reasons. First, education is a readily available statistic in many datasets. Second, education may often be a better proxy of a household's permanent component of income than their income in any given year. Finally, all surveys that we consider have many observations within the education groups we define, mitigating sampling error concerns. We additionally compute life-cycle profiles of consumption, wages and asset holdings within each education group. We consider households whose head is at least 25 years old, and combine all households between 75 and 80 years old into one group. Our baseline approach is to measure the cross-sectional consumption, portfolio and labor income variables using 2019 data and to assume 2019 represents a steady state. We test robustness of our results to violations of this assumption in Section 7. We use monthly time series to estimate our impulse responses in order to maximize power, but consider a period t to be a quarter, in keeping with our available consumption data.

3.1 Consumption Data

We use monthly consumer price indices published by the Bureau of Labor Statistics (BLS) as our measure of goods prices p_{jt} . The BLS publishes price indices for a variety of goods. Some of these goods have been introduced recently: for instance, the BLS only began tracking the price of "Medical Equipment and Supplies" separately from "Medical Care: Commodities" in 2006. Since we need long time series to estimate our regressions, we only track categories that satisfy three criteria. First, they must be available at least back to 1990. Second, they must represent a sufficiently large share of the aggregate consumption bundle. Finally, they must add up to 100% of consumption. This leaves us with 25 consumption goods, roughly at the level of the BLS' CPI categories.^{14,15}

¹⁴Food at home, Food away from home, Alcoholic beverages, Shelter, Fuels and utilities, Education, Apparel, New Vehicles, Used Vehicles, Other Vehicles, Motor fuel, Public transportation, Personal care, Motor vehicle insurance, Motor vehicle fees, Motor vehicle parts/equipment, Motor vehicle maintenance/repair, Medical care services, Recreation, Medical care commodities, Postage and delivery services, Information and information processing, Information technology, hardware/services, Tobacco and smoking products, and Household furnishings/operations. See Appendix B for details.

¹⁵Prior work has found that households at different income levels experience different trend inflation in consumption prices, and that this difference is driven by differences within fine product groups (Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2019). Producing price indices for such narrow product groups that are suitable for time series analysis is challenging, as high quality data at this level is only recently available. Furthermore, there is no *a priori* reason to think inflation rates of finer product categories should be differentially responsive to short-run shocks. We therefore limit attention to the 25 CPI groups.

We use data from the Consumer Expenditure Survey (CEX) to measure group-specific life cycle consumption of goods. The CEX is a nationally representative quarterly survey run by the Bureau of Labor Statistics (BLS) that provides data on income, expenditures and demographic characteristics of U.S. consumers at the household level. It is principally used to update the relative importance of goods and services in the market basket for the Consumer Price Index (CPI). Its broad coverage of all components of household expenditure makes it uniquely well-suited to our exercise.¹⁶

The CEX is comprised of two surveys. The first is a series of interviews conducted by BLS field economists which ask households about their expenditure in a number of categories over the last quarter. There are 512 different expenditure categories listed in the “Interview” portion of the CEX, and they include both non-durable consumption, rent paid on housing, investment in new durable goods and mortgage interest payments.

The second survey asks households to record their daily expenditures over the course two weeks. This “Diary” survey offers much more granular information about household expenditure, particularly on food products, but lacks information on infrequently-purchased products. The Diary contains information on 284 unique products.

We use microdata from the BLS for the Interview survey from 2019 and group the expenditure categories to 25 groups for which we have a CPI price series, calculating the share of household expenditure on each category.¹⁷ We then calculate the average expenditure of households in group g , weighting each household i by its survey weight, following the BLS procedure for computing representative consumption baskets as closely as possible. These average expenditures form our estimate of $p_{j0}c_{j0}^a$: initial consumption of good j by group g households of initial age a . In our baseline scenario, we assume a constant life cycle profile of consumption of each good, so that $p_{jt}c_{jt}^a = p_{j0}c_{j0}^{a+t}$ absent any shocks. For example, 25 year old households will have the same baseline expenditure on each good j in period $t = 4$ as did 26 year old households in period $t = 0$. That is, different cohorts have the same life cycle path of utility functions and consumption aggregators. Section 7 tests sensitivity of this assumption.

Finally, we calculate depreciation for vehicles using data from the National Household Travel Survey (NHTS). We assume vehicle usage is well approximated by miles driven, and that depreciation is a linear function of mileage. We compute the average annual miles driven by education group in the NHTS and multiply it by the effect of mileage on used car prices estimated by Dexheimer (2003).

¹⁶While other consumption datasets, such as the Nielsen HomeScan dataset or JPMorgan Chase Institute data offer larger sample sizes for household consumption, the CEX remains the only dataset which accounts for all of household’s expenditure.

¹⁷Note that households do not, in general, report healthcare expenditures covered by Medicare or Medicaid. However, inflation in the cost of medical care covered by these services does not affect household well-being as they are covered by the government. Thus we do not include them in our welfare calculation.

3.2 Labor Income Data

We use monthly wage and employment information from the Current Population Survey (CPS). The CPS is a monthly rotating panel, in which households are sampled for four consecutive months, not surveyed for the following 8 months, then sampled again for the following four months. The data contain information on each household member’s employment status, occupation, educational attainment, and demographic information. It is primarily used to construct the civilian unemployment rate and can be linked to form a panel using the procedure of Flood and Pacas (2008). We define households’ group membership by the education and age of the household head. We use the full CPS to construct monthly age \times education group specific unemployment rates, which we consider our estimate of u_t^a in equation (7).

Households are asked about their earnings in the fourth and eighth month of being surveyed in the so-called “Outgoing Rotations Group” (ORG) component. Our benchmark approach is to consider respondents’ weekly earnings. We calculate average life-cycle profiles of weekly earnings by regressing individual wages on age \times education fixed effects in the 2019 CPS, treating household heads and spouses separately. That is, we regress

$$(10) \quad W_{it} = \alpha_{a(i,t),g(i,t)} + \varepsilon_{it}$$

where W_{it} is a measure of the wage (weekly earnings in our baseline), while $a(i, t)$ and $g(i, t)$ indicate the age and education, respectively of individual i in period t . We only run this regression for those who are employed. The fixed effects $\alpha_{a,g}$ constitute our estimate of the life-cycle profile of wage income absent shocks, or $W_0^a L_0^a$ in equation (7). As with consumption data, our baseline scenario assumes a constant life cycle profile of earnings absent any shocks, so that $W_t^a L_t^a = W_0^{a+t} L_0^{a+t}$ for all t and a .

To construct time series wage indices, we first estimate a version of regression (10) in 2019, using log weekly earnings as the outcome variable. Our wage index for group g is the average residual ε_{it} from this regression amongst individuals who are employed both in month t and month $t - 12$. Conditioning on such “job-stayers” reduces the influence of composition effects on wage indices, which is known to be important at high frequencies (Solon, Barsky, and Parker, 1994; Grigsby, 2022).

3.3 Portfolio Data

We use information from the Survey of Consumer Finances (SCF) to calculate the value of both assets and liabilities on household balance sheets by age and education level. The SCF is a triennial nationally representative survey conducted by the Federal Reserve Board to study household balance sheets. Respondents are asked a series of questions about their income, assets, and liabilities, as well as some basic demographic information. We use information on the following balance sheet categories: housing, equity holdings, bond holdings, business wealth, retirement accounts, vehicles, and other financial and non-financial assets. Vehicles

and housing are both treated as durable goods in our calculations. Our baseline sample uses only the 2019 SCF. We additionally include information on mortgage payments from the CEX.

We estimate each group's quarterly accumulation of each asset class k using a synthetic cohort approach. Specifically, we calculate the value of holdings of asset k of group members who are of age a and age $a - 1$. Next, we perform a Locally-Weighted Scatterplot Smoothing (LOWESS) procedure to minimize large swings in asset holdings caused by measurement error. Finally, we approximate $Q_{k0}\Delta N_{k0}^a$ with the estimated change in asset holdings between adjacent ages implied by the LOWESS $Q_{k0}\hat{N}_{k0}^a - Q_{k0}\hat{N}_{k0}^{a-1}$. We again assume a constant life cycle of portfolio so that $Q_{kt}\Delta N_{kt}^a = Q_{k0}\Delta N_{k0}^{a+t}$.

This approach has the added benefit of filtering out movements in the value of assets that arise from price fluctuations. Because we construct implied changes in asset values using life-cycle changes in cross-sectional data, we hold fixed asset prices at the point the survey is administered. This approach may more accurately reflect changes in the *quantity* of asset holdings ΔN_{kt} , which is what is required in our framework.

A sufficient assumption for this approach to be valid is that households born in year t will have the same asset accumulation path as households of the same type born in year $t - 1$. While this assumption is strong, relaxing it would require panel data on asset holdings and trades for various household groups, which is seldom available.¹⁸

The SCF data directly give us the value of asset holdings $Q_{k0}N_{k0}^a$. To recover the no-shock dividend income of each asset class, we use data on dividend yields, which report D_{k0}/Q_{k0} . Multiplying the value of the asset holding from the SCF by the dividend yield returns $D_{k0}N_{k0}^a$ as desired. We assume that dividends do not move for pre-purchased nominal bonds: a fixed coupon will not respond to the economic shock; rather, the asset price will adjust. What's more, we assume that the (money) dividend stream for durable goods is zero: the benefit of durable goods is captured by the durable consumption component as suggested by equation (9). Thus the dividend component is only relevant for equities. We proxy the dividend yield for equities using the dividend yield of the S&P500. Finally, data on mortgage interest payments come from the CEX. We estimate impulse response functions of mortgage payments by education group and include this as part of the portfolio channel.

We use a variety of price indices for our analysis. Equity price returns, estimated dividend yields, and dividend growth are computed from the value-weighted indices (including and excluding dividends) from the Center for Research in Security Prices (CRSP). The S&P CoreLogic Case-Shiller Home Price Index is used to compute house price responses. Interest rates are evaluated using the effective federal funds rate and market Treasury yields of various maturities (1, 2, 3, 5, 7, and 10 years). For corporate bonds, we use Moody's Aaa and Baa corporate bond yields. Bond prices are assumed to be the reciprocal of the yield. We use the 1-year Treasury Yield for discounting purposes, so that $R_{0 \rightarrow t} = (1 + yield_0^{TBill})^{t/4}$.

¹⁸Fagereng et al. (2022) use Norwegian tax records to solve this problem and find that long run asset price increases are regressive.

3.4 Transfer Income

We use data from the Survey of Income and Program Participation (SIPP) to measure transfer income. Transfer income is defined as the sum of means-tested transfer income and social insurance payments. The former component includes payments from the following means-tested programs: TANF, SSI, GA, veterans pension, and pass-through child support. The latter includes other payments from Veterans Affairs, Social Security, unemployment insurance and the G.I. Bill.

Unfortunately, it is difficult to estimate impulse response functions for transfer income by group, since the SIPP does not have a long time series. We therefore assume that the response of transfer income mirrors that of the CPI every four quarters. This is a reasonable assumption for two reasons. First, Social Security payments, which form the bulk of transfer income, are explicitly indexed to the CPI. As this indexation happens only once a year, we cumulate the IRF over four quarters, and produce a step-wise IRF that moves transfer income only in the first quarter of the year.

Second, transfer income is small for the majority of the population. Appendix Figure A1 shows average transfer income over the life cycle in the SIPP. Until the age of 65, almost every household type receives less than \$100 per month in transfer income. Labor and asset income for these “prime-age” households is over 20 times larger. This suggests that the transfers received by young households have only a small effect on the total welfare effect of inflationary shocks.

3.5 Shocks and Time Series Controls

Our first application considers responses to the Känzig (2021) oil supply news shock. High frequency identification techniques are used to address the challenges stemming from the endogeneity of OPEC decisions and global economic conditions. The author uses changes in oil price futures on OPEC announcement dates as an instrument in a six-variable oil price SVAR-IV (real oil price, world oil production, world industrial production, U.S. industrial production, and the U.S. CPI). Under a sufficiently tight window, global economic conditions are unlikely to change within the window and the impact of news about future oil supply is isolated.

Our second application examines impulse responses to the Gertler and Karadi (2015) monetary policy shock. The authors use surprises in federal funds rate futures on FOMC announcement days in tight windows (30 minutes) to isolate the impact of news about monetary policy. These surprises are used as instruments in a monetary SVAR-IV and can capture shocks to forward guidance.

The oil supply news shock and monetary policy shock series are identified from the SVAR-IV specifications of Känzig (2021) and Gertler and Karadi (2015) respectively. In the case of the latter, we extend the VAR data and federal funds rate futures surprises through June 2019 using an updated version of the Gürkaynak, Sack, and Swanson (2005) surprises dataset featured

in Gürkaynak, Karasoy-Can, and Lee (2022). Given these identified shocks, we construct impulse responses to the variables affecting the household’s budget constraint with an “internal instrument” SVAR framework (described in Section 4).

Our monetary application features the set of controls used in Gertler and Karadi (2015), which uses control variables typically included in monetary VARs. Notably, we include the U.S. industrial production index as a measure of real activity. We also use the Gilchrist-Zakrajšek excess bond premium as a measure of the spread between yields on private and public debt due to financial market frictions.¹⁹ Similarly, our oil application mimics the set of controls used in Känzig (2021), which uses variables typically featured in oil VARs. The controls include the West Texas Intermediate crude oil price deflated by the CPI as a measure for real oil prices, the Baumeister and Hamilton (2019) world industrial production measure, Energy Information Administration world oil production, and the Kilian and Murphy (2014) measure of world oil inventories.

3.6 Summary Statistics

Table 1 reports descriptive statistics for our four survey datasets. Columns 1 through 3 report the statistics for households whose head has a high school degree, some college, or a college degree, respectively. Column 4 reports statistics for the full sample of households whose head is at least 25 years old. The four datasets all have similar education and age mixes. Approximately 31-34% of households have just a high school degree, 28-30% have some college, while 37-40% have a college degree. The average age, conditional on being at least 25, is 51 years old in all of our datasets. College-educated households are slightly younger than their less-educated counterparts, reflecting increased educational attainment across cohorts. The full age distribution, included in Appendix Table A1, also matches well across all datasets.

Each dataset has a substantial sample within our three education groups. Considering only households led by individuals over 25 years old in 2019, the CEX has 23,927 observations, the outgoing rotation groups of the CPS has 138,270, the SCF has 26,750 and the SIPP has 14,429 observations. This is encouraging in that sampling error is unlikely to drive our results.

The average annual consumption expenditures in the CEX is \$56,138. However, this is large variance across the educational groups, with high school households consuming \$39,495 and college-educated households consuming \$73,665 per year. This partly reflects differences in income: the CPS reports average weekly earnings of \$786 for high-school household heads and \$1,450 for college-educated households. Likewise, unemployment rates for high school educated households (4.3%) are nearly double that of college-educated households (2.2%). The increased income also translates into larger wealth for the highly-educated: the net worth is around \$1.5 million for college-educated households, but just \$260 thousand for those with less than a high school degree. The asset holdings numbers reported here mirror well those found elsewhere in the literature (Bartscher et al., 2021). These differences in portfolios, income and

¹⁹We use an updated version of the Gilchrist and Zakrajšek (2012) excess bond premium maintained by the Federal Reserve Board.

TABLE 1: Descriptive Statistics: Cross-Sectional Survey Data

	HS or Less	Some College	College+	Full Sample
<i>Panel (a): Consumer Expenditure Survey (CEX)</i>				
Share of sample (%)	31.51	30.14	38.35	100.00
Average age	52.8	51.1	49.4	51.0
Annual Expenditure	\$39495	\$51170	\$73665	\$56138
Motor Fuel Consumption	\$2505	\$2769	\$2854	\$2719
Food at Home Consumption	\$6932	\$7031	\$8249	\$7467
Shelter Consumption	\$9855	\$11988	\$18614	\$13862
Observations	7424	7248	9255	23927
<i>Panel (b): Current Population Survey (CPS)</i>				
Share of sample (%)	32.5	28.4	39.1	100.0
Average age	54.0	51.9	50.1	51.9
Married Rate (%)	44.1	46.7	55.3	49.2
Unemployment Rate, Household Head (%)	4.3	3.3	2.2	3.3
Unemployment Rate, Spouse (%)	2.3	1.8	1.5	1.9
Av. Weekly Earnings, Household Head	\$786	\$932	\$1450	\$1125
Av. Weekly Earnings, Spouse	\$845	\$1044	\$1375	\$1146
Observations	45486	39896	52888	138270
<i>Panel (c): Survey of Consumer Finances (SCF)</i>				
Share of sample (%)	34.49	28.32	37.18	100.00
Average age	52.4	51.2	50.8	51.5
Total Asset Holdings (1000s)	\$294	\$431	\$1581	\$811
Equity Holdings (1000s)	\$42	\$72	\$446	\$200
Bond Holdings (1000s)	\$16	\$36	\$178	\$82
Housing Holdings (1000s)	\$206	\$263	\$514	\$353
	(58.72%)	(63.70%)	(76.15%)	(66.61%)
Net wealth (1000s)	\$260	\$391	\$1548	\$776
Observations	7803	6457	12490	26750
<i>Panel (d): Survey of Income Program Participation (SIPP)</i>				
Share of sample (%)	32.09	28.68	39.83	100.00
Average age	53.0	51.1	48.6	50.7
Annual Transfer Income	\$7264	\$7527	\$5754	\$6737
Means-based Programs	\$618	\$327	\$141	\$345
Social Insurance	\$6646	\$7200	\$5612	\$6392
Observations	5028	4103	5298	14429

Notes: All dollar units are 2019 dollars. In Panel (c), Total Asset Holdings includes equity, bonds, housing, vehicles, liquid assets, business wealth, and other financial and non-financial assets. Additionally, Housing holdings are the average over households with positive holdings. In parenthesis —below the housing holdings— is presented the share of households with positive holdings in this asset class. Age and education correspond to that of the household head in every sample. All numbers average over all of 2019. CPS data correspond to the outgoing rotation groups (ORG) sample of the CPS. Only households whose head is at least 25 years old are included.

consumption patterns will be crucial for the differential welfare effects of shocks.

4. ESTIMATING IMPULSE RESPONSE FUNCTIONS (IRFS)

To construct impulse responses, we take an “identified shock” view following the sizable empirical literature.²⁰ For our oil application, we begin by replicating the baseline SVAR-IV featured in Känzig (2021). The 12-lag log-level VAR includes the real price of oil, world oil production, world oil inventories, world industrial production, US industrial production, and the US consumer price index (CPI) using monthly data from 1974:M1 to 2017:M12. We use the Känzig (2021) oil futures surprises series as an instrument, which uses variation in oil futures prices around OPEC production announcements. Given the reduced form VAR parameters and instrument, the shorter sample 1983:M4 to 2017:M12 (corresponding to the instrument’s sample period) is used to identify the column of the VAR impact matrix corresponding to the oil supply news shock. Finally, the oil supply news shock is itself identified under invertibility using the procedure described in Section 2.1.4 of Stock and Watson (2018).

Equipped with the estimated oil supply news shock, we next estimate impulse responses for each of our outcome variables. We view the estimated oil supply news shock as being sound; the SVAR identification procedure plausibly satisfies invertibility, the relevance condition, and the exclusion restriction. Plainly, unexpected OPEC supply announcements are exogenous to other fundamental drivers of our outcomes, conditional on the SVAR controls. However, the estimated shock is potentially measured with error arising from estimation uncertainty.

We therefore treat the estimated oil supply news shock as an “internal instrument” in separate recursive SVARs for each of the outcome variables (Plagborg-Møller and Wolf, 2021). Let \mathbf{y}_t^{oil} contain the set of variables included in the baseline Känzig (2021) oil SVAR, y_t^n give the n^{th} outcome variable, and z_t contain the estimated oil supply news shock. For constant c and coefficients A_j , we estimate the following SVAR for $\mathbf{y}_t = [z_t, \mathbf{y}_t^{oil}, y_t^n]'$ where the estimated shock z_t is ordered first

$$\mathbf{y}_t = \mathbf{c} + \mathbf{A}_1 \mathbf{y}_{t-1} + \dots + \mathbf{A}_{12} \mathbf{y}_{t-12} + \mathbf{H} \epsilon_t.$$

Then, from the “internal instrument” recursive causal ordering, the first column of \mathbf{H} (denoted as $\mathbf{H}_{\cdot 1}$) identifies the impact response (where the horizon $h = 0$) of the oil supply news shock. We store the element of $\mathbf{H}_{\cdot 1}$ corresponding to the response of outcome variable y_t^n as $\Psi_{n,0}$. The impulse responses for subsequent horizons $\Psi_{n,h}$ can be computed by propagating the oil supply news shock through the VAR model.

Standard errors are computed using the moving block bootstrap featured in Jentsch and Lunsford (2019). To help address the well-known small sample bias of VAR estimation, we report bias-corrected point estimates with percentile confidence intervals.

The “identified shock” view is a product of our setting’s differing outcome variable sample

²⁰See Ramey (2016) for a detailed discussion.

lengths. While most outcome variables have long samples, some are shorter—like CPI: Postage and Delivery Services (1998:M12-2019:M12). Separating shock estimation from outcome variable impulse response computation allows for all available information to be exploited; the estimated oil supply news shock is created using data corresponding to the *entire* sample length of the oil futures surprises series. In contrast, a procedure that combines shock estimation and impulse response function computation is constrained by the outcome variable’s available sample. This procedure is also transparent in the sense that identification step for computing the estimated oil supply news shock uses the same controls as Känzig (2021).

We follow an identical approach for the monetary application. We replicate and update the Gertler and Karadi (2015) baseline 12-lag log-level monetary SVAR-IV. The VAR contains the one-year government bond rate, industrial production, Gilchrist and Zakrajšek (2012) excess bond premium, and the Consumer Price Index. Updates to the excess bond premium are maintained by the Federal Reserve Board, described in the preceding section. The instrument for the one-year government bond rate—the three month ahead monthly fed futures surprises—is updated using data from Gürkaynak et al. (2022). Mirroring the Gertler and Karadi (2015) baseline specification, the reduced form VAR is estimated using data from 1979:M1-2019:M6 while the shorter sample 1990:M1-2019:M6 (corresponding to the availability of the fed futures surprises series) is used to identify the column of the VAR impact matrix corresponding to the monetary policy shock and the shock itself. Just as in the oil shock application, we view the estimated monetary policy shock as being measured with error. The estimated monetary policy shock is then used as an instrument in an “internal instrument” recursive SVAR. This 12-lag VAR augments the initial monetary VAR with the outcome variable and the estimated monetary policy shock (ordered first) and is estimated using the largest available sample.

We project impulse response functions for four years in all of our applications and assume the marginal welfare effects are zero after that four year horizon.²¹ Estimation of effects over longer horizons is challenging in time series contexts. Our exercise thus seeks to study the short-run effects of inflationary shocks.

5. ESTIMATED IRFS

This section reports the impulse response functions to monetary and oil price shocks, estimated using the local projection IV approach laid out in the prior section. Subsection 5.1 presents impulses to oil supply news shocks, while subsection 5.2 presents impulses from monetary policy shocks.

²¹We do not estimate IRFs for business wealth as there is scant reliable data available on these assets’ returns or prices.

5.1 Oil Supply News Shocks (Känzig, 2021)

We estimate impulses to an increase in the price of oil, caused by announced contractions in OPEC oil supply following Känzig (2021). We scale the size of the shock to represent a 10% increase in the West Texas Intermediates (WTI) crude oil price, since the standard deviation of monthly oil price growth is around 10%.²² Oil supply shocks may be thought of as a form of aggregate supply or cost-push shock, which may be especially pronounced for industries for which energy is a large share of costs.

Figure 1 plots the impulse response function of various consumption prices in response to this oil price increase. Panel A plots the path of the WTI oil price in response to the supply news shock – this is the path of the “oil price shock” that we consider. By construction, crude oil prices jump by 10% on impact due to the shock. Over the course of the following four years, the crude oil price converges back to its pre-shock level. This increase in the price of oil leads to the aggregate CPI-U rising by 15.5 basis points on impact (Panel B), which grows to 35 basis points after two quarters before converging back to the pre-shock path for the aggregate price index. This is consistent with Känzig (2021)’s findings for the aggregate economy.²³

The aggregate path of inflation masks rich heterogeneity in the inflation experience of different goods. Panels C and D of Figure 1 plot the response of the disaggregated CPI categories for motor fuel and fuels and utilities (e.g. the cost of heating a home), respectively. After the oil price shock, the price of motor fuel rises by 2.2 percent on impact, rising to a 4.7 percent increase after one quarter. Meanwhile, the price of fuels and utilities (e.g., the cost of heating one’s house) rises by 0.3 percent on impact, rising to 1.2 percent one year after the shock. The crude oil price shock is associated with large movements in motor fuel prices but more modest, if still substantial, movements in fuel and utilities.

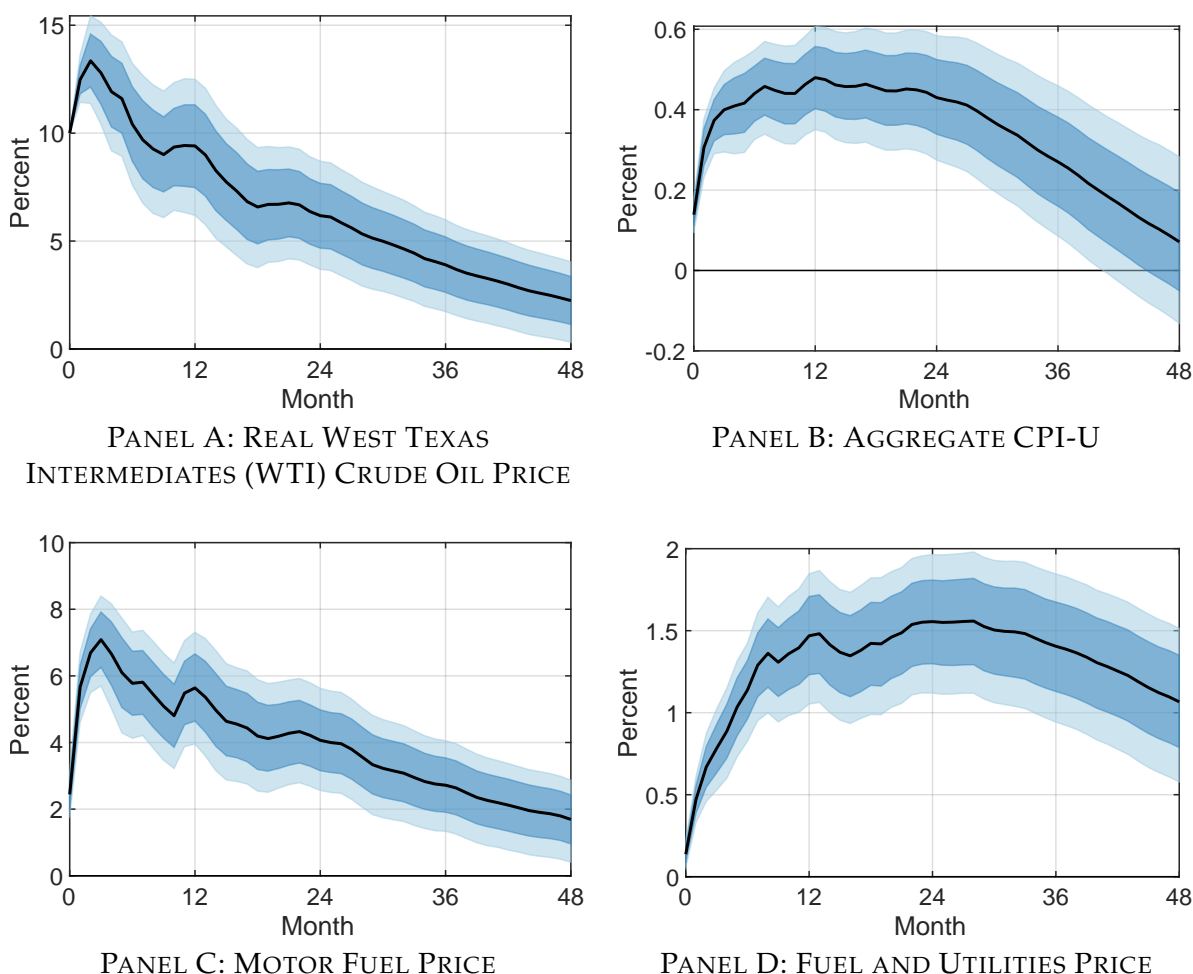
To visualize the effect of oil price shocks on disaggregated goods prices, Figure 2 presents coefficient plots of impulse responses for all of our disaggregated CPI subcategories measured at different horizons. Panel A plots the response of each subcategory’s price on impact, while Panel B plots the response after 12 months. For instance, the fuels and utilities coefficient shows the aforementioned 0.3% increase on impact in Panel A, and the 1.2% increase after one year in Panel B. These responses principally arise through the confluence of 1) the extent to which oil price movements constitute cost-push shocks for each good and 2) the elasticity of demand for each product.

Figure 2 shows that, intuitively, motor fuel experiences by far the largest increase in response to the crude oil price shock, both on impact and after a year. The next largest price movements come from “fuel and utilities,” “information technology, hardware and services,” and “public

²²Propositions 1 and 2 are related to first-order welfare effects of small price movements. Our time series regressions identify the effect of shocks up to scale. To estimate the effect of a smaller than 10% oil price shock—say a 1% shock—one simply needs to rescale all our results by a factor of 1/10. We therefore choose a 10% shock to be able to interpret our results as the effects of a one standard deviation oil price movement, but note that our propositions should strictly-speaking be applied to smaller shocks.

²³Känzig (2021) also finds that the oil price shock reduces aggregate US industrial production and consumption, a decline in the S&P500 index and a rise in aggregate unemployment rates, which we verify below.

FIGURE 1: Impulse Responses for an Oil Price Shock

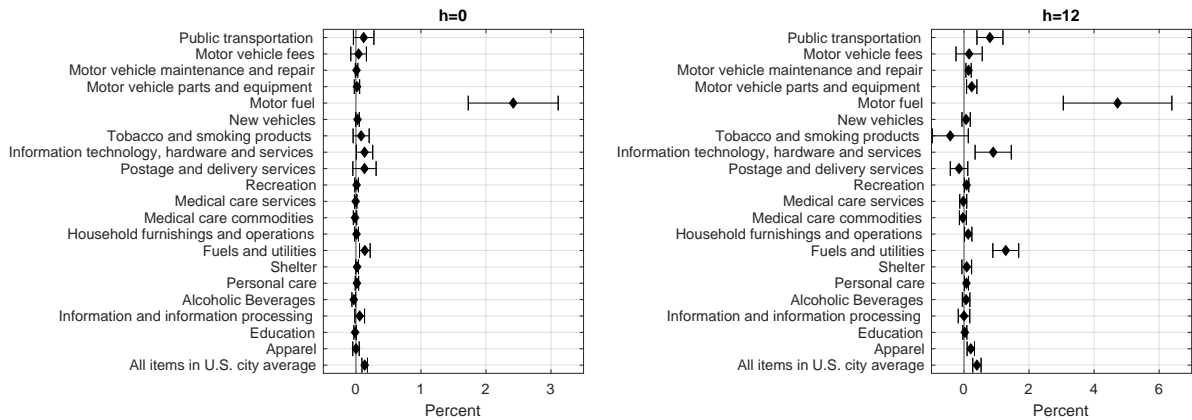


Notes: Figure plots cumulative impulse response functions (IRFs) to inflationary oil supply news shocks constructed by Känzig (2021). Shocks normalized to represent a 10% increase in the Real West Texas Intermediates (WTI) crude oil price in high frequency windows around OPEC supply announcements. IRFs estimated using the “internal instrument” SVAR procedure explained in section 4. Panel A plots the IRF of the real WTI crude oil price: it is the IRF of the directly shocked variable. Panel B plots the IRF of aggregate CPI-U, while Panels C and D report the IRFs of the CPI categories for Motor Fuel and Fuel and Utilities. All regressions control for industrial production in the US and the world, world oil production and world oil inventories. The SVAR is specified with 12 lags. Dark blue regions specifies 68% confidence interval, and light blue regions the 90% confidence interval.

transportation.” All of these goods rely heavily on energy in production. In contrast, the price of goods such as medical care, recreation or education—which do not have a large energy cost share in production—show essentially no response to the oil price shock. Thus the extent to which inflationary oil price shocks affect household well-being through the consumption channel will be primarily determined by household expenditures on motor fuel. We explore the cross-sectional patterns of consumption in section 6.1 below.

Turning to the labor income channel, Figure 3 plots the estimated response of unemployment (Panels A through C) and nominal weekly earnings conditional on being employed (Panels D through F) to the 10% increase in the price of crude oil for our three education groups. We aggregate all age groups together for these plots.

FIGURE 2: Response of Disaggregated CPI Prices to an Oil Price Shock



PANEL A: IMPACT RESPONSE

PANEL B: RESPONSE AFTER 12 MONTHS

Notes: Figure plots impulse response functions (IRFs) of CPI consumption goods to inflationary oil supply news shocks constructed by Känzig (2021). Shocks normalized to represent a 10% increase in the Real West Texas Intermediates (WTI) crude oil price in high frequency windows around OPEC supply announcements. IRFs estimated using the “internal instrument” SVAR procedure explained in section 4. Panel A plots the IRF on impact (i.e. at 0-horizon) of each consumption good, while Panel B plots the cumulative IRF 12 months after the shock impulse. All regressions control for industrial production in the US and the world, world oil production and world oil inventories. The SVAR is specified with 12 lags. Error bars represent 90% confidence intervals.

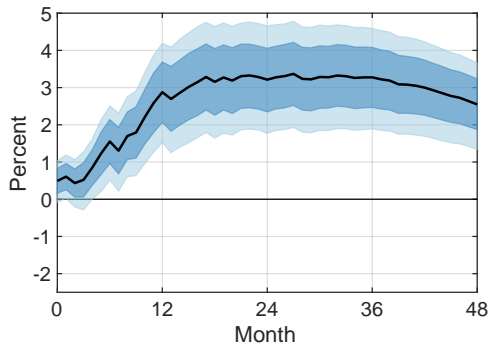
The figure shows that households with limited education experience relatively large increases in (log) unemployment in response to oil price increases. The oil price shock leads to an unemployment rate increase of 0.15 percent for these households after two years, which continues to grow over time. Thus, given the standard deviation of quarterly log unemployment rates for this group is 0.22, the one standard deviation oil price increase leads to an increase in low-education unemployment of around 0.68 standard deviations. In contrast, those with a bachelor’s degree or higher experience very little job loss due to the oil shocks, while those with some college have a muted response.

Conditional on remaining employed, however, those with a high school education or less experience relatively muted declines in weekly earnings compared with those with some college. The wages of college-educated households again appear somewhat insulated from the effects of an oil price shock, except perhaps after 13 months. Overall, the 10% oil price shock leads to between a 0.2 and 0.6 log point reduction in earnings, depending on the group, that manifests after around three years. The lack of response for low-education households may in part be due to institutional downward rigidities, such as the minimum wage or unions. These rigidities may be part of the reason why unemployment rates respond so much amongst low-education workers.²⁴

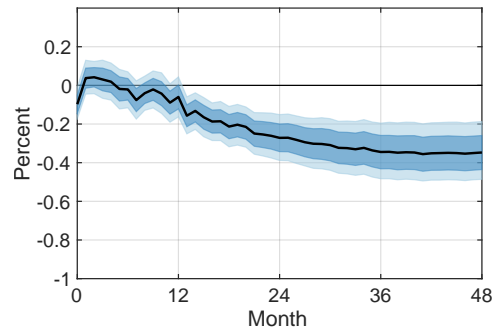
Finally, Figure 4 shows a coefficient plot of oil price responses for various asset prices and dividend responses. We assume that nominal bond coupon payments are fixed and do not respond to the shock. This is a reasonable assumption for the U.S.: adjustable rate mortgages

²⁴See, for instance, Grigsby, Hurst, Yildirmaz, and Zhestkova (2021) for evidence of differential wage rigidity at the bottom of the distribution and Faia and Pezone (2023) for evidence that wage rigidity affects employment responses to monetary shocks.

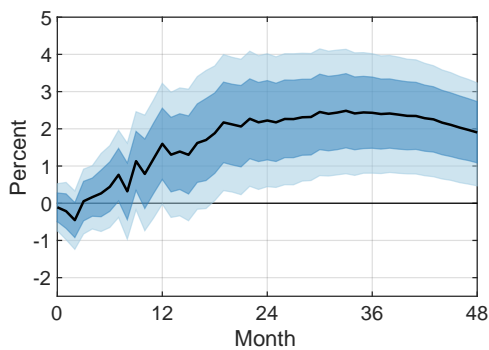
FIGURE 3: Impulse Response of Labor Income to an Oil Price Shock



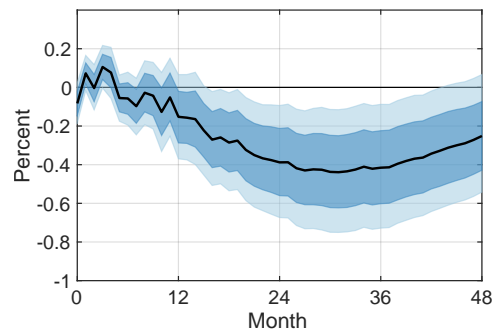
PANEL A: LOG UNEMPLOYMENT RATE, HIGH SCHOOL OR LESS



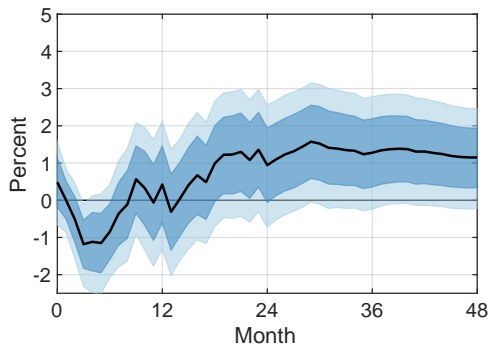
PANEL D: LOG WEEKLY EARNINGS, HIGH SCHOOL OR LESS



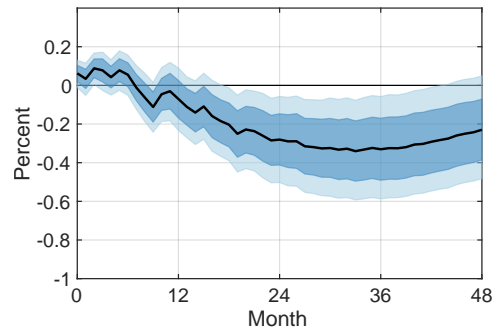
PANEL B: LOG UNEMPLOYMENT RATE, SOME COLLEGE



PANEL E: LOG WEEKLY EARNINGS, SOME COLLEGE



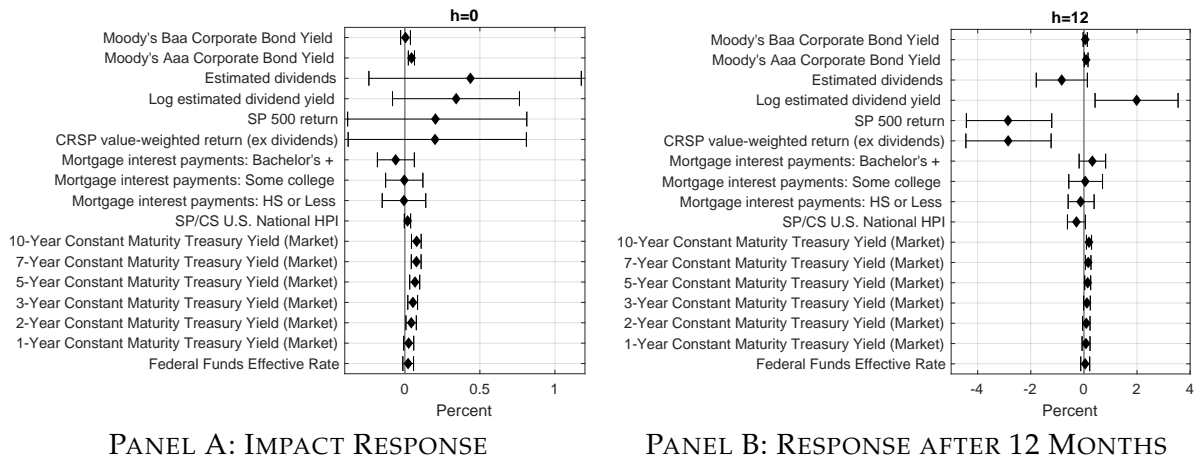
PANEL C: LOG UNEMPLOYMENT RATE, BACHELOR'S DEGREE OR MORE



PANEL F: LOG WEEKLY EARNINGS, BACHELOR'S DEGREE OR MORE

Notes: Figure plots cumulative impulse response functions (IRFs) to inflationary oil supply news shocks constructed by Känzig (2021). Shocks normalized to represent a 10% increase in the Real West Texas Intermediates (WTI) crude oil price in high frequency windows around OPEC supply announcements. IRFs estimated using the “internal instrument” SVAR procedure explained in section 4. Panels A, C and E plot the IRFs of log unemployment rates for those with high school or less (HS), some college, or at least a bachelor’s degree, respectively. Panels B, D and F report the IRFs of log weekly earnings for households in our three education groups, controlling for group-specific life cycle profiles and composition effects following section 3.2. Unemployment data constructed from the Current Population Survey (CPS) and weekly earnings constructed using the the Outgoing Rotations Group (ORG) of the CPS. All regressions control for industrial production in the US and the world, world oil production and world oil inventories. The SVAR is specified with 12 lags. Dark blue regions specifies 68% confidence interval, and light blue regions the 90% confidence interval.

FIGURE 4: Response of Asset Prices and Dividends to an Oil Price Shock



Notes: Figure plots impulse response functions (IRFs) of asset prices and dividend yields to inflationary oil supply news shocks constructed by Känzig (2021). Shocks normalized to represent a 10% increase in the Real West Texas Intermediates (WTI) crude oil price in high frequency windows around OPEC supply announcements. IRFs estimated using the “internal instrument” SVAR procedure explained in section 4. Panel A plots the IRF on impact (i.e. at 0-horizon) of each asset good, while Panel B plots the cumulative IRF 12 months after the shock impulse. All regressions control for industrial production in the US and the world, world oil production and world oil inventories. The SVAR is specified with 12 lags. Error bars represent 90% confidence intervals.

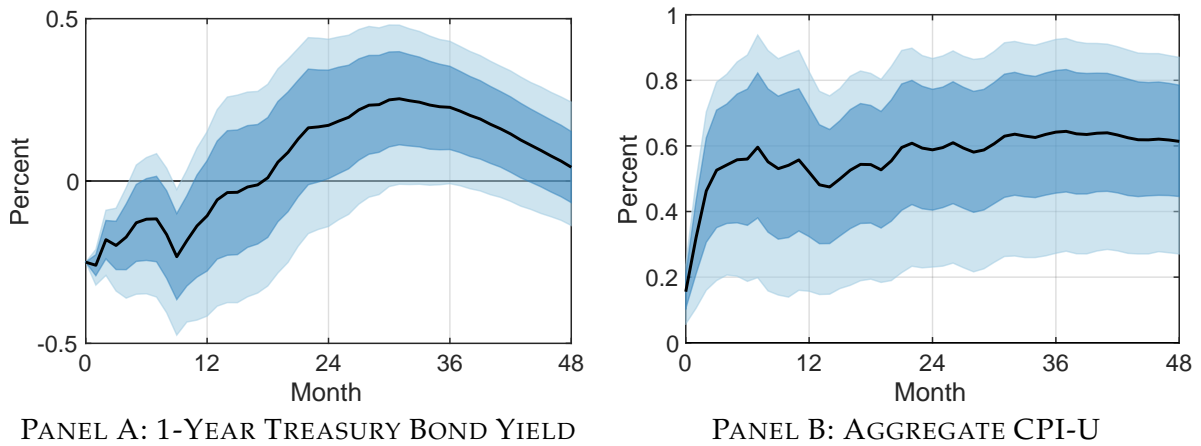
made up less than 10% of new mortgage originations in 2010 (Moench, Vickery, and Aragon, 2010), credit card debt often has a fixed APR, and both U.S. treasury bonds and corporate bonds usually pay a fixed coupon. Since we use 2019 dollars as the numeraire, this implies that bond income is pre-determined and does not respond to shocks.

Figure 4 shows limited impact on most asset prices on impact. However, the value of the S&P500 declines by 2% a year after the initial oil price shock. This decline is partially accounted for by dividend payments, which fall by around 82 basis points. As the cost of inputs rises, firms earn lower profits and pay lower dividends. The expectation of this continued low dividend payout leads to lower equity prices. Meanwhile, we find essentially no effect on any other asset price: house prices, treasury bonds, and corporate bonds are all largely unresponsive to the oil price shock. This suggests that our results are principally driven by the oil shock itself, rather than policy responses to the shock: were policy to significantly respond to the shock, one should expect to see meaningful movements in treasury bond yields.

This implies that equity holders lose from the oil price shock because they receive lower dividend income. Importantly, however, the oil price increase is beneficial to those who were planning to accumulate equity, because it is now cheaper to do so. Thus the strength of the portfolio channel of oil prices welfare effects depends critically on who holds and is accumulating equities, but is largely unaffected by other household asset decisions.

In summary, the welfare loss associated with rising oil prices will be strongest for those whose consumption expenditures focus on motor fuel and fuel and utilities, those who own equity, those who plan to be net sellers of equity, and low-education households who experience the largest increases in unemployment. As we will soon see, many of these patterns are reversed

FIGURE 5: Impulse Responses to an Expansionary Monetary Policy Shock



Notes: Figure plots cumulative impulse response functions (IRFs) to inflationary monetary policy shocks constructed by Gertler and Karadi (2015). Shocks normalized to represent a 25 basis point decrease in the one-year treasury bond yield in 30-minute windows around FOMC announcements. IRFs estimated using the “internal instrument” SVAR procedure explained in section 4. Panel A plots the IRF of the 1-year treasury yield: it is the IRF of the directly shocked variable. Panel B plots the IRF of aggregate CPI-U. All regressions control for US industrial production, the excess bond premium (Gilchrist and Zakrajšek, 2012) and aggregate CPI. The SVAR is specified with 12 lags. Dark blue regions specifies 68% confidence interval, and light blue regions the 90% confidence interval.

for monetary expansions.

5.2 Monetary Policy Shocks (Gertler and Karadi, 2015)

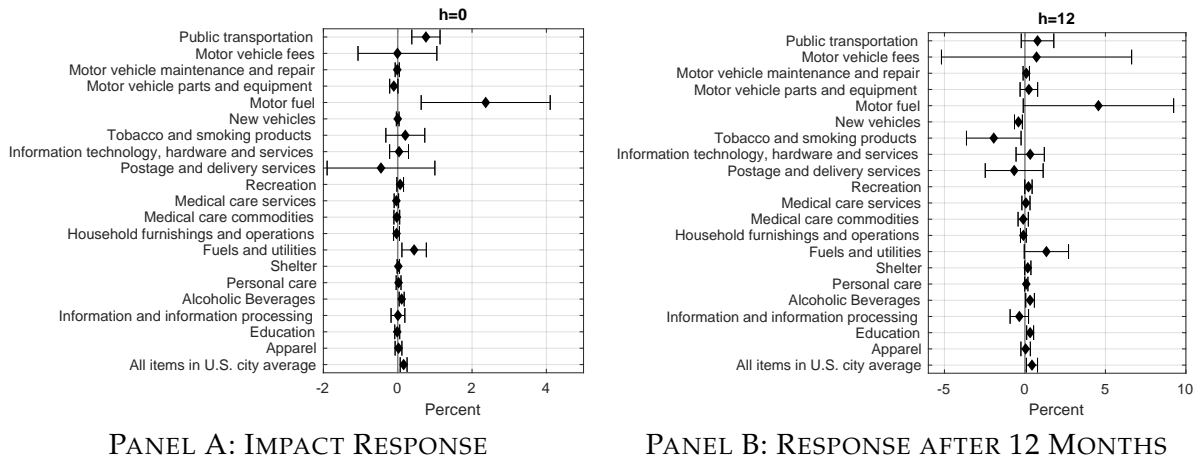
Following Gertler and Karadi (2015), we estimate impulses to a monetary policy shock identified by high-frequency responses in the one-year treasury yield around Federal Reserve announcements. We scale the shock to represent a 25 basis point decline in the one-year treasury yield. This is a natural scaling both because it represents a common adjustment in the Federal Funds Rate and because, as we will soon see, it generates the same aggregate inflation response as a 10% oil price shock.

Figure 5 plots the estimated response of the one-year treasury yield (Panel A) and aggregate CPI-U (Panel B) in response to the monetary shock. The initial 25 basis point decline in treasury yields gradually dissipates over the subsequent two years. We have scaled the monetary shock such that it generates a similar impact response of aggregate inflation as the 10% oil price shock: on impact, the 25 basis point decline in nominal interest rates generates an increase in the aggregate CPI-U of 16.2 basis points, which rises to 34 basis points after two quarters.

This aggregate effect on CPI masks heterogeneity across consumption goods. Figure 6 plots the response of consumption goods prices in the quarter of the shock (Panel A) and after one year (Panel B). As in the oil shock, the impact response is largest for motor fuel, fuel and utilities, and public transport. After four quarters, our point estimates suggest somewhat less variation in price responses, however the standard errors are large. Like oil prices, therefore, the consumption channel will be dominated by motor fuel expenditures.

While the consumption channel is similar across inflationary oil price shocks and monetary

FIGURE 6: Response of Disaggregated CPI Prices to an Expansionary Monetary Policy Shock



PANEL A: IMPACT RESPONSE **PANEL B: RESPONSE AFTER 12 MONTHS**

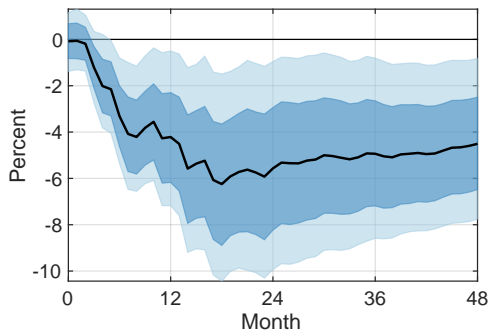
Notes: Figure plots cumulative impulse response functions (IRFs) to inflationary monetary policy shocks constructed by Gertler and Karadi (2015). Shocks normalized to represent a 25 basis point decrease in the one-year treasury bond yield in 30-minute windows around FOMC announcements. IRFs estimated using “internal instrument” SVAR procedure explained in section 4. Panel A plots the IRF on impact (i.e. at 0-horizon) of each consumption good, while Panel B plots the cumulative IRF 12 months after the shock impulse. All regressions control for US industrial production, the excess bond premium (Gilchrist and Zakrajšek, 2012) and aggregate CPI. The SVAR is specified with 12 lags. Error bars represent 90% confidence intervals.

shocks, the labor income channel is extremely different. Figure 7 plots the response of unemployment rates and log weekly earnings across the three education groups. While oil price increases lead to increases in unemployment and declines in log wages, monetary shocks have, if anything, the opposite effect. A 25 basis point cut in interest rates leads to a 0.83 log point decline in unemployment for those with less than high school education after one year. There is a similarly-sized response for those with some college education. Those with at least a bachelor’s degree also see unemployment rate declines of 1.96 log points.²⁵ There is limited response for log wages of employed workers for any of our groups. However, the unemployment responses indicate that the labor income channel will push towards a welfare gain from expansionary monetary shocks, which is opposite to the oil price shock both overall and distributionally.

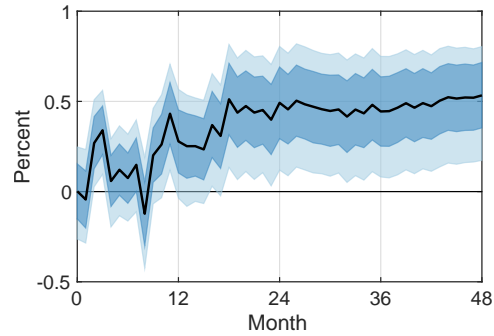
Turning to the portfolio channel, Figure 4 plots impulse response functions for four large asset classes. Panel A plots the cumulative return on the CRSP stock market index, excluding dividends. This represents the change in stock prices induced by the monetary shock. A 25 basis point decline in interest rates leads to a 3 percentage point increase in stock prices on impact. The reason for this is twofold. First, dividend payments increase substantially in the subsequent periods, as shown by Panel B. Second, a no-arbitrage argument implies that a decline in bond yields necessitates a decline in expected return of stock prices, which happens if stock prices rise on impact. After four quarters, the stock price and dividend increases both level out 4 percent higher than they would be absent the shock.

²⁵Note that we estimate the response of log unemployment $\ln u$ to be consistent with prior literature (Broer et al., 2022). However, the framework requires the response of $\ln(1 - u)$. To convert unemployment IRFs to those needed by the framework, we multiply the estimated IRFs by $-u/(1 - u)$, since $d \ln(1 - u) = -du/(1 - u) = -u/(1 - u)d \ln u$.

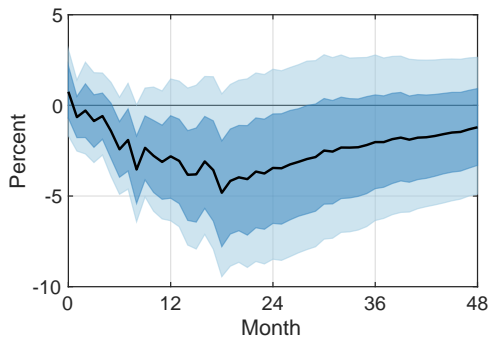
FIGURE 7: Impulse Response of Labor Income to an Expansionary Monetary Policy Shock



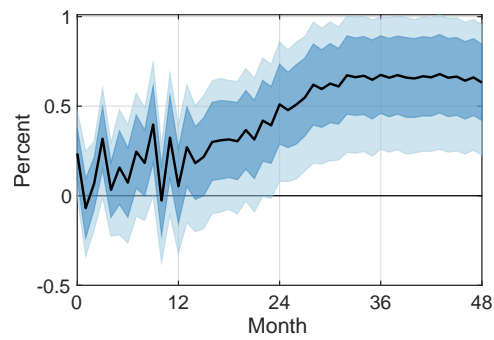
PANEL A: LOG UNEMPLOYMENT RATE, HIGH SCHOOL OR LESS



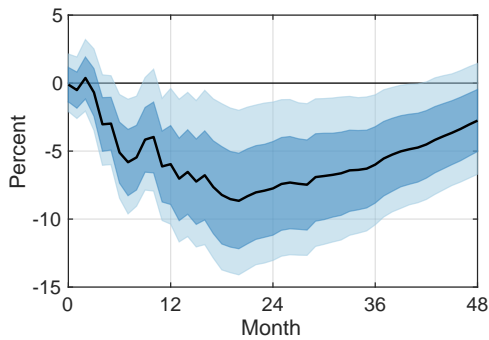
PANEL D: LOG WEEKLY EARNINGS, HIGH SCHOOL OR LESS



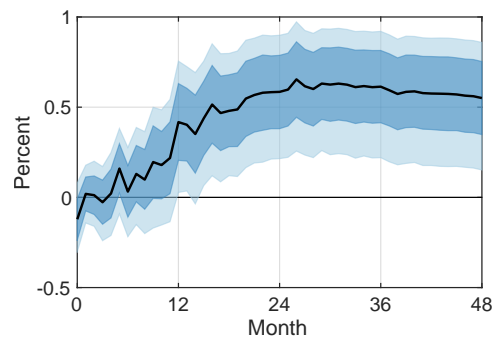
PANEL B: LOG UNEMPLOYMENT RATE, SOME COLLEGE



PANEL E: LOG WEEKLY EARNINGS, SOME COLLEGE



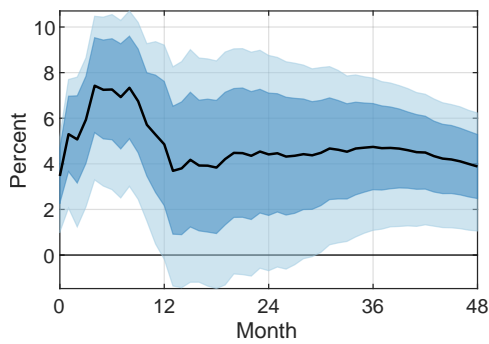
PANEL C: LOG UNEMPLOYMENT RATE, BACHELOR'S DEGREE OR MORE



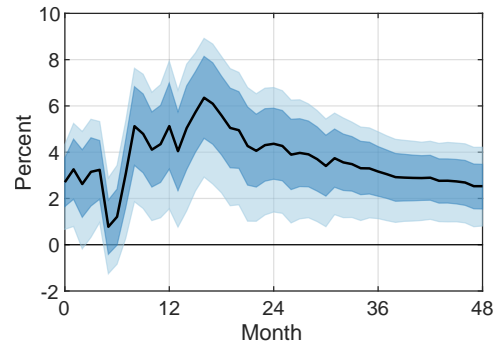
PANEL F: LOG WEEKLY EARNINGS, BACHELOR'S DEGREE OR MORE

Notes: Figure plots cumulative impulse response functions (IRFs) to inflationary monetary policy shocks constructed by Gertler and Karadi (2015). Shocks normalized to represent a 25 basis point decrease in the one-year treasury bond yield in 30-minute windows around FOMC announcements. IRFs estimated using the “internal instrument” SVAR procedure explained in section 4. Panels A, C and E plot the IRFs of log unemployment rates for those with high school or less (HS), some college, or at least a bachelor’s degree, respectively. Panels B, D and F report the IRFs of log weekly earnings for households in our three education groups, controlling for group-specific life cycle profiles and composition effects following section 3.2. Unemployment data constructed from the Current Population Survey (CPS) and weekly earnings constructed using the the Outgoing Rotations Group (ORG) of the CPS. All regressions control for US industrial production, the excess bond premium (Gilchrist and Zakrajšek, 2012) and aggregate CPI. The SVAR is specified with 12 lags. Dark blue regions specifies 68% confidence interval, and light blue regions the 90% confidence interval.

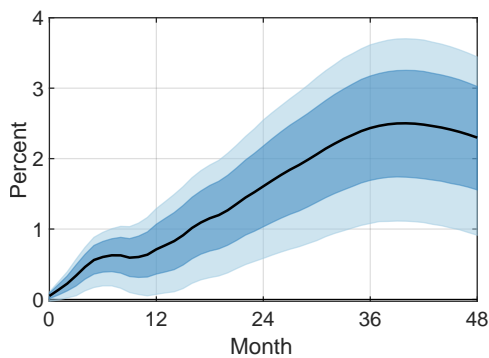
FIGURE 8: Impulse Response of Asset Prices to an Expansionary Monetary Policy Shock



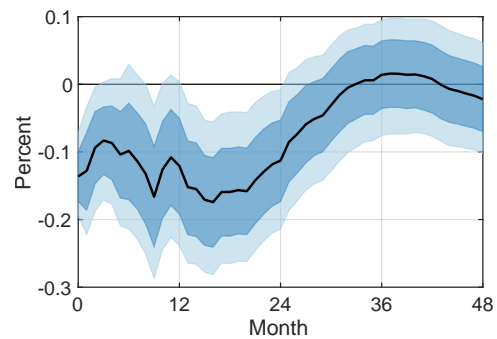
PANEL A: CUMULATIVE STOCK MARKET RETURN, EXCL. DIVIDENDS



PANEL B: DIVIDEND PAYOUTS



PANEL C: CASE-SHILLER HOME PRICE INDEX (HPI)



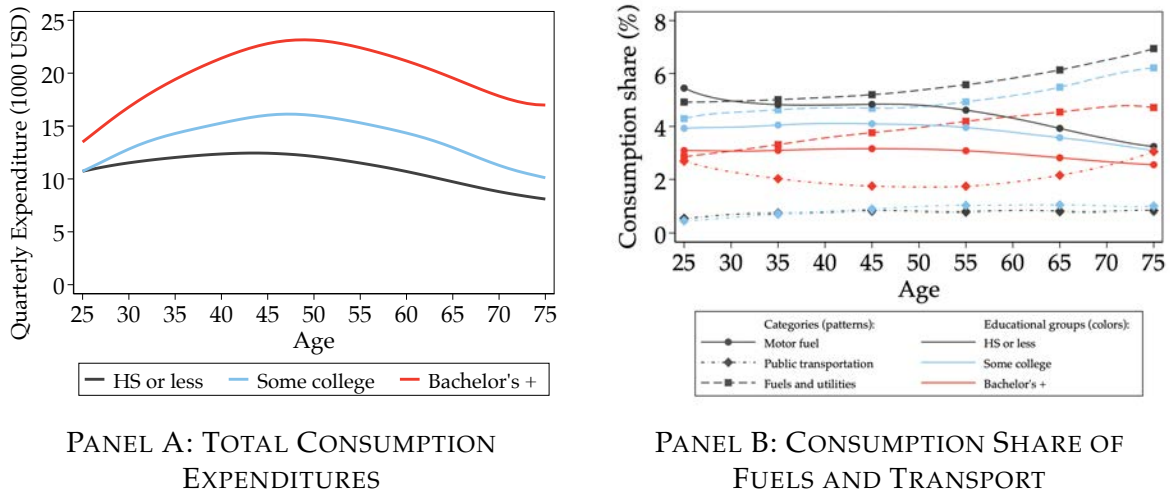
PANEL D: MOODY'S AAA CORPORATE BOND YIELDS

Notes: Figure plots cumulative impulse response functions (IRFs) to inflationary monetary policy shocks constructed by Gertler and Karadi (2015). Shocks normalized to represent a 25 basis point decrease in the one-year treasury bond yield in 30-minute windows around FOMC announcements. IRFs estimated using “internal instrument” SVAR procedure explained in section 4. Panel A plots the IRF of S&P500 returns, excluding dividends, while Panel B plots the IRF of dividend payouts on the S&P500. Panel C plots the IRF of the Case-Shiller Home Price Index, while Panel D plots the IRF of Moody’s AAA Corporate Bond Yields. All regressions control for US industrial production, the excess bond premium (Gilchrist and Zakrajšek, 2012) and aggregate CPI. The SVAR is specified with 12 lags. Dark blue regions specifies 68% confidence interval, and light blue regions the 90% confidence interval. Error bars represent 90% confidence intervals.

Panel C shows the effect of nominal rate cuts on the Case-Shiller home price index (HPI). The figure shows that declines in interest rates lead to gradual increases in house prices, which peak at around 1.5 percent increases after 3 years. This could occur either because there is pass-through from treasury bonds to mortgage rates which takes a few years to realize, because the increase in aggregate demand caused by expansionary monetary policy shifts housing demand, or because a decline in interest rates benefits long-duration assets such as housing. The delayed response of house prices suggest the first two forces may be the strongest. Finally, Panel D shows a limited effect of monetary policy on corporate bond yields. Thus, the portfolio channel principally benefits those who are selling equities or their home. This is in stark contrast to oil shocks, which benefit those who would *buy* equities.

Overall, inflation caused by monetary policy shocks hurts those who consume motor fuel or fuel and utilities, and those who would purchase housing or equity, but benefits those who

FIGURE 9: Life Cycle Consumption Shares



Notes: Panel A plots total quarterly expenditure by group, and Panel B plots shares of expenditure on fuel and transport by group. Data is from the Consumer Expenditure Survey for 2019. In Panel A, expenditure is averaged within group and age, and then a LOWESS smoother is applied across age. Panel B averages expenditure shares across 10 year age bins. Note that public transportation includes air travel.

hold equities and the labor income of the college-educated. In sharp contrast, inflation caused by oil price shocks hurts those who hold equities and the labor income of the less-educated, those who would purchase equities, and those whose consumption basket overweights motor fuel and utilities. We next present evidence from our cross-sectional data to see which households are most affected by these patterns.

6. MONEY-METRIC WELFARE CALCULATIONS

This section aggregates the impulse response functions estimated in the last section into money metric welfare effects of inflationary oil price and monetary shocks. We first present cross-sectional patterns from U.S. survey data to study who holds assets whose prices respond to shocks, and whose expenditures concentrate on the highly-responsive motor fuel. We then compute money metric welfare changes from oil supply and monetary shocks, and decompose the overall welfare effects into the consumption, labor income and portfolio channels.

6.1 Features of Cross-Sectional Consumption, Wage and Asset Data

Figure 9 plots consumption patterns over the life cycle for the three education groups. Panel A plots total quarterly consumption in the 2019 CEX, smoothed over the life cycle using a LOWESS smoother. We estimate a hump-shaped life cycle consumption profile for all three groups, consistent with that found elsewhere in the literature (Browning and Lusardi, 1996; Attanasio, 1999). Red lines correspond to those with at least a bachelor's degree, the light blue lines represent those with some college, while the black lines show the patterns for those with

a high school education or less. We find that college-educated households spend, on average, \$21412 per quarter at the age of 40, while those with a high school education or less spend just \$12360 per quarter. This naturally creates a scale dependence for the money-metric welfare calculations to come: college-educated households will have larger money-metric welfare movements simply because they have higher expenditures, asset holdings and labor income. Thus, for some analysis, we will normalize the money metric welfare effects of the shock by four-year consumption expenditures of each household type.²⁶

Panel B shows the share of household consumption expenditures that are accounted for by the three categories most responsive to oil and monetary shocks. The figure shows that those with a high school education or less spend a larger share of their income on motor fuel and fuel and utilities. Motor fuel has a hump-shaped life cycle profile: the oldest households drive less because they do not commute to work. Fuel and utilities expenditure are rising through the life cycle as households move into larger houses. Public transport, which includes air travel, occupies a relatively small share of consumption, but is largest among young and old college-educated households. The patterns here suggest that the consumption channel will be largest—as a share of consumption—for those with less than a high school education. Therefore, we might expect the consumption channel to push towards regressive inflation regardless of its source.

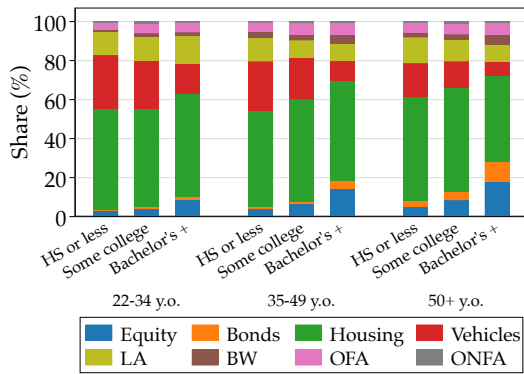
Turning to the portfolio channel, Figure 10 plots asset holdings and accumulation patterns over the life cycle for our three education groups. Panel A plots the share of total assets held in equities (blue bars), bonds (yellow bars), housing (green bars), vehicles (red bars), liquid assets such as checking/savings accounts and cash (yellow bars), business wealth (brown bars), and other financial (pink bars) or non-financial assets (gray bars). By far the largest share of assets for most households is housing. We see, however, that equities constitute a larger share of assets for older households and those with at least a college education. This suggests that the dividend responses documented above will have a larger impact on older college educated households. Note that the dominance of housing and vehicles in all portfolios illustrates the importance of accounting for their durable nature as both a consumption good and store of value, as outlined in Section 2.3.

Panel B plots the accumulation of equity over the life cycle. Again, we smooth accumulation profiles using a LOWESS smoother. All education groups accumulate equity during middle age, before slowing accumulation or decumulating after the retirement age.²⁷ This hump-shaped accumulation pattern is especially strong for those with a college education, however. This implies that middle-aged college-educated households will realize large welfare gains if equity prices fall, and losses if equity prices rise. Likewise, Panel C shows that college-educated households accumulate more housing than do low-education households, at least until around age 60. This means that reductions in housing prices will benefit younger col-

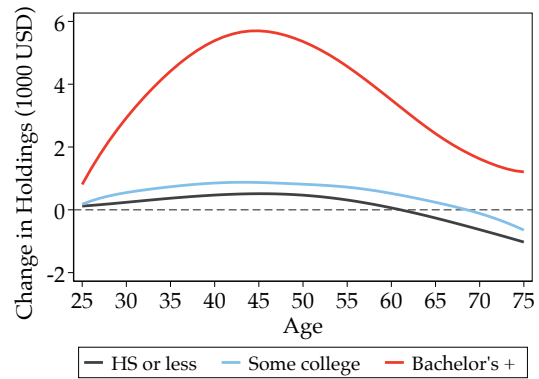
²⁶To calculate these four-year consumption profiles, our baseline exercise maintains the assumption that 2019 is a steady state and projects household expenditures forward using life cycle profiles assuming $p_{j0}c_{jt}^a = p_{j0}c_{j0}^{a+t}$.

²⁷The fact that households continue saving past retirement has been documented by, among others, De Nardi, French, Jones, and McGee (2021). This is often attributed to bequest motives or uncertain longevity. We capture this by allowing asset holdings to enter the utility function.

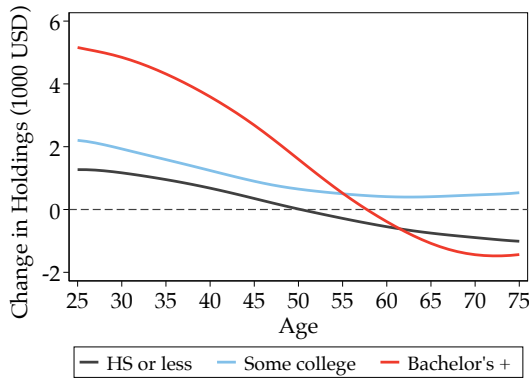
FIGURE 10: Life Cycle Asset Portfolios and Accumulation



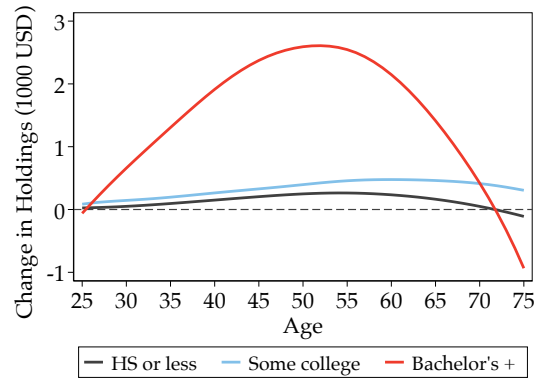
PANEL A: PORTFOLIO SHARES



PANEL B: EQUITY ACCUMULATION



PANEL C: HOUSING ACCUMULATION



PANEL D: NON-CORP. BOND ACCUMULATION

Notes: Panel A plots asset share by group. LA stands for Liquid Assets (mainly cash and checking/savings accounts). BW stands for Business Wealth. OFA stands for Other Financial Assets. OFNA stands for Other Non-Financial Assets. Panel C does the same for the change in housing wealth, and Panel D for non-corporate Bond Accumulation. Panel B averages the change in household's equity portfolio between age bins by group, and applies a LOWESS smoother. Panel C does the same for housing wealth, and Panel D for non-corporate bond holdings. Data for all panels is from the Survey of Consumer Finances for 2019.

lege educated households on average. Older households, however, decumulate housing; thus house price increases are beneficial for them. Finally Panel D shows a similar hump-shaped profile in non-corporate bond accumulation.

6.2 Welfare Effects of Oil Price Shocks

Given the relatively short time series of our data, we choose to estimate impulse response functions out to a horizon of 16 quarters. Thus, in calculating the welfare formula in equation (7), we limit ourselves to the cumulative effects over a truncated, four-year horizon. In principle, nothing prevents one from extending the analysis to further years. In practice however, there is a tradeoff between the precision of the estimates and the length of period studied.

Table 2 reports the money-metric welfare change in response to an oil shock and a monetary

policy shock for each of our three education groups and across three age groups. Columns (1) through (5) report the money-metric welfare effects of a 10% oil price increase, while columns (6) through (10) report the money-metric welfare effects of a 25 basis point cut in interest rates. Within a shocked group, each column reports a different source of the overall money-metric welfare change. Column (11) reports estimated total consumption over the four-years starting the fourth quarter of 2019 for each household group. In addition, we normalize the money-metric welfare losses from oil shocks by four years of total consumption for each household in Figure 11. Since this Figure plots welfare losses, negative numbers represent welfare gains.²⁸

Column (1) reports the consumption channel of oil price shocks. We find that young and middle-aged households suffer the most from the consumption channel in dollar terms, but this is relative similar across our three education groups. This is largely due to a scale effect: middle-aged high-education households spend 81% more per quarter than their low-education counterparts, as evidenced by column (11). This implies that the money-metric welfare effect on these individuals is large: they must be paid more to be made whole for a given price movement. However, Panel A of Figure 11 shows that the consumption channel is larger as a share of total consumption for low-education households, especially at younger ages. Such low-education households lose the equivalent of 0.25% of consumption from the oil price shock, compared with less than 0.1% for college-educated households. This is due to the patterns discussed above: low-education household have a larger share of expenditure on motor fuel and fuels and utilities, which are the most responsive to oil price shocks.

Column (2) reports the labor income channel of oil price shocks. Overall, the estimated impulse responses in section 5.1 show only small differences in labor income losses between education groups. Thus Panel B of Figure 11 shows that the labor income channel is of similar size across the three education groups as a share of total consumption: each group's welfare declines by approximately 0.25% of consumption during working age. The labor income channel is also of similar size over the life cycle, until workers start to retire at older ages and labor income becomes less important.

Column (3) reports the portfolio channel of oil shocks. Strikingly, we find that the portfolio channel is negligible for young low-education workers, but large and *positive* for middle-aged high-education workers. Because high-education households accumulate equity in the middle of their life, temporarily falling equity prices are beneficial to them. This is offset partially by the decline in dividend payouts as a result of the shock, which disproportionately hurt those who hold equities. This offsetting dividend effect is thus particularly pronounced among older college-educated households. However, this dividend effect is significantly smaller than the asset accumulation effects, since dividend income is smaller than equity accumulation for college-educated households. Panel C of Figure 11 shows that middle-aged college-educated households gain almost 0.5% of consumption from the portfolio channel, while old households with just a high school education lose a little less than 0.25% of consumption. This is a major force towards regressivity of inflation induced by oil price shocks.

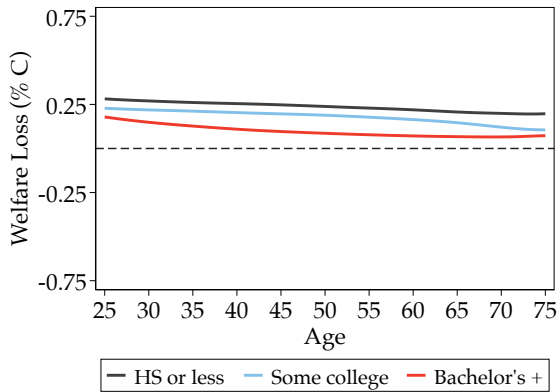
²⁸We plot the full life cycle profiles of raw money-metric welfare effects in Appendix C.

TABLE 2: Money-Metric Welfare Effects of Inflationary Oil Supply and Monetary Policy Shocks, by education and age group (2019 \$)

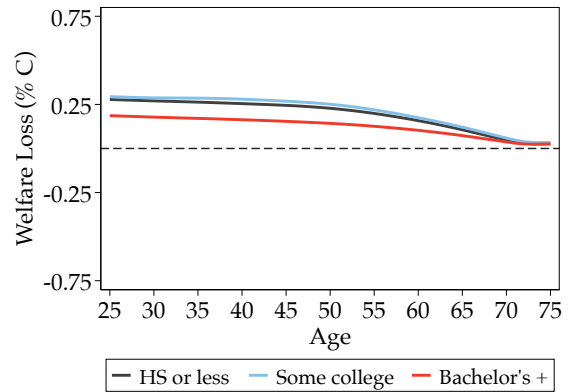
Household Group	Oil Supply Shock					Monetary Policy Shock					Total Consumption (11)	
	Consumption Channel (1)	Labor Income Channel (2)	Portfolio Channel (3)	Transfers Channel (4)	Total Welfare (5)	Consumption Channel (6)	Labor Income Channel (7)	Portfolio Channel (8)	Transfers Channel (9)	Total Welfare (10)		
<i>HS or Less</i>												
22-35 y.o.	-\$535	-\$536	+\$125	+\$14	-\$932	-\$1115	+\$1498	-\$403	+\$22	+\$2	\$198 702	
35-50 y.o.	-\$522	-\$517	+\$142	+\$26	-\$870	-\$1211	+\$1503	-\$418	+\$40	-\$85	\$208 874	
51+ y.o.	-\$345	-\$190	-\$159	+\$181	-\$513	-\$1010	+\$580	+\$311	+\$283	+\$164	\$162 487	
<i>Some College</i>												
22-35 y.o.	-\$495	-\$652	+\$274	+\$18	-\$856	-\$1129	+\$1139	-\$1533	+\$28	-\$1496	\$226 562	
35-50 y.o.	-\$531	-\$723	+\$251	+\$32	-\$972	-\$1356	+\$1328	-\$1375	+\$50	-\$1354	\$266 893	
51+ y.o.	-\$320	-\$298	-\$38	+\$206	-\$449	-\$1307	+\$569	-\$480	+\$322	-\$895	\$214 457	
<i>College+</i>												
22-35 y.o.	-\$444	-\$537	+\$158	+\$5	+\$572	-\$1474	+\$1777	-\$3813	+\$8	-\$3501	\$303 290	
35-50 y.o.	-\$388	-\$595	+\$1804	+\$13	+\$833	-\$2102	+\$2155	-\$4124	+\$20	-\$4051	\$378 697	
51+ y.o.	-\$239	-\$270	+\$232	+\$208	-\$69	-\$2073	+\$1064	-\$170	+\$327	-\$851	\$330 710	

Notes: Table shows the estimated welfare effects of oil and monetary policy shocks by group and three age bins. Welfare effects and total consumption are cumulated over 16-quarters. All rows average over the corresponding age bins, weighting by the age distribution of the CPS.

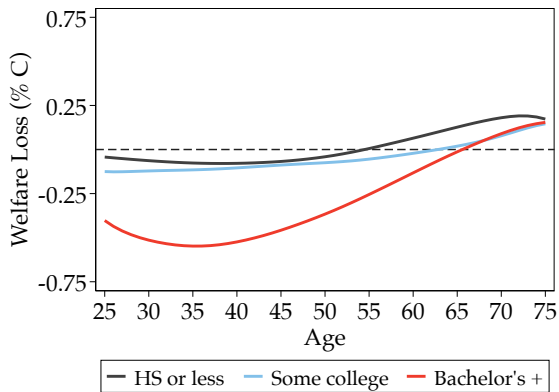
FIGURE 11: Welfare Losses From Inflationary Oil Price Shocks: Scaled by Consumption



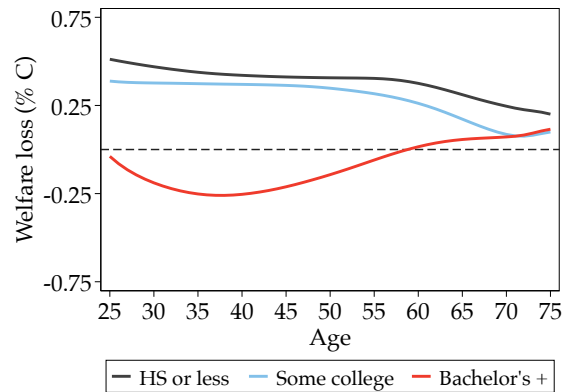
PANEL A: CONSUMPTION CHANNEL



PANEL B: LABOR INCOME CHANNEL



PANEL C: PORTFOLIO CHANNEL



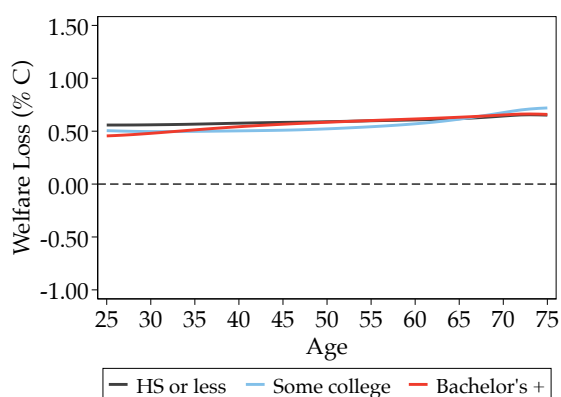
PANEL D: TOTAL WELFARE CHANGE

Notes: Figure shows the estimated welfare loss from a 10% oil price shock. We normalise the figures by total four year consumption by age and group, using projected lifecycle consumption patterns in 2019. Panels A-C split the effect into the consumption, labor income and portfolio channels. A negative number represents a welfare gain.

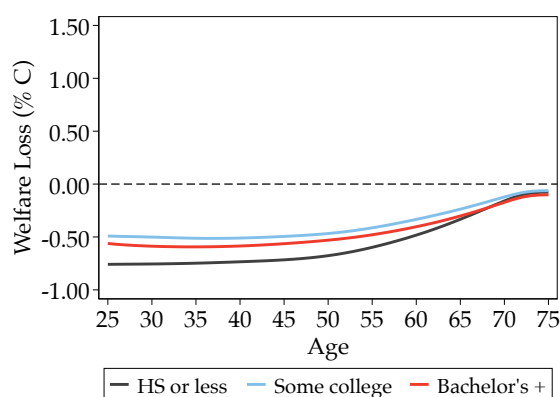
The transfer channel, in contrast, is small compared to the other three channels (Column (4)). It reflects the formulaic upward adjustment to Social Security payments induced by the oil shock's effect on the CPI, and exclusively benefits the old. We present estimates of this channel by age in Appendix Figure A6.

Finally, Column (5) shows the total money-metric welfare change induced by a supply-driven 10% increase in the price of oil. We find that middle-aged households with no more than a high school degree must be paid \$870 to achieve the same utility level as was attainable absent the oil price shock. In stark contrast, college-educated households of the same age would need to lose \$833 in order to return to their pre-shock indifference curve. Younger college-educated households, who are still net equity-accumulators, also gain around \$572, but older high-education households lose \$69. Panel D of Figure 11 shows that these total welfare effects result in a loss of around 0.5% of consumption for lower-education households which is relatively flat across the life cycle, with the exception that young households lose more

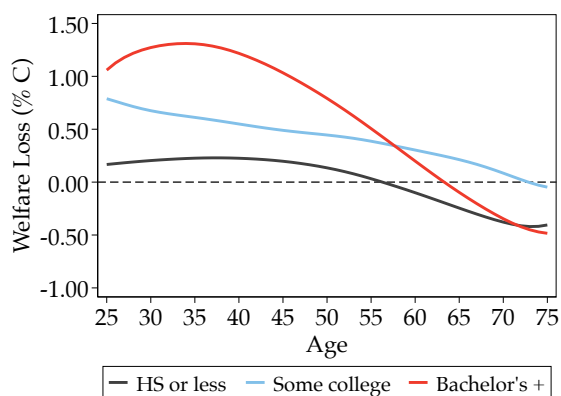
FIGURE 12: Welfare Losses From of Inflationary Monetary Policy Shocks: Scaled by Consumption



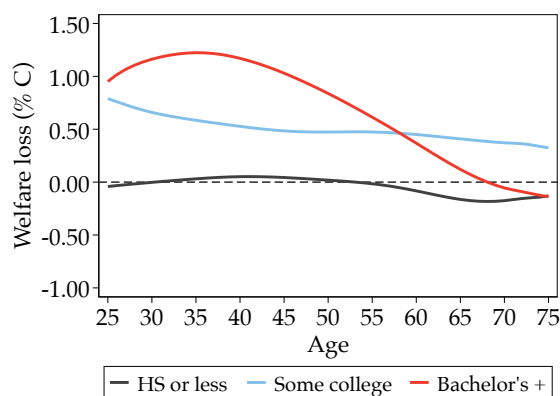
PANEL A: CONSUMPTION CHANNEL



PANEL B: LABOR INCOME CHANNEL



PANEL C: PORTFOLIO CHANNEL



PANEL D: TOTAL WELFARE CHANGE

Notes: Figure shows the estimated welfare loss from a 25 basis point cut to the federal funds rate. We normalise the figures by total four year consumption by age and group, using projected lifecycle consumption patterns in 2019. Panels A-C split the effect into the consumption channel, the labor income channel. A negative number represents a welfare gain.

from labor income channel. High-education households have a U-shaped life cycle profile, gaining around 0.1% of consumption when young, rising to 0.3% when middle-aged and losing around 0.25% when elderly. Inflation caused by oil supply contractions therefore appear highly regressive, mostly due to the consumption and portfolio channels.²⁹

6.3 Welfare Effects of Monetary Shocks

Columns (6) through (10) of Table 2 show the money-metric welfare effects over 16 quarters of a 25 basis point reduction in interest rates caused by monetary policy announcements. Figure 12 shows these welfare changes normalized by four year total consumption.

²⁹Note that, due to discounting, equity price increases would still hurt young accumulators if the price response lasted forever and they could resell equities when old.

Column (6) shows that the consumption channel is larger for households with more consumption. However, Panel A of Figure 12 shows that in proportional terms, the consumption channel affects all household types equally. Less educated households gain the most from the labor income channel in proportional terms, matching the patterns found in Germany (Broer et al., 2022) and Sweden (Coglianese et al., 2022). The labor income and consumption channels thus combine to push towards progressivity of inflationary monetary policy.

The portfolio channel likewise pushes towards monetary policy being progressive. Column (8) shows that young and middle-aged high-education households suffer money-metric welfare losses of over \$3800 from this channel, while similarly-aged low-education households lose less than \$420. This is large as a share of initial consumption: Panel C of Figure 12 shows that the portfolio channel leads to a loss of 1.2% of consumption for middle-aged college-educated households. Again, we see an important life cycle profile to the portfolio effect. This is driven by three forces. First, middle-aged college educated households accumulate equity, and thus are hurt by rising equity prices. Second, while all households accumulate housing throughout much of their life cycle, college-educated households accumulate at a faster rate and earlier than households with a high school education. Thus, rising house prices especially hurt younger households with a college education. Finally, there is a countervailing force through asset income: older college-educated households own more equities and benefit from the increased dividend payouts as a result of expansionary monetary policy. These effects combine to generate the U-shaped life-cycle profile of welfare effects among college-educated households, and much larger welfare losses among college-educated households than lower-education households through the portfolio channel.

Combining these effects together reveals that inflationary shocks to monetary policy are on balance progressive. Column (10) shows that those with a high school degree or less must be paid between \$85 (if middle-aged) or lose \$164 (if middle-aged) in order to achieve the same pre-shock utility. In contrast, middle-aged college-educated households must be paid \$4051 to achieve their pre-shock utility, while older college-educated households must be paid \$851. As a share of consumption, this implies that the movements in labor income, consumption prices and asset prices net to approximately zero for households with at most a high-school education, while young and middle-aged high-education households see losses of over 1% of consumption, which decline in the later stages of life (Panel D). While all households are hurt by the rise in consumption prices caused by monetary policy, this loss is exacerbated by movements in asset prices for high-education households. Meanwhile, low-education households are compensated by rising labor income. Thus rate cuts disproportionately benefit low-education households. On the flip side, rate increases would disproportionately harm these households.

Note that this has an important policy implication – if the monetary authority responds to oil-price induced inflation by unexpectedly raising interest rates, it may exacerbate the distributional consequences of the initial oil price shock. While our methodology does not allow us to address whether changes in the policy *rule* would have a similar distributional impact, the regressive nature of disinflationary monetary policy shocks is noteworthy.

Our results also provides reduced form evidence that monetary policy has scope to differentially affect households at different points of the distribution. This is a precondition for optimal policy to incorporate inequality considerations as forcefully argued by McKay and Wolf (2023a). Again, however, there remains the caveat that our results concern policy shocks rather than rules.

7. DISCUSSION: ROBUSTNESS AND EXTENSIONS

Although our framework has a number of strengths, it is not without limitations. This section discusses some of these limitations, and explores how robust the paper’s conclusions are to alternative model specifications and estimation strategies.

Before exploring deviations from our baseline analysis, it is worth reiterating the strengths of our approach. The framework we have developed permits the estimation of the first-order welfare effects of identified macro shocks without specifying a full structural model. This “reduced-form” approach to welfare calculation is computationally tractable, thereby permitting a great deal of heterogeneity across households. Its use is primarily limited by the availability of data: the framework can in principal accommodate an arbitrarily large variety of goods, assets and household types as well as non-parametric differentiable utility functions. It does not necessitate specifying a production side of the economy to solve for general equilibrium: rather, we estimate general equilibrium effects that identified macro shocks have on prices.

Nevertheless, the framework has a few core limitations. First, we have thus far ignored borrowing constraints and short-selling constraints on assets. Such constraints have a long tradition in macro and finance and may be welfare relevant. We are able to accommodate such constraints in our framework with additional parametric assumptions, as we explain below. In addition, we have assumed that the “no-shock” choices are observable as choices in 2019, implicitly assuming that 2019 constitutes a steady state. We relax this in a robustness test below. Finally, our estimation remains an approximation in that it considers first order welfare effects from a small-noise expansion.

Borrowing Constraints. We now consider the role of borrowing constraints. Suppose now that each agent faces an additional constraint on net wealth, such that

$$\sum_k Q_{kt} N_{kt}^a \geq \underline{b}.$$

We show in Appendix D that the change in welfare from an impulse is then

$$(11) \quad (dV^a) / \frac{U_C}{P_0^a} = \sum_t R_{0 \rightarrow t}^{-1} \overbrace{\prod_{s=0}^t (1 + \tau_s^a)^{-1}}^{\text{Discount Wedge}} \left(- \sum_j p_{j,t} c_{jt}^a \Psi_{i,t}^{p,j} + W_t^a L_t^a \Psi_{i,t+h}^W + T_t^a \Psi_{i,t}^T \right) + \sum_k \left[N_{kt-1}^a D_{kt} \Psi_{i,t}^{D,k} - Q_{kt} \Delta N_{kt}^a \Psi_{i,t}^{Q,k} \right] + \underbrace{\frac{\tau_t^a}{1 + \tau_t^a} \sum_k Q_{kt} N_{kt} \Psi_{kt}^Q}_{\text{Value of Relaxed Constraint}}$$

There are two additional terms when we compare (11) to (6). First, an additional term $\tau \geq 0$ in the discount factor for each s reflects a wedge in the consumer's Euler equation, which is active where they are at the constraint on total borrowing. Specifically, τ_t is an increasing function of the Lagrange multiplier on the period t net worth constraint and solves

$$(12) \quad \frac{P_{t+1}^a U_C(C_t^a, \{N_{kt}^a\}, L_t^a)}{P_t^a U_C(C_{t+1}^a, \{N_{kt+1}^a\}, L_{t+1}^a)} = \beta \left(\frac{\delta_{t+1}^a}{\delta_t^a} \right) \cdot R_t (1 + \tau_t^a)$$

If the constraint does not bind, $\tau_t^a = 0$ and there is no wedge in the Euler equation. In this case, households set the marginal rate of substitution between consumption in t and $t + 1$ equal to the real interest rate. The more binding the constraint, the larger is τ_t^a and the greater the difference in discount rates.

The second additional term reflects the extra value from rising asset prices after the impulse in relaxing the borrowing constraint of the consumer. Again, this term is growing the larger is τ_t^a or, equivalently, the more binding is the constraint.

Obtaining values for τ_t^a from the data requires imposing parametric structure on the utility function of each consumer and so is not without loss. However, doing so is a useful exercise to explore the role of borrowing constraints. We here assume that

$$(13) \quad U(C_t, \{N_{kt}\}, L_t) = \log(C_t) + \hat{U}(\{N_{kt}\}, L_t)$$

so that the LHS of (12) is simply growth in consumption expenditures. Were there no net worth constraints, households would perfectly intertemporally smooth their consumption: deviations from this perfect smoothing are thus informative about the bindingness of constraints.

We calculate τ_t^a by using annual consumption data from the CEX to compute yearly consumption growth by age for the years 1990 to 2019. We assume $\beta = (0.98)^{\frac{1}{4}}$ and test robustness to assuming a constant death rate or empirical death rates taken from the Social Security Administration.³⁰ The average value of τ over the life cycle is plotted by education group in Appendix Figure A8. These wedges decline with age, reflecting growth in consumption. They are also larger for those with at least a Bachelor's degree, who experience sharper lifetime consumption growth (see Figure 9) and do not seem able to perfectly smooth consumption. This may reflect both the presence of student debt and the steeper life cycle profile of income for

³⁰See <https://www.ssa.gov/oact/STATS/table4c6.html>, retrieved February 1, 2023.

TABLE 3: Robustness of Estimated Welfare Effects of Oil and Monetary Shocks

Specification	Oil Supply News Shock			Monetary Policy Shock		
	\leq HS (1)	Some College (2)	College+ (3)	\leq HS (4)	Some College (5)	College+ (6)
Baseline	-\$798	-\$816	+\$494	+\$23	-\$1278	-\$3055
Borrowing Constraints (no Death)	-\$845	-\$893	+\$76	+\$202	-\$977	-\$1718
Borrowing Constraints (incl. Death)	-\$881	-\$959	-\$394	+\$298	-\$804	-\$1036
UI Replacement Rates	-\$778	-\$804	+\$504	-\$17	-\$1302	-\$3103
Project No-Shock Choices	-\$824	-\$810	+\$720	-\$40	-\$1399	-\$3505
No CPI Controls	-\$811	-\$837	+\$439	+\$12	-\$1272	-\$3009

Notes: This Table reports robustness of our baseline welfare results to a variety of specifications. Each dollar value is the weighted average for each educational group of the welfare effects over the lifecycle, restricting only ages between 25 and 65. Columns (1)-(3) report robustness of oil supply shock impacts, while columns (4)-(6) report robustness of monetary shock impacts. Columns (1) and (4) consider those with high school education or less, (2) and (5) consider those with some college, while (3) and (6) consider those with at least a college degree. All columns average over the life cycle, weighting by the age distribution of the CPS. The first row shows our baseline estimates. The second and third rows consider an extension of the model in which households are subject to a net worth constraint. The second row assumes constant death rates, while the third row uses empirical life cycle death rates. The fourth row downweights the response of unemployment by 16% to account for the role of unemployment insurance. The fifth row assumes the continuation of product-, asset-, and household group-specific log-linear trends in consumption, labor income and portfolio choice, as described in detail in Appendix D.2. The sixth row omits controls for aggregate CPI from the estimation of the impulse response functions.

these households.

Table 3 shows the results of our robustness exercises. The first row shows our baseline estimates of the welfare effect of oil shocks (columns 1-3) and monetary shocks (columns 4-6). Columns 1 and 4 show the effect on the average household with a high school education or less, columns 2 and 5 show the effect on those with some college, while columns 3 and 6 show the effect on those with at least a college degree.

The second row and third shows our results after accounting for wedges arising from borrowing constraints. The second assumes a constant death rate in equation (12), while the third calculates the death rates using data from the Period Life Table of the Social Security Administration for 2019. The patterns are qualitatively similar to the baseline, and our conclusions about the regressivity of inflationary shocks continue to hold. However, the portfolio change is dampened considerably, especially for young and middle-aged college-educated households. Now, reductions (increases) in asset prices tighten (loosen) the net worth constraint. The bindingness of this constraint is inferred by deviations from the permanent income hypothesis: that is, excess growth in consumption expenditures. Young and middle-aged college-educated households have rapidly growing consumption expenditures, which imply a binding borrowing constraint. What's more, they both hold and accumulate a relatively large amount of assets. Thus, when asset prices increase, the welfare loss from more expensive accumulation is partially offset by a relaxation of the borrowing constraint. This reduces the magnitude of the overall portfolio effect, thus reducing the magnitude of the differences in welfare effects between high and low educated households. We stress that this calculation relies on specific

assumptions regarding the functional forms of utility, discount rates and death rates. It should thus be viewed as a complement for, and not dominant over, our baseline results. Nevertheless, the patterns shown here suggest that the relaxation of these constraints could be quantitatively important for the welfare effects of shocks.

Unemployment Insurance. Our baseline estimation ignores the existence of unemployment insurance (UI). This is largely due to the difficulty of measuring UI accurately throughout the distribution given the complexity of the UI system. The UI program replaces a percentage of lost nominal wages after a job loss. This percentage is around roughly 40%. However, UI has a cap on benefits leading the actual average replacement rate to be substantially lower, especially for high-education households. What’s more, only 40% of unemployed workers actually receive unemployment insurance, either because some idiosyncracies of the UI program render them ineligible or because eligible individuals do not take advantage of the benefit. Given these considerations, Chodorow-Reich and Karabarbounis (2016), whence many of the above numbers originate, measure the true average replacement rate to be 6%. This implies that unemployment insurance likely has a negligible effect on our estimates.

Nevertheless, we consider the importance of the UI program by assuming that 40% of unemployed workers receive UI, and those who receive it recover 40% of their wage income. This is an upper bound on the role of UI. Under these assumptions, one need simply to deflate the impulse response of unemployment rates by 16% ($40\% \times 40\%$), since an unemployment spell leads the household to only lose 84% of their income, rather than all of it. The results of this exercise are reported in the fourth row of Table 3. Incorporating UI replacement rates has only a small impact on our results.

Alternative Assumptions for No-shock Choices. Our baseline estimation assumes that the economy is in steady state in 2019. In reality, there has been differences in trend inflation across different product and asset prices (Argente and Lee, 2021; Jaravel, 2019). To account for this, we compute a robustness test in which we assume that, absent any shocks, households’ consumption expenditure on each product, portfolio values, and labor incomes would all have grown according to their long-run log-linear trends. Our approach to doing so is detailed in Appendix D.2. The results are included in the fifth row of Table 3. Again, this makes little qualitative difference to our estimated welfare effects. This is because the trend growth in the factors leading to declines in utility are offset by trend growth in the factors leading to increases in welfare.

Alternative estimation strategies. We consider an alternative to our baseline “internal instrument” SVAR procedure for estimating impulse response functions in which we drop the aggregate CPI-U price level from our SVAR model. The baseline estimation includes aggregate CPI to mirror the estimation procedures of Känzig (2021) and Gertler and Karadi (2015) as closely as possible. This increases our confidence in the exogeneity of the constructed shocks: the Fed may partially react to aggregate inflation, for instance. However, one might be concerned that including aggregate inflation in the VAR alongside price indices for other specific products may lead to problems due to the cointegration of aggregate CPI with a full set of

product-specific VARs. The sixth row of Table 3 shows, however, that omitting aggregate CPI from our VAR estimation does not meaningfully alter our results.

8. CONCLUSION

This paper estimates the incidence of inflationary macroeconomic shocks by proposing a new methodology to account for movements of all aspects of the budget constraint. We propose a framework in which the money-metric welfare effect of a shock on a given household may be estimated by aggregating empirical impulse response functions for consumption prices, labor income, asset prices, and dividend payouts using cross-sectional consumption, asset portfolio and labor income data. In effect, the incidence of a shock depends on whether the choices a household would make absent the shock become more or less expensive relative to their income.

Using U.S. survey data and standard time series techniques, we find that the source of inflation matters for its distributional consequences. While oil price shocks are regressive, expansionary monetary policy shocks are progressive. This discrepancy is primarily driven by the different effects that these shocks have on households labor income and asset portfolios. While monetary policy raises labor income, dividends, equity prices and house prices, oil price shocks do the opposite. Rising asset prices benefit those who would *sell* the asset; thus middle-aged college-educated households who accumulate equity benefit from oil price shocks and are hurt by monetary policy. Low-education households who principally rely on labor income benefit from monetary policy's positive impact on the labor market, but are hurt by the weaker labor market caused by oil supply contractions. These qualitative conclusions do not depend on specific functional form assumptions for utility or general equilibrium, and robust to a variety of alternate estimation strategies and to allowing for binding borrowing constraints.

Note that our framework is a price theoretic approach to valuing arbitrary movements in prices, wages, and portfolios. The response to oil supply and monetary shocks are two particularly interesting combinations of price movements to study given their perceived importance in driving U.S. inflation. However, future empirical research could apply our framework to study other price movements, such as those induced by as fiscal policy shocks, exchange rate shocks or supply chain disruptions, or to study the impact of oil and monetary shocks in different countries and contexts. Theoretically, our framework does not easily incorporate uncertainty shocks nor preference shocks; doing so would be fruitful ground for future work.

REFERENCES

- AGUIAR, M., M. AMADOR, AND C. ARELLANO (2021): "Micro Risks and Pareto Improving Policies," Working paper.
- AMBERG, N., T. JANSSON, M. KLEIN, AND A. ROGANTINI PICCO (2021): "Five Facts About

- the Distributional Income Effects of Monetary Policy,” *American Economic Review: Insights*, 4, 289–304.
- ARGENTE, D. AND M. LEE (2021): “Cost of Living Inequality during the Great Recession,” *Journal of the European Economic Association*, 19, 913–952.
- ATTANASIO, O. P. (1999): “Consumption,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Elsevier Science B.V., vol. 1, chap. 11, 741–812.
- AUCLERT, A. (2019): “Monetary Policy and the Redistribution Channel,” *American Economic Review*, 109, 2333–67.
- BAQAEE, D. AND A. BURSTEIN (forthcoming): “Welfare and Output with Income Effects and Taste Shocks,” *Quarterly Journal of Economics*.
- BAQAEE, D., A. BURSTEIN, AND Y. KOIKE-MORI (2022): “A New Method for Measuring Welfare with Income Effects using Cross-Sectional Data,” Working paper.
- BAQAEE, D. R. AND A. BURSTEIN (2022): “Aggregate Welfare and Output with Heterogeneous Agents,” Working paper.
- BARTSCHER, A. K., M. KUHN, M. SCHULARICK, AND P. WACHTEL (2021): “Monetary Policy and Racial Inequality,” *FRB of New York Staff Report*.
- BAUMEISTER, C. AND J. D. HAMILTON (2019): “Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks,” *American Economic Review*, 109, 1873–1910.
- BROER, T., J. KRAMER, AND K. MITMAN (2022): “The Curious Incidence of Monetary Policy Shocks Across the Income Distribution,” *Working Paper*.
- BROWNING, M. AND A. LUSARDI (1996): “Household Saving: Micro Theories and Micro Facts,” *Journal of Economic Literature*, 34, 1797–1855.
- BURNSIDE, C., M. EICHENBAUM, AND S. REBELO (1995): “Capital Utilization and Returns to Scale,” *NBER Macroeconomics Annual*, 10.
- CARDOSO, M., C. FERREIRA, J. M. LEIVA, G. NUÑO, Á. ORTIZ, T. RODRIGO, AND S. VAZQUEZ (2022): “The Heterogeneous Impact of Inflation on Households’ Balance Sheets,” Working paper.
- CHARI, V. V., P. J. KEHOE, AND E. R. MCGRATTAN (2007): “Business Cycle Accounting,” *Econometrica*, 75, 781–836.
- CHODOROW-REICH, G. AND L. KARABARBOUNIS (2016): “The cyclicity of the opportunity cost of employment,” *Journal of Political Economy*, 124, 1563–1618.
- COGLIANESE, J., M. OLSSON, AND C. PATTERSON (2022): “Monetary Policy and the Labor Market: A Quasi-Experiment in Sweden,” *Working Paper*.

- DAVILA, E. AND A. SCHAAB (2022): “Welfare Assessments with Heterogeneous Individuals,” *NBER Working Paper Series*.
- DE NARDI, M., E. FRENCH, J. B. JONES, AND R. MCGEE (2021): “Why Do Couples and Singles Save During Retirement?” Working Paper 28828, National Bureau of Economic Research.
- DEXHEIMER, V. (2003): “Hedonic Methods of Price Measurement for Used Cars,” .
- DOEPKE, M. AND M. SCHNEIDER (2006): “Inflation and the Redistribution of Nominal Wealth,” *Journal of Political Economy*, 114, 1069–1097.
- EROSA, A. AND G. VENTURA (2002): “On Inflation as Regressive Consumption Tax,” *Journal of Monetary Economics*, 49, 761–795.
- FAGERENG, A., M. GOMEZ, E. GOUIN-BONENFANT, M. HOLM, B. MOLL, AND G. NATVIK (2022): “Asset-Price Redistribution,” Tech. rep., Working paper, LSE.
- FAIA, E. AND V. PEZONE (2023): “The Cost of Wage Rigidity,” *The Review of Economic Studies*.
- FANG, X., Y. LIU, AND N. ROUSSANOV (2022): “Getting to the Core: Inflation Risks Within and Across Asset Classes,” Working Paper 30169, National Bureau of Economic Research.
- FLOOD, S. M. AND J. PACAS (2008): “Using the Annual Social and Economic Supplement as Part of a Current Population Survey Panel,” *Census Bureau Working Paper*, 1–36.
- GAGLIARDONE, L. AND M. GERTLER (2023): “Oil Prices, Monetary Policy and Inflation Surges,” Tech. rep., Working paper.
- GERTLER, M. AND P. KARADI (2015): “Monetary Policy Surprises, Credit Costs, and Economic Activity,” *American Economic Journal: Macroeconomics*, 7, 44–76.
- GILCHRIST, S. AND E. ZAKRAJŠEK (2012): “Credit spreads and business cycle fluctuations,” *American Economic Review*, 102, 1692–1720.
- GLOVER, A., J. HEATHCOTE, D. KRUEGER, AND J. V. RÍOS-RULL (2020): “Intergenerational redistribution in the great recession,” *Journal of Political Economy*, 128, 3730–3778.
- GREENWOOD, B. J., Z. HERCOWITZ, AND G. W. HUFFMAN (1988): “Investment, Capacity Utilization and the Real Business Cycle,” *The American Economic Review*, 78, 402–417.
- GRIGSBY, J. (2022): “Skill Heterogeneity and Aggregate Labor Market Dynamics,” Working Paper 30052, National Bureau of Economic Research.
- GRIGSBY, J., E. HURST, A. YILDIRMAZ, AND Y. ZHESTKOVA (2021): “Nominal Wage Adjustments During the Pandemic Recession,” *American Economic Review: Papers & Proceedings*, 111.
- GÜRKAYNAK, R., H. G. KARASOY-CAN, AND S. S. LEE (2022): “Stock Market’s Assessment of Monetary Policy Transmission: The Cash Flow Effect,” *The Journal of Finance*, 77, 2375–2421.

- GÜRKAYNAK, R., B. SACK, AND E. 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.
- JARAVEL, X. (2019): “The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector,” *Quarterly Journal of Economics*, 134, 715–783.
- (2021): “Inflation Inequality: Measurement, Causes, and Policy Implications,” *Annual Review of Economics*, 13, 599–629.
- JARAVEL, X. AND D. LASHKARI (2022): “Nonparametric Measurement of Long-Run Growth in Consumer Welfare,” *Working Paper*.
- JENTSCH, C. AND K. G. LUNSFORD (2019): “The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States: Comment,” *American Economic Review*, 109, 2655–78.
- KÄNZIG, D. R. (2021): “The macroeconomic effects of oil supply news: Evidence from opec announcements,” *American Economic Review*, 11, 1092–1125.
- KAPLAN, G. AND S. SCHULHOFER-WOHL (2017): “Inflation at the household level,” *Journal of Monetary Economics*, 91, 19–38, the Swiss National Bank/Study Center Gerzensee Special Issue: “Modern Macroeconomics: Study Center Gerzensee Conference in Honor of Robert G. King” Sponsored by the Swiss National Bank and the Study Center Gerzensee.
- KILIAN, L. AND D. P. MURPHY (2014): “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil,” *Journal of Applied Econometrics*, 29.
- KIM, R. AND J. VOGEL (2020): “Trade and Welfare (Across Local Labor Markets),” Working Paper 27133, National Bureau of Economic Research.
- LEE, M., C. MACALUSO, AND F. SCHWARTZMAN (2022): “Minority Unemployment, Inflation, and Monetary Policy,” Working Paper.
- MCKAY, A. AND C. WOLF (2023a): “Optimal Policy Rules in HANK,” .
- (2023b): “What Can Time-Series Regressions Tell Us About Policy Counterfactuals?” .
- MCKAY, A. AND C. K. WOLF (2023c): “Monetary Policy and Inequality,” *Journal of Economic Perspectives*, 37.
- MOENCH, E., J. I. VICKERY, AND D. ARAGON (2010): “Why is the Market Share of Adjustable-Rate Mortgages so Low?” *New York Fed Current Issues in Economics and Finance*, 16, 1–11.
- OBERFIELD, E. (2023): “Inequality and Measured Growth,” Working paper.
- ORCHARD, J. (2022): “Cyclical Demand Shifts and Cost of Living Inequality,” Working Paper.
- OYAMA, D. AND T. TAKENAWA (2018): “On the (Non-)Differentiability of the Optimal Value Function When the Optimal Solution Is Unique,” *Journal of Mathematical Economics*, 76, 21–32.

- PLAGBORG-MØLLER, M. AND C. K. WOLF (2021): "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 89, 955–980.
- RAMEY, V. (2016): "Macroeconomic Shocks and Their Propagation," Elsevier, vol. 2, chap. Chapter 2, 71–162.
- SIDRAUSKI, M. (1967): "Rational Choice and Patterns of Growth in a Monetary Economy," *American Economic Review*, 57, 534–544.
- OLON, G., R. BARSKY, AND J. PARKER (1994): "Measuring the Cyclicity of Real Wages: How Important is Composition Bias," *Quarterly Journal of Economics*, 109, 1–25.
- STOCK, J. H. AND M. W. WATSON (2018): "Identification and Estimation of Dynamic Causal Effects in Macroeconomics using External Instruments," *The Economic Journal*, 128, 917–948.
- TZAMOURANI, P. (2021): "The Interest Rate Exposure of Euro Area Households," *European Economic Review*, 132, 103643.
- VAN DEN HEUVEL, S. J. (2008): "The welfare cost of bank capital requirements," *Journal of Monetary Economics*, 55, 298–320.
- YANG, Y. (2022): "Redistributive Inflation and Optimal Monetary Policy," Working paper.

A. THEORY APPENDIX

A.1 Proof of Proposition 1

This result is a straightforward application of the envelope theorem. Suppressing dependence on group and age for notational parsimony, the Lagrangian associated with (2) is

$$\begin{aligned} \mathcal{L} = \sum_t \beta^t \delta_t & \left(U(C_t, \{N_{kt}\}_k, L_t) - \lambda_t \left[\sum_j p_{jt} c_{jt} \right. \right. \\ & \left. \left. - \sum_k [N_{kt-1} D_{kt} - Q_{kt} \Delta N_{kt} - \chi_k(\Delta N_{k,t})] - W_t L_t - T_t \right] \right). \end{aligned}$$

The envelope theorem gives the response of welfare to a perturbation affecting prices, wages and transfers as

$$\begin{aligned} \frac{dV}{d \ln z} = \sum_t \beta^t \delta_t \lambda_t & \left[- \sum_j p_{jt} c_{jt} \frac{d \ln p_{jt}}{d \ln z} + W_t L_t \frac{d \ln W_t}{d \ln z} \right. \\ (A1) \quad & \left. + \sum_k \left(N_{kt-1} D_{kt} \frac{d \ln D_{kt}}{d \ln z} - Q_{kt} \Delta N_{kt} \frac{d \ln Q_{kt}}{d \ln z} \right) + T_t \frac{d \ln T_t}{d \ln z} \right]. \end{aligned}$$

The first order condition for the one period bond is

$$\lambda_t Q_{0t} = \Delta \hat{\delta}_t \beta \lambda_{t+1},$$

where $\hat{\delta}_t \equiv \delta_{t+1}/\delta_t$. Inserting these into (A1) (and assuming $\delta_0 = 1$) gives

$$\begin{aligned} \frac{dV}{d \ln z} = \lambda_0 \sum_t \prod_{\tau=0}^t Q_{0\tau} & \left[- \sum_j p_{jt} c_{jt} \frac{d \ln p_{jt}}{d \ln z} + W_t L_t \frac{d \ln W_t}{d \ln z} \right. \\ (A2) \quad & \left. + \sum_k \left(N_{kt-1} D_{kt} \frac{d \ln D_{kt}}{d \ln z} - Q_{kt} \Delta N_{kt} \frac{d \ln Q_{kt}}{d \ln z} \right) + T_t \frac{d \ln T_t}{d \ln z} \right]. \end{aligned}$$

The price of the riskless bond is $Q_{0t} = 1/R_t$. Substituting this into the above equation yields the result.

A.2 Proof of Proposition 2

Here we prove Proposition 2: the stochastic version of our framework. For this, recall that we assume prices, wages and dividends follow stochastic processes given by

$$(A3) \quad D_{kt} = \bar{D}_{kt} \exp(v_{kt}^D)^\sigma, \quad Q_{kt} = \bar{Q}_{kt} \exp(v_{kt}^Q)^\sigma, \quad p_{jt} = \bar{p}_{jt} \exp(v_{jt}^p)^\sigma, \quad W_t^a = \bar{W}_t \exp(v_t^W)^\sigma,$$

where, given matrices $\theta(L)$,

$$v_{kt}^D = \theta_k^D(L)\epsilon_t, \quad v_{kt}^Q = \theta_k^Q(L)\epsilon_t, \quad v_{jt}^p = \theta_j^p(L)\epsilon_t, \quad v_t^W = \theta^W(L)\epsilon_t,$$

Letting s_t denote the state of the world in period t , summarizing realizations of ϵ , write the household problem in sequence form:

$$(A4) \quad V(\sigma, \epsilon_0) = \max_{\{c_{jt}(s_t)\}_j, C_t(s_t), \{N_{kt}(s_t)\}_k}_{t=0}^{\infty} \sum_{t=0}^{\infty} \beta^t \delta_t \sum_{s_t} \pi_t(s_t) U(C_t(s_t), \{N_{kt}(s_t)\}_k, L_t(s_t))$$

where $\pi_t(s_t)$ is the probability of realizing state s_t in period t , subject to state-by-state budget constraints for all t

$$\begin{aligned} \sum_j p_{jt}(s_t) c_{jt}(s_t) &= \sum_k [N_{kt-1}(s_{t-1}) D_{kt}(s_t) - Q_{kt}(s_t) \Delta N_{kt}(s_t) - \chi_k(\Delta N_{k,t}(s_t))] \\ &\quad + W_t(s_t) L_t(s_t) + T_t(s_t), \end{aligned}$$

the consumption aggregator in (1), an initial set of assets $\{N_{k0}\}_k$ and a set of no-Ponzi conditions

$$\lim_{T \rightarrow \infty} \sum_{s_T} R_{0 \rightarrow T}^{-1} \pi_T(s_T) Q_{kT}(s_T) N_{kT}(s_T) = 0 \quad \forall k, t.$$

Note we have suppressed the age and group superscripts for notational convenience. Note further that we have written the household's value function such that it depends on the initial realization of ϵ_0 . This problem is parameterized by $\sigma \in [0, 1]$, which indexes a perturbation from a completely deterministic economy. We approximate $V(\sigma)$ using a Taylor approximation around $\sigma = 0$:

$$(A5) \quad V(\sigma, \epsilon_0) \approx V(0) + \frac{dV(0, \epsilon_0)}{d\sigma} \sigma.$$

Observe that if $\sigma = 0$, $V(\sigma, \epsilon)$ does not depend on ϵ because then shocks have zero variance and hence, from (A3), prices are the same in all realisation of ϵ . Now we require an expression for $dV(0, \epsilon_0)/d\sigma$. To that end, note the Lagrangian associated with (A4) is

$$\begin{aligned} \mathcal{L} &= \sum_t \beta^t \delta_t \sum_{s_t} \pi_t(s_t) \left(U(C_t(s_t), \{N_{kt}\}_k, L_t) - \lambda(s_t) \left[\sum_j p_{jt}(s_t) c_{jt}(s_t) \right. \right. \\ &\quad \left. \left. - \sum_k [N_{kt-1}(s_{t-1}) D_{kt}(s_t) - Q_{kt}(s_t) \Delta N_{kt}(s_t) - \chi_k(\Delta N_{k,t}(s_t))] \right. \right. \\ &\quad \left. \left. - W_t(s_t) L_t(s_t) - T_t(s_t) \right] \right). \end{aligned}$$

Plugging in for the stochastic processes for prices, wages, dividends and transfers, and taking

the derivative of this Lagrangian with respect to σ gives an expression for $dV(\sigma, \epsilon_0)/d\sigma$:

$$(A6) \quad \frac{dV}{d\sigma} = \frac{\partial \mathcal{L}}{\partial \sigma} = \sum_t \beta^t \delta_t \sum_{s_t} \pi_t(s_t) \lambda_t(s_t) \left(- \sum_j p_{j,t}(s_t) c_{jt}(s_t) v_{jt}^p + W_t(s_t) v_t^W + T_t(s_t) v_t^T \right. \\ \left. + \sum_k \left[N_{kt-1}(s_{t-1}) D_{kt}(s_t) v_{kt}^D - Q_{kt}(s_t) \Delta N_{kt}(s_t) v_{kt}^Q \right] \right).$$

Note that the first equality invokes the envelope theorem of Oyama and Takenawa (2018) so that total derivatives may become partial derivatives. The value of σ affects this expression implicitly through the stochastic processes.

The above expression contains a Lagrange multiplier. To account for this, note that the first order condition for the riskless bond, which does not enter utility and is subject to no adjustment costs by assumption, is

$$\lambda_t(s_t) Q_{Rt}(s_t) = \hat{\delta}_{t+1} \beta \mathbb{E}[\lambda_{t+1}(s_{t+1}) | s_t].$$

We take the limit of this expression as $\sigma \rightarrow 0$, which becomes

$$(A7) \quad \lambda_t Q_{Rt} = \hat{\delta}_{t+1} \lambda_{t+1}.$$

Plugging equation (A7) into equation (A6) and defining the operator $\mathbb{E}_0[x_t]$ to be the expectation of a variable x taken over possible realizations of s_t given the time zero probabilities, one can see

$$(A8) \quad \frac{dV(0, \epsilon_0)}{d\sigma} = \lambda_0 \sum_t R_{0 \rightarrow t}^{-1} \left(- \sum_j p_{j,t} c_{jt} \mathbb{E}_0[v_{jt}^p] + W_t \mathbb{E}_0[v_t^W] + T_t \mathbb{E}_0[v_t^T] \right. \\ \left. + \sum_k \left[N_{kt-1} D_{kt} \mathbb{E}_0[v_{kt}^D] - Q_{kt} \Delta N_{kt} \mathbb{E}_0[v_{kt}^Q] \right] \right).$$

Now define the change in welfare from an impulse to element n of the structural shock vector, for any value of σ , as

$$dV \equiv V(\sigma, \epsilon_0^n = 1, \cdot) - V(\sigma, \epsilon_0^n = 0, \cdot).$$

where \cdot includes all other values of ϵ . Using (A5), note that the $V(0)$ term drops out of this expression, so that:

$$dV = \left(\frac{dV(0, \epsilon_0^n = 1, \cdot)}{d\sigma} - \frac{dV(0, \epsilon_0^n = 0, \cdot)}{d\sigma} \right) \sigma.$$

Setting $\sigma = 1$ and plugging with equation (A8), this can be written as

$$dV = \lambda_0 \sum_t R_{0 \rightarrow t}^{-1} \left(- \sum_j p_{j,t} c_{jt} \Psi_{n,t}^{p,j} \right. \\ \left. + \sum_k \left[N_{kt-1} D_{kt} \Psi_{n,t}^{D,k} - Q_{kt} \Delta N_{kt} \Psi_{n,t}^{Q,k} \right] + W_t \Psi_{n,t+h}^W + T_t^a \Psi_{n,t}^T \right). \quad \square$$

B. DATA APPENDIX

This section describes the data used in our analysis in more detail. It describes the Consumer Expenditure Survey, Survey of Consumer Finances, Current Population Survey, Survey of Income and Program Participation, and National Household Travel Survey and outlines our approach to cleaning them. In addition, Table A1 presents a detailed age \times education distribution in our samples of each dataset.

B.1 Consumer Expenditure Survey (CEX)

Data on households' consumption is obtained from the Public Use Micro Data (PUMD) from Interview section of the CEX. This data is available directly from the Bureau of Labor Statistics (BLS) website. The Interview survey collects expenditures on goods and services, grouped into Universal Classification Codes (UCCs). Following (and augmenting) the crosswalk from Orchard (2022) we map the UCCs to 25 categories of consumption to compute quarterly household expenditures in each of these groups in 2019. The categories are: Food at home, Food away from home, Alcoholic beverages, Shelter, Fuels and utilities, Education, Apparel, New Vehicles, Used Vehicles, Other Vehicles, Motor fuel, Public transportation, Personal care, Motor vehicle insurance, Motor vehicle fees, Motor vehicle parts/equipment, Motor vehicle maintenance/repair, Medical care services, Recreation, Medical care commodities, Postage and delivery services, Information and information processing, Information technology, hardware/services, Tobacco and smoking products, and Household furnishings/operations.

We also extract total quarterly expenditures from the survey. For this particular calculation, we use an existing variable in the survey, which records the total expenditure of the household on each quarter. Then, we compute the average at the annual level to obtain mean household expenditures in 2019. In particular, this procedure allows us to remove the seasonality expected in quarterly expenditures.

For our purposes, we focus only on the sample of households whose reference person is between 25 and 80 years old, and top-code age at 75.³¹ Educational attainments for each reference person can also be obtained from the survey, which we use to define our three educational groups: HS or less, Some college, and Bachelor's +.

³¹We do this in order to improve our estimates for 75 years old households. The implicit assumption is that consumption patterns are similar between people aged between 75 and 80.

Combining consumption data and household characteristics we estimate average expenditures in each of our 25 categories, for all our demographic groups in 2019. Next, to minimize jumps in consumption patterns caused by measurement error, we run a Locally Weighted Scatterplot Smoothing (LOWESS) for each of the categories and each of the demographic groups. We also do this smoothing for the mean annual expenditures.³²

Finally, we also extract mortgage interest payments from the CEX. Households indicate how much they pay in mortgage interest each of the three previous months in the quarter they are interviewed. We sum these payments and record this quantity as the quarterly expenditure in interest payments by each household. We then follow similar smoothing and average procedures for this variable as we did with consumption patterns.

B.2 Survey of Consumer Finances (SCF)

Data on households' portfolio holdings and asset accumulation patterns is obtained from the SCF. We use both the Full Public Data Set, as well as the Summary Extract Public Data, which can be downloaded directly from the Federal Reserve webpage.

For equity holdings, we include directly held stocks, and indirectly holdings from mutual funds or retirement accounts. Similarly, for bonds we account for direct holdings, as well as contributions from mutual funds, annuities, trusts, and retirement accounts. In both cases, indirect holdings are estimated using the information on the percentage of the financial instrument invested in the corresponding asset class. For example, combination mutual fund holdings are split evenly between bonds and stocks. We further split bonds into corporate and non-corporate bonds. The former are obtained from the "Corporate and Foreign bonds" variable, while the latter are all other bonds.³³ For vehicles, we use the value of all owned vehicles from the Summary Extract. Similarly, for houses we use the value of the primary residence, also from the Summary Extract.

In terms of demographics, we mimic the definitions used for the consumption data. In particular, we focus on households whose reference person is between 25 and 80 years old, and top-code age at 75. The educational groups we define are the same as above: HS or less, Some college, and Bachelor's +. Consequently, we are able obtain average asset holdings for each of our asset classes at the age-education attainment level.

For our framework, we are interested in accumulation patterns and previous holdings by quarter. To this end, we start by linearly interpolating asset holdings at the quarter level. Then, we define the accumulation of each of our asset classes as the difference between the holdings at age a and the holdings at age $a - 1/4$. Previous holdings at age a are naturally defined as the holdings at age $a - 1/4$. Then, as we proceeded with consumption data, we run a LOWESS for each asset class over the lifecycle, to reduce jumps due to measurement error.

³²Unless otherwise stated, we use a smoothing bandwidth of 0.8 for all LOWESS procedures.

³³It might be possible that some portion of indirectly held bonds are corporate. However, the SCF does not allow us to know this.

In parallel, we also obtain the home ownership rate from the SCF. This variable is defined as the share of households in each age-education group that have positive housing holdings. As with the holdings of each asset class, we interpolate and smooth the home ownership share to avoid jumps due to measurement error. This variable weights the welfare effects of housing and renting within each demographic group. Explicitly, welfare effects from owner-occupied housing are multiplied by this rate, while welfare effects from rent are multiplied by one minus this rate.

We follow the IRS in supposing a house fully depreciates in 27.5 years. This yields an annual depreciation rate of 3.636% which we convert to quarterly by dividing by 4.

B.3 Current Population Survey (CPS)

Our data on labor income is from the Current Population Survey (CPS) provided by IPUMS. The CPS is a survey jointly sponsored by the U.S. Census Bureau and Bureau of Labor Statistics (BLS). It is designed to be nationally representative of the population and is used for a variety of official labor market statistics. Most famously, it is used to construct the civilian unemployment and labor force participation rates.

The CPS is a rotating panel of household addresses. Households are sampled for a period of four months, before being dropped from the sample for eight months, and included again for an additional four months. Thus a household may be included from January through April in 2005, excluded from May to December in 2005, and included again from January through April in 2006. Each of these four-month spells in the sample are known as “rotations.”

Households provide information on all household members. The “Basic” CPS, administered each month, contains information on demographics such as age, race and sex, as well as education, geography, employment status, occupation and industry. Household groups are defined by the age and education of their household head. We use the basic monthly CPS to construct estimates of unemployment rates for each of our three education groups, using the provided sampling weights. Since all our asset and consumption data are measured at the household level, we separately compute these unemployment rates for household heads and, if present, their spouse. These unemployment rates are used in the estimation of IRFs $\Psi_{i,t}^u$. We additionally compute 2019 employment-to-population ratios for each group g and age a , which we denote $e_0^{a,g}$.

In addition to the basic CPS, households are asked an additional set of questions in the final month of each rotation. This is known as either the “Outgoing Rotations Group” (ORG) or “Earner Study.” In this supplemental survey, households are asked whether they are paid hourly and, if so, the usual hours worked per week and their hourly wages. Wage/salary workers are additionally asked about their weekly earnings. To clarify our approach, we now write labor earnings as $W_{it}h_{it}e_{it}$, for h_{it} hours worked of employed workers and e_{it} an indicator for individual i being employed. That is, we have decomposed the quantity of labor L_{it} into an intensive margin (h_{it}) and extensive margin e_{it} . Proposition 2 requires both a time series of

wages on which to estimate impulse responses and an estimate of 2019 earnings $W_0^{a,g} h_0^{a,g} e_0^{a,g}$.

We proceed in two steps. First, we use the nominal weekly earnings variable *EARNWEEK* as our measure of earnings $W_{it} h_{it}$. We drop all individuals with missing earnings and estimate the following regression model for all individuals in 2019

$$(A9) \quad W_{it} h_{it} = \alpha_{a(i,t),g(i,t)} + \varepsilon_{it}$$

where $\alpha_{a,g}$ is an age \times education group fixed effect. These $\alpha_{a,g}$ serve as our estimate of $W_0^{a,g} h_0^{a,g}$. We then combine this estimate with the $e_0^{a,g}$ from the Basic CPS to get our estimate of 2019 earnings $W_0^{a,g} h_0^{a,g} e_0^{a,g}$.

Finally, we use the average of the residuals ε_{it} from regression (A9) as our time series of wages for group g in our estimation of Ψ_t^w . This average is taken only over those workers who are employed in both period t and either month $t + 12$ or $t - 12$. This focus on “job-stayers” reduces the role that composition effects plays in our wage index. To construct the time series variable, we use log weekly earnings as our measure of the wage in regression (A9). The residuals are calculated in every month of the sample, though the regression estimating education-specific life cycle profiles is run only including 2019 observations. The wage index therefore has the units of 2019 dollars. Thus, the residuals will include wage inflation driven by a rising price level. This justifies our inclusion of aggregate inflation as a control in our IRF estimation.

B.4 Survey of Income and Program Participation (SIPP)

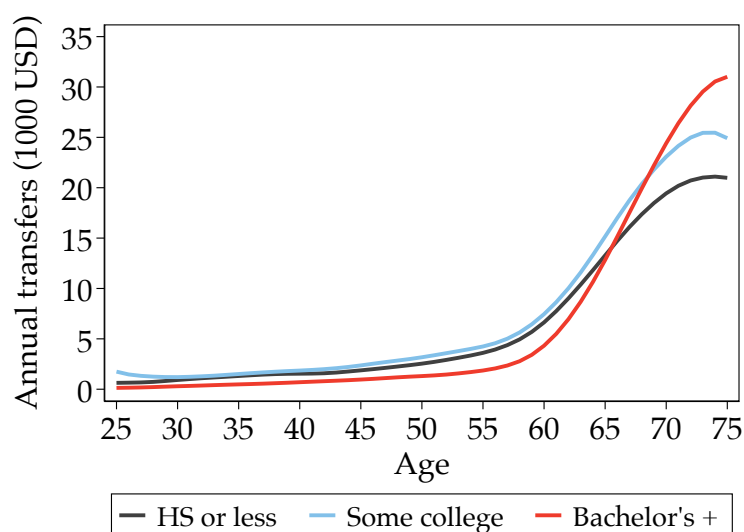
Data on income from transfers is obtained from the Survey of Income and Program Participation (SIPP). This survey collects information from different households on a monthly basis. Each panel is active for 4 consecutive years. For our purposes, we use the second wave of the 2018 SIPP Panel to obtain estimates for 2019. The data can be obtained from the US Census Bureau website.

As with consumption we focus on households whose reference person is between 25 and 80 years old, and top-code age at 75. Again, we also use educational attainments to identify our three educational groups: HS or less, Some college, and Bachelor’s +. From the survey, we compute total monthly income from transfers as the sum of means-tested transfer income and social insurance payments. The former component includes payments from the following means-tested programs: TANF, SSI, GA, veterans pension, and pass-through child support. The latter includes other payments from Veterans Affairs, Social Security, unemployment compensations, and G.I. Bill.³⁴ We then accumulate this income at the household level for each year, and finally compute the mean annual income from transfers for each of our demographic groups.³⁵ After computing this average, we estimate quarterly income from transfers and interpolate transfer income between quarters. Lastly, we smooth these transfer income patterns

³⁴Explicitly, we use the variables TPTRNINC and TPSCININC. A description of these variables as well as the sources of income in the SIPP can be found in the [SIPP webpage](#).

³⁵As suggested by the US Census Bureau, the annual average is computed using the weights of each household in December of the corresponding year. See the 2018 SIPP User’s guide.

FIGURE A1: Transfer income over the lifecycle



Notes: Figure shows annual income from transfers by group. Data is from the second wave of the 2018 panel of the Survey of Income and Program Participation. Income is averaged within group and age, and then a LOWESS smoother is applied across age.

over the lifecycle with a LOWESS smoother.³⁶ The result is shown in Figure A1.

B.5 National Household Travel Survey (NHTS)

Data on usage and characteristics of vehicles is obtained from the National Household Travel Survey (NHTS). This survey is conducted by the Federal Highway Administration and collects information on travel behaviors of US residents by all modes of transport and all purposes.³⁷ For our purposes we focus on the 2017 NHTS, the most recent survey.

As with the previous surveys, we focus on households whose reference person is between 25 and 80 years old, and top-code age at 75, and construct our educational attainment groups: HS or less, Some college, and Bachelor's +.³⁸ Of the vehicle related variables, we focus only on mileage: the annual number of miles driven per month of age of the car. We compute it using the `bestmile` variable in the NHTS, divided by the age of the vehicle in months.³⁹ We compute the average of this variable by each educational attainment group and then calculate the annual depreciation parameter due to vehicle use by expressing this average in kilometers per month of age and multiplying it by 0.000117.⁴⁰ Finally, we divide the depreciation parameter by 4 to obtain a quarterly estimate.

³⁶We use a LOWESS bandwidth of 0.5 for transfer income in order to better capture the jump at 65.

³⁷More details can be found in the [NHTS website](#).

³⁸We use imputed age, instead of the reported one. However, 0.21% of observations in the person dataset of the NHTS differ between reported and imputed age. We drop these observations.

³⁹The `bestmile` variable is an alternative measure of annual miles that accounts for vehicles that do not have a readable odometer or for which no self-report is provided. Details about the methodology used in the NHTS to obtain the variable can be found in the [NHTS documentation](#).

⁴⁰This is the estimate for the relative mileage effect on car price reported in Figure 5 of Dexheimer (2003)

TABLE A1: Detailed Demographic Statistics: Cross-Sectional Survey Data

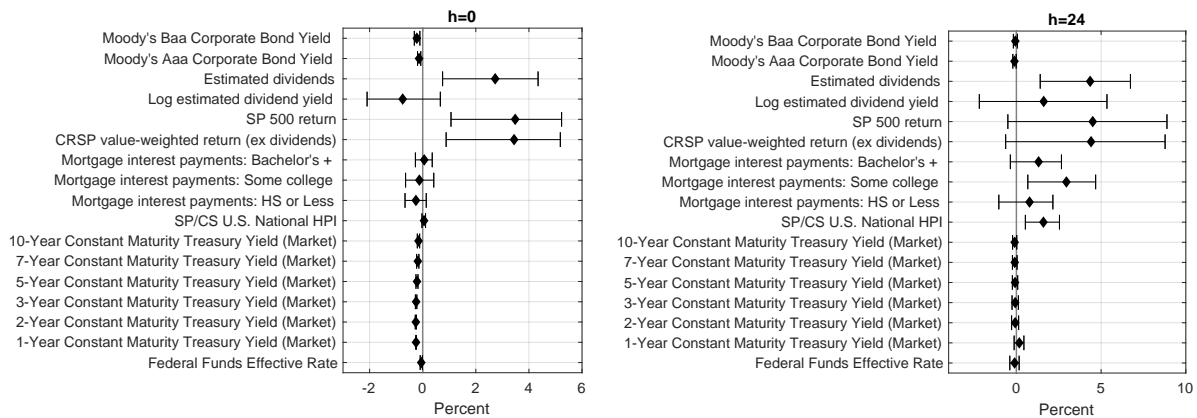
	HS or Less	Some College	College+	Full Sample
<i>Panel (a): Consumer Expenditure Survey (CEX)</i>				
25-34 y.o. (%)	15.48	17.30	21.19	18.22
35-44 y.o. (%)	17.33	19.49	20.23	19.09
45-54 y.o. (%)	18.24	18.99	19.73	19.04
55+ y.o. (%)	48.94	44.22	38.85	43.65
<i>Panel (b): Current Population Survey (CPS)</i>				
25-34 y.o. (%)	15.1	17.8	19.9	17.7
35-44 y.o. (%)	16.1	17.6	20.6	18.3
45-54 y.o. (%)	17.5	18.6	19.5	18.6
55+ y.o. (%)	51.3	46.0	40.0	45.4
<i>Panel (c): Survey of Consumer Finances (SCF)</i>				
25-34 y.o. (%)	15.74	18.75	18.66	17.68
35-44 y.o. (%)	17.36	16.92	20.30	18.33
45-54 y.o. (%)	18.56	20.29	18.21	18.92
55+ y.o. (%)	48.34	44.04	42.83	45.07
<i>Panel (d): Survey of Income Program Participation (SIPP)</i>				
25-34 y.o. (%)	15.46	17.63	23.90	19.44
35-44 y.o. (%)	16.04	19.12	20.90	18.85
45-54 y.o. (%)	18.66	19.37	19.44	19.17
55+ y.o. (%)	50.18	44.67	36.56	43.18

Notes: Age and education correspond to that of the household head in every sample. All numbers average over all of 2019. CPS data correspond to the outgoing rotation groups (ORG) sample of the CPS. Only households whose head is at least 25 years old are included.

C. ADDITIONAL RESULTS

This section presents a number of additional results. It first presents full impulse response functions for asset prices and coefficient plots at different horizons. It then reports money metric welfare effects of oil supply and monetary shocks (compared with welfare effects normalized by consumption in the main text). Finally, it presents the transfer channel of welfare effects and life cycles of asset accumulation for total assets.

FIGURE A2: Estimated Response of Disaggregated Asset Prices to a 25 basis point decline in the 1-year treasury yield: Impact and two years out

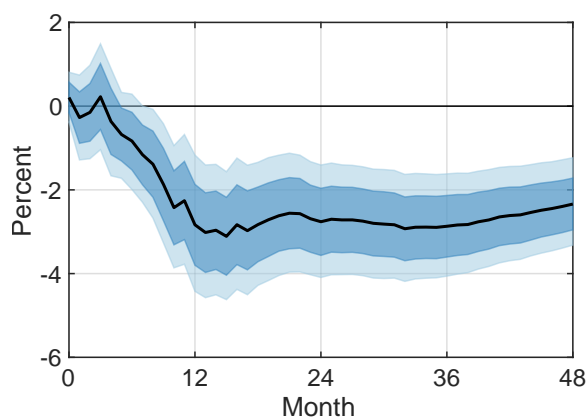


PANEL A: IMPACT RESPONSE

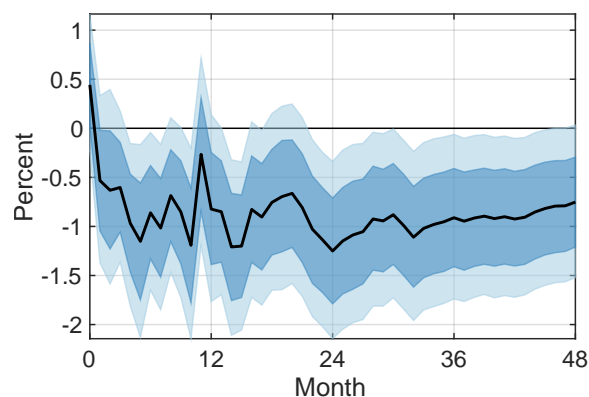
PANEL B: RESPONSE AFTER 24 MONTHS

Notes: Figure plots impulse response functions (IRFs) of asset prices and dividend yields to inflationary monetary shocks constructed following Gertler and Karadi (2015). Shocks normalized to represent a 25 basis point cut in one-year treasury yields caused by FOMC announcements. IRFs estimated using the “internal instrument” SVAR procedure explained in section 4. Panel A plots the IRF on impact (i.e. at 0-horizon) of each asset good, while Panel B plots the cumulative IRF 24 months after the shock impulse. The SVAR is specified with 12 lags. Error bars represent 90% confidence intervals.

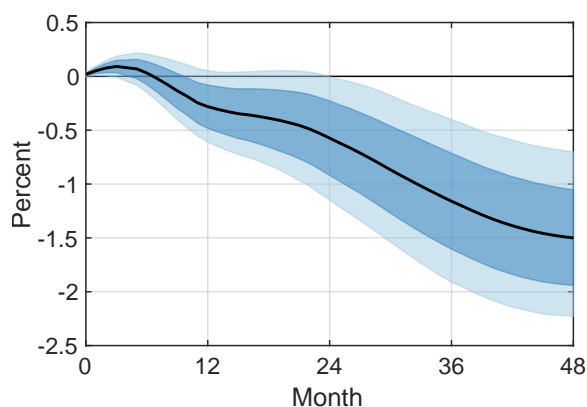
FIGURE A3: Estimated Impulse Response to a 10% increase in crude oil prices: Assets



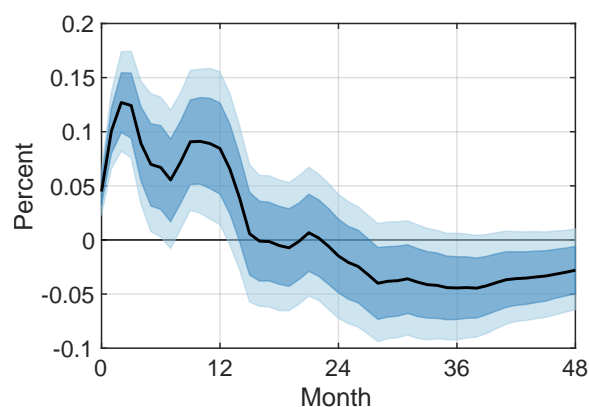
PANEL A: CUMULATIVE STOCK MARKET RETURN, EXCL. DIVIDENDS



PANEL B: DIVIDEND PAYOUTS



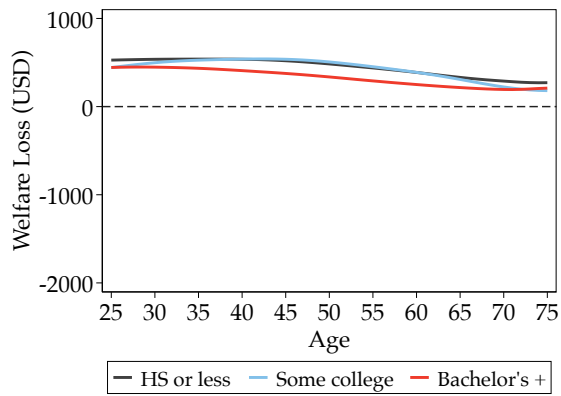
PANEL C: CASE-SHILLER HOME PRICE INDEX (HPI)



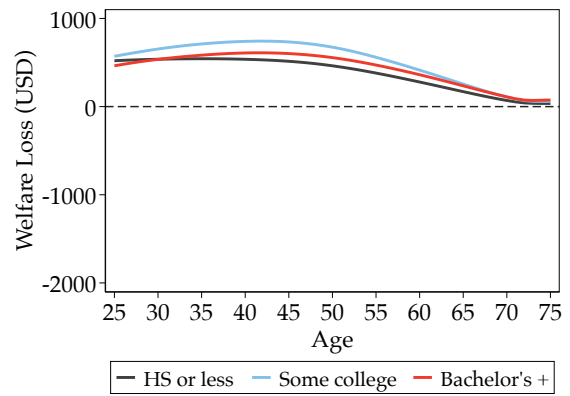
PANEL D: MOODY'S AAA CORPORATE BOND YIELDS

Notes: Figure plots impulse response functions (IRFs) of asset prices and dividend yields to inflationary oil supply news shocks constructed by Känzig (2021). Shocks normalized to represent a 10% increase in the West Texas Intermediates Crude Oil price driven by announced reductions in OPEC oil supply. IRFs estimated using the “internal instrument” SVAR procedure explained in section 4. Panel A plots the IRF of the S&P500 stock return, excluding dividends. Panel B reports dividend payouts from the S&P500. Panel C plots the response of the Case-Shiller Home Price Index (HPI). Panel D reports the response of the Moody’s Aaa Corporate Bond Yield. The SVAR is specified with 12 lags. Error bars represent 90% confidence intervals.

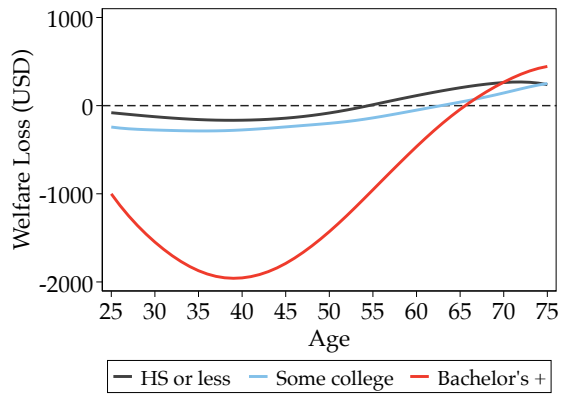
FIGURE A4: Money-Metric Welfare Loss of Inflationary Oil Price Shocks over the Life Cycle and by Household Education



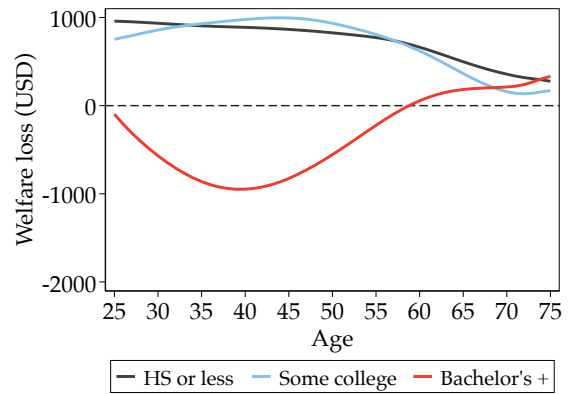
PANEL A: CONSUMPTION CHANNEL



PANEL B: LABOR INCOME CHANNEL



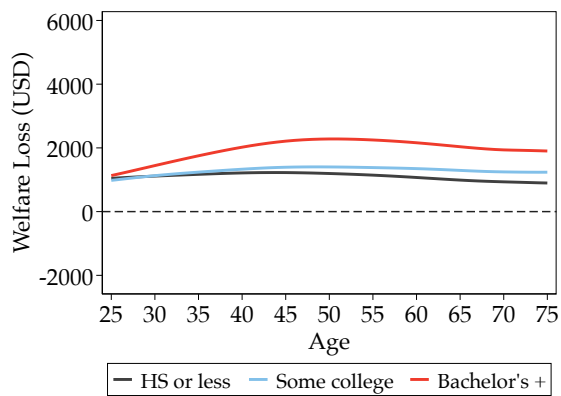
PANEL C: PORTFOLIO CHANNEL



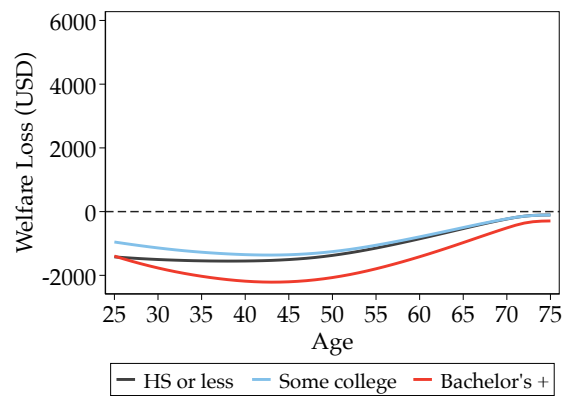
PANEL D: TOTAL WELFARE CHANGE

Notes: Figure shows the estimated money-metric welfare effects of a 10% increase in the West Texas Intermediates Crude Oil price driven by announced reductions in OPEC oil supply. Panels A-C split the effect into the consumption channel, the labor income channel. A negative number represents a welfare gain.

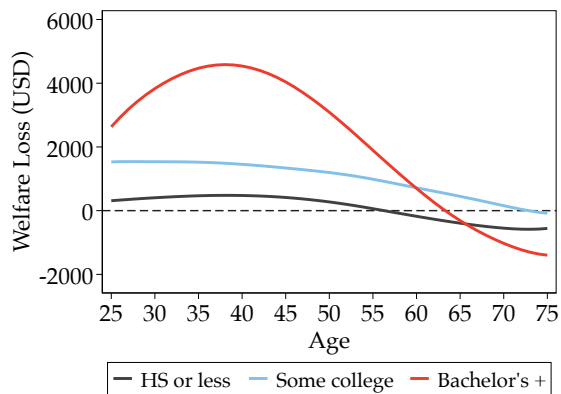
FIGURE A5: Money-Metric Welfare Loss of Inflationary Monetary Shocks over the Life Cycle and by Household Education



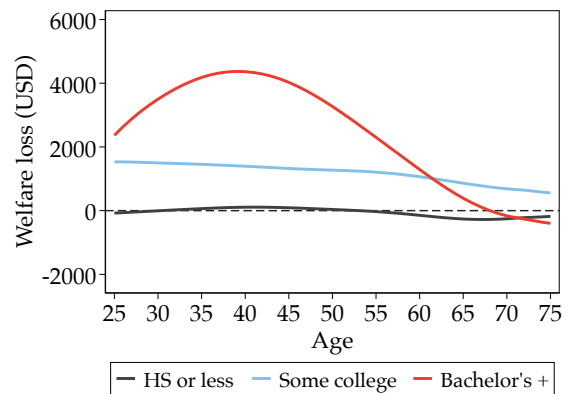
PANEL A: CONSUMPTION CHANNEL



PANEL B: LABOR INCOME CHANNEL



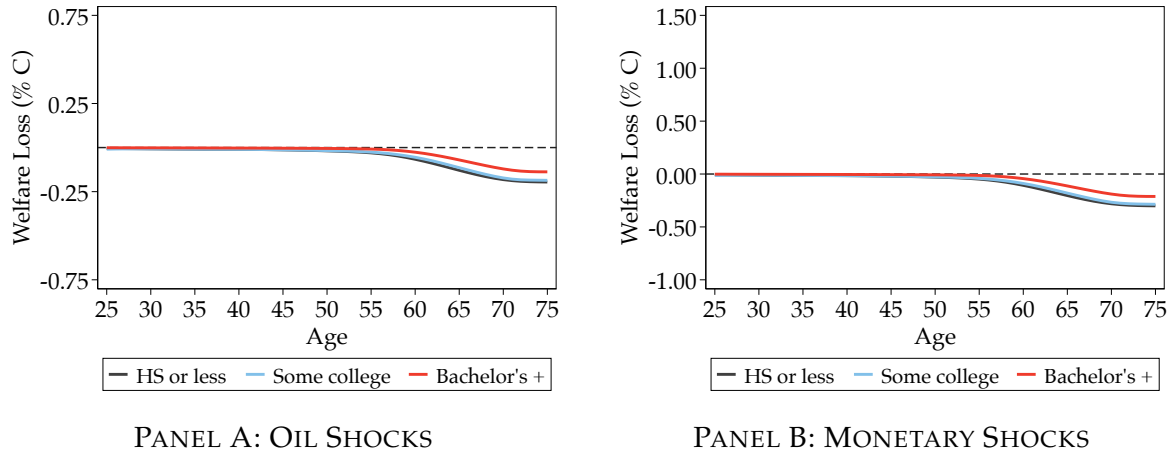
PANEL C: PORTFOLIO CHANNEL



PANEL D: TOTAL WELFARE CHANGE

Notes: Figure shows the estimated money-metric welfare effects of a 25 basis point cut to the federal funds rate. Panels A-C split the effect into the consumption channel, the labor income channel. A negative number represents a welfare gain.

FIGURE A6: Government Transfer Channel of an Oil Price and Monetary Shocks



Notes: Figure reports the welfare loss arising from changes in government transfer income that result from inflationary oil price shocks (Panel A) or monetary shocks (Panel B). Negative numbers represent welfare gains. Oil shocks constructed by Känzig (2021) and scaled to represent 10% movements in the real WTI crude oil price. Monetary shocks constructed following Gertler and Karadi (2015) and scaled to represent a 25 basis point reduction in one-year treasury yields. Vertical axis is normalized to be a share of consumption.

FIGURE A7: Life Cycle Accumulation of all Assets

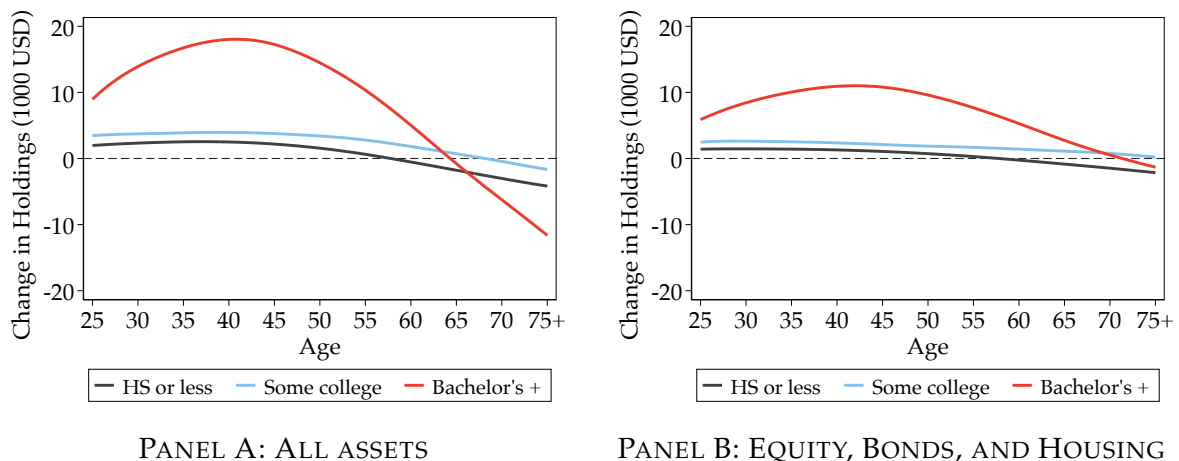


Figure reports the year-over-year accumulation in assets across the life cycle for our three education groups. Vertical axis scaled to be in units of thousands of 2019 dollars. All assets include equity, (corporate and non-corporate) bonds, housing, vehicles, liquid assets, business wealth, and other financial and non-financial assets.

D. ROBUSTNESS TESTS

D.1 Constraints

Net Worth Constraint. In this section we show how our results change if we have an additional, state-by-state net worth constraint. The proof follows that in A closely. Suppose that the problem of the household is to solve

$$(A10) \quad V(\sigma) = \max_{\{\{c_{jt}(s_t)\}_j, L_t(s_t), \{N_{kt}(s_t)\}_k\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \delta_t \mathbb{E}_t[U(C_t(s_t), \{N_{kt}(s_t)\}, L_t(s_t))]$$

subject to state by state budget constraints for all t ,

$$\begin{aligned} \sum_j p_{jt}(s_t) c_{jt}(s_t) &= \sum_k [N_{kt-1}(s_{t-1})] D_{kt}(s_t) - Q_{kt}(s_t) \Delta N_{kt}(s_t) - \chi(\Delta N_{k,t}(s_t)) \\ &\quad + W_t(s_t) L_t(s_t) + T_t(s_t), \end{aligned}$$

state by state net-worth constraints,

$$\sum_k Q_{kt}(s) N_{kt}(s) \geq \underline{b},$$

the consumption aggregator in (1), an initial set of assets $\{N_{k0}\}$ and a set of no-Ponzi conditions

$$\lim_{T \rightarrow \infty} \mathbb{E}_T R_{0 \rightarrow T}^{-1} N_{kT} Q_{kT} \geq 0 \quad \forall k.$$

The Lagrangian associated with (A10) is

$$\begin{aligned} \mathcal{L} &= \sum_t \beta^t \delta_t \sum_{s_t} \pi_t(s_t) \left(U((C_t(s_t), \{N_{kt}(s_t)\}, L_t(s_t)) \right. \\ &\quad - \lambda_t(s_t) \left[\sum_j p_{jt}(s_t) c_{jt}(s_t) - \sum_k [N_{kt-1}(s_{t-1})] D_{kt}(s_t) - Q_{kt}(s_t) \Delta N_{kt}(s_t) - \chi_k(\Delta N_{k,t}(s_t)) \right. \\ &\quad \left. \left. - W_t(s_t) L_t(s_t) - T_t(s_t) \right] - \mu_t(s_t) \left[\bar{b} - \sum_k Q_{kt}(s_t) N_{kt}(s_t) \right] \right). \end{aligned}$$

where $\lambda_t(s_t)$ and $\mu_t(s_t)$ are the Lagrange multipliers on the period t budget and net worth constraints, respectively, given a realization of shocks s_t . The first order conditions for the riskless bond in state s_t are

$$N_1 : \quad \lambda_t Q_{1t}(s_t) = \beta \hat{\delta}_{t+1} \mathbb{E}_0[\lambda_{t+1}(s_{t+1})] + \mu_t(s_t) Q_{1t}(s_t)$$

We again pursue a small noise expansion. We can write the derivative of the value function

with respect to σ as

$$\begin{aligned} \frac{dV(\sigma)}{d\sigma} = \frac{\partial \mathcal{L}}{\partial \sigma} = \sum_t \beta^t \delta_t \sum_{s_t} \pi_t(s_t) & \left(\lambda_t(s_t) \left[- \sum_j p_{j,t}(s_t) c_{jt}(s_t) v_{jt}^p(s_t) \right. \right. \\ & + \sum_k \left[N_{kt-1}(s_{t-1}) D_{kt}(s_t) v_{kt}^D(s_t) - Q_{kt}(s_t) \Delta N_{kt}(s_t) v_{kt}^Q(s_t) \right] \\ & \left. \left. + W_t(s_t) L_t(s_t) v_t^W + T_t(s_t) v_t^T \right] + \mu_t(s_t) \sum_k Q_{kt}(s_t) N_{kt}(s_t) v_{kt}^Q(s_t) \right). \end{aligned}$$

and as $\sigma \rightarrow 0$ this approaches

$$\begin{aligned} \frac{dV(0)}{d\sigma} = \sum_t \beta^t \delta_t \sum_{s_t} \pi_t(s_t) & \left(\lambda_t \left[- \sum_j p_{j,t} c_{jt} v_{jt}^p(s_t) \right. \right. \\ & + \sum_k \left[N_{kt-1} D_{kt} v_{kt}^D(s_t) - Q_{kt} \Delta N_{kt} v_{kt}^Q(s_t) \right] \\ (A11) \quad & \left. \left. + W_t L_t v_t^W(s_t) + T_t v_t^T(s_t) \right] + \mu_t \sum_k Q_{kt} N_{kt} v_{kt}^Q(s_t) \right). \end{aligned}$$

or, combining probability measures $\pi_t(s_t)$ into expectation operators:

$$\begin{aligned} & = \sum_t \beta^t \delta_t \left(\lambda_t \left[- \sum_j p_{j,t} c_{jt} \mathbb{E}_0[v_{jt}^p] + W_t L_t \mathbb{E}_0[v_t^W] + T_t \mathbb{E}_0[v_t^T] \right] \right. \\ & \left. + \sum_k \left[N_{kt-1} D_{kt} \mathbb{E}_0[v_{kt}^D] - Q_{kt} \Delta N_{kt} \mathbb{E}_0[v_{kt}^Q] + \mu_t \sum_k Q_{kt} N_{kt} \mathbb{E}_0[v_{kt}^Q] \right] \right). \end{aligned}$$

where we have invoked the envelope theorem to abstract from movements in choice variables.

Now write the FOC for the riskless bond as $\sigma \rightarrow 0$ as

$$\lambda_t(1 - \tilde{\mu}_t) = \beta \hat{\delta}_{t+1} R_t \lambda_{t+1}$$

where $\tilde{\mu}_t \equiv \mu_t / \lambda_t$. We can then write

$$\begin{aligned} \frac{dV(0)}{d\sigma} = \lambda_0 \sum_t R_{0 \rightarrow t}^{-1} \prod_{s=0}^t (1 + \tau_s)^{-1} & \left(- \sum_j p_{j,t} c_{jt} \mathbb{E}_t[v_{jt}^p] + W_t L_t \mathbb{E}_0[v_t^W] + T_t \mathbb{E}_0[v_t^T] \right. \\ & \left. + \sum_k \left[N_{kt-1} D_{kt} \mathbb{E}_0[v_{kt}^D] - Q_{kt} \Delta N_{kt} \mathbb{E}_0[v_{kt}^Q] + \tilde{\mu}_t \sum_k Q_{kt} N_{kt} \mathbb{E}_0[v_{kt}^Q] \right] \right). \end{aligned}$$

for $1 + \tau_t \equiv (1 - \tilde{\mu}_t)^{-1}$ is a ‘‘wedge’’ in the Euler Equation, in the spirit of Chari, Kehoe, and McGrattan (2007). Lastly, using a first order approximation to the value function around $\sigma = 0$, we get the change in welfare from an impulse at time zero to element n of the fundamental shock vector ϵ as

$$\begin{aligned}
dV &\equiv V(1, \epsilon_0^n = 1) - V(1, \epsilon_0^n = 0) \approx \frac{dV(0, \epsilon_0^n = 1, \cdot)}{d\sigma} - \frac{dV(0, \epsilon_0^n = 0, \cdot)}{d\sigma} \\
&= \lambda_0 \sum_t R_{0 \rightarrow t}^{-1} \prod_{s=0}^t (1 + \tau_s)^{-1} \left(\left[- \sum_j p_{j,t} c_{jt} \Psi_{n,t}^{p,j} + \sum_k \left[N_{kt-1} D_{kt} \Psi_{n,t}^{D,k} - Q_{kt} \Delta N_{kt} \Psi_{n,t}^{Q,k} \right] \right. \right. \\
\text{(A12)} \quad &\left. \left. + W_t L_t \Psi_{n,t}^W(s_t) + T_t \Psi_{n,t}^T \right] + \tilde{\mu}_t \sum_k Q_{kt} N_{kt} \Psi_{n,t}^{Q,k} \right).
\end{aligned}$$

Incorporating borrowing constraints thus has two effects relative to our baseline framework. First, it changes the effective discount rate through Euler equation wedges τ . Second, it introduces a term representing the relaxation of the borrowing constraint caused by asset price movements induced by the shock. This relaxation is valued by the Lagrange multiplier $\tilde{\mu}_t = 1 - (1 + \tau_t)^{-1}$. Thus, given an estimate of τ_t (which implies a value of $\tilde{\mu}_t$), one can account for borrowing constraints within our framework.

If one is willing to impose an assumption on the utility function one can use the Euler Equation to get an estimate of τ_t . For instance, suppose that $U_C(\cdot) = C^{-1}$, so that utility of consumption is separable from utility derived from asset holdings or leisure and has a log form. Furthermore, suppose that the consumption aggregator is homothetic. In this case, the Euler Equation may be written as

$$\text{(A13)} \quad \frac{P_{t+1} C_{t+1}}{P_t C_t} = \beta \frac{\delta_{t+1}}{\delta_t} \cdot R_t (1 + \tau_t)$$

Thus, given estimates of the nominal interest rate R_t , effective discount rates β , death rates $\hat{\delta}_{t+1}$ and growth in consumption expenditures, one can form an estimate of τ_t as needed.

Figure A8 reports the estimated wedges in the Euler equation τ over the life cycle for our three education groups. Higher values of τ indicate more binding constraints. We estimate that young highly-educated households are borrowing constrained, reflecting their rapid growth in consumption expenditures. For all education groups, the wedge in the Euler equation falls throughout working age. The elderly also appear to have a larger wedge in their Euler equation, possibly through misperceptions about death probabilities. For this reason, we only consider those aged 25-65 in our robustness exercises.

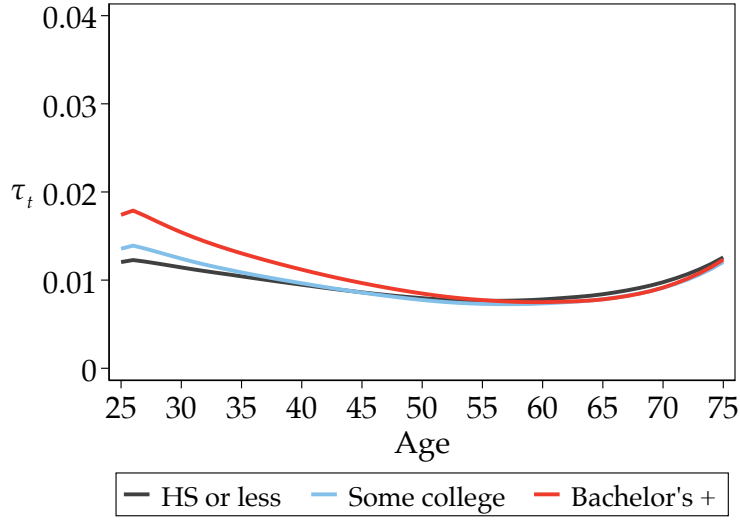
Short-selling Constraints. It is straightforward to introduce additional short-selling constraints on assets $k \neq 0$. Such constraints are of the form

$$\text{(A14)} \quad Q_{kt} N_{kt} \geq 0$$

where we let μ_{kt} be the Lagrange multiplier on the short-selling constraint for asset k .⁴¹ Since, by assumption, this constraint does not affect the riskfree bond ($k = 0$), such constraints do not distort the Euler equation and so do not affect effective discount rates. However, following the

⁴¹Note further that, by dividing both sides by Q_{kt} , this formulation also accounts for constraints of the form $N_{kt} \geq 0$.

FIGURE A8: Estimated Euler Equation Wedges



Notes: Figure plots the estimated τ from (A13) in the CEX for our three groups of consumers by age. Death rates are taken from the Period Life Table for 2019 from the Social Security Administration.

same steps as above, one can show that constraints on short-selling leads to one extra term in the expression for utility response to shocks, given by

$$\mu_{kt} Q_{kt} N_{kt} \Psi_{n,t}^{Q,k}.$$

Indeed, this is exactly the extra term in equation (A12) induced by the net worth constraint if there is only one asset. Note, however, that this term is necessarily zero if the constraint takes the no-short-selling form of equation (A14): if the constraint binds, then $Q_{kt} N_{kt} = 0$, but if it does not bind, then complementary slackness guarantees that $\mu_{kt} = 0$. Thus short-selling constraints do not affect our formula for the welfare response to shocks.

D.2 Projecting “No-shock” Consumption, Wage and Asset Holdings

Our framework requires projections for each of the components in Proposition 2 forward through time from the moment of the identified impulse. In our baseline results, we assume that 2019 is a steady state: that is, absent the shock, prices, wages, and dividends streams would all have remained fixed at their 2019 levels, conditional on age. However, different products have experienced trend inflation, different assets have seen different trend returns and the college wage premium has changed over time. We therefore perform a robustness test where we assume that, absent the shock, all prices (of both assets and consumption goods), wages and dividend streams would follow their own log-linear trend.

Since agents age over the course of the shock impact, we have to account not only for the evolution of each component over time, but also over the lifecycle. Below, we describe the procedure to estimate each component.

For consumption, we first estimate a log-linear trend, over time, in the expenditure for each category and each combination of age-group, which we denote $\pi_{C,j}^{a,g}$. That is, we estimate the following regression for each good j and each age and education group using data from the CEX

$$\ln(p_{jt}c_{jt}^{a,g}) = \pi_{C,j}^{a,g} \cdot t + \varepsilon_{j,t}^{a,g}.$$

Then, taking the consumption in the last quarter of 2019 as $t = 0$, we follow the synthetic cohort approach to project the expenditures

$$(A15) \quad p_{jt}c_{jt}^{a,g} = p_{j0}c_{j0}^{(a+t)g} \cdot (1 + \pi_{C,j}^{a,g})^t$$

In other words, to compute consumption in t of a household that was 30-year old on impact, we take the consumption in $t = 0$ of a $(30 + t)$ -year old household and project that quantity over time using the log-linear growth rate of the corresponding category.

For wages and unemployment rates we follow a similar approach. Given a lifecycle profile of wages at $t = 0$, we compute wages at t as

$$W_t^{a,g} = W_0^{(a+t)g} (1 + \pi_W^{a,g})^t$$

where $\pi_W^{a,g}$ is trend wage inflation of age a households in education group g , estimated similarly to equation (A15).

Finally, we compute unemployment rates as

$$u_t^{a,g} = u_0^{(a+t)g}$$

where $t = 0$ is the last quarter of 2019. That is, we assume that there is no trend in unemployment rates over time, which is largely true over our sample period.

Finally, for assets we proceed as follows. The observed variables in the SCF are $Q_{k0}N_{k0}^{a,g}$. Following the approach mentioned above, we can estimate $Q_{k0}\Delta N_{k0}^{a,g}$. To obtain the $D_{k0}N_{k0}^{a,g}$ series for equity, note that we can rewrite

$$D_{k0}N_{k0} = \underbrace{\frac{D_{k0}}{Q_{k0}}}_{\text{Dividend yield}} \cdot \underbrace{Q_{k0}N_{k0}}_{\text{Observed value of holdings}}$$

Then, for a group g household which was a -years old on impact we compute dividends in time t as

$$D_{kt}N_{kt}^{a,g} = D_{k0}N_{k0}^{(a+t)g} \cdot (1 + \pi_k^D)^t$$

The dividend yield for bonds is simply the bond yield, while the dividend yield for equities is publicly available information. Again, we calculate the trend in dividends for each asset class

as π_k^D similarly to equation (A15). For changes in asset holdings we proceed similarly:

$$Q_{kt} \Delta N_{kt}^{ag} = Q_{k0} \Delta N_{k0}^{(a+t)g} \cdot (1 + \pi_k^Q)^t$$

for π_k^Q the estimated log-linear trend in the price index for asset k .