Are Inflationary Shocks Regressive? A Feasible Set Approach

Felipe Del Canto John Grigsby Eric Qian

Conor Walsh*

Latest draft can be found here.

Monday 30th December, 2024

Abstract

We develop a framework to measure the welfare impact of macroeconomic 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. 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 all cases, the dominant channel is the effect of the shock on the cost of accumulating assets, not movements in goods prices or labor income.

JEL Codes: E3, E2

*Del Canto: Harvard University, felipe_delcanto@g.harvard.edu;
Grigsby (*corresponding author): Princeton University and NBER, john.grigsby@princeton.edu;
Qian: Princeton University, ericqian@princeton.edu;
Walsh: Columbia Business School, conor.walsh.a@gmail.com.
Acknowledgments: Corina Boar, Greg Kaplan, Alisdair McKay, Simon Mongey, and Martin Schneider gave discussions which have greatly benefited the final version of the paper. We also thank Emilien Gouin-Bonenfant, Andre Kurmann, Nicola Pavoni, Richard Rogerson, Gianluca Violante, Christian Wolf and numerous seminar participants for insightful comments.

1. INTRODUCTION

The recent inflationary episode has renewed interest in understanding the distributional incidence of inflationary macroeconomic shocks. Whether an inflationary episode is regressive may appear to be a simple question, but confronting it requires overcoming two challenges. First, its distributional consequences may depend on the inflationary shock that drives it: 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 shocks. 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 at the top.

This paper studies the first-order welfare impact of inflationary shocks on heterogeneous households. We develop a new empirical framework that accounts for movements in all pieces of the budget constraint in response to macroeconomic shocks. 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 include additional constraints on the household, such as borrowing or net-worth constraints, and hence we 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 appropriately discounted 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, (5) transfer income from the government and (6) the direct effect of the shock on constraints. This holds without needing to specify the general equilibrium structure of the economy and is robust to allowing for general forms of idiosyncratic risk and borrowing constraints.

Our methodology requires two measurable inputs. First, we require empirical impulse response functions (IRFs) for the elements of an agent's feasible set, 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, the likes of which are readily available from household surveys.

We apply the framework to study two inflationary shocks that appear important in recent periods: oil supply shocks and monetary shocks. Using "internal instrument" Structural Vector Autoregression (SVAR) techniques, we estimate IRFs of disaggregated CPI price indices, labor-income series, and asset-price and dividend indices 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 different sources of inflation carry radically different distributional consequences. Oil supply contractions appear to be regressive, whereas nominal interest rate cuts are progressive. After a one standard deviation oil price increase, an average household with less than a high school education must be paid around \$800 (around 1.6% of one year's consumption) in 2019 to be able to afford their pre-shock level of utility. Meanwhile, college-educated households *gain* the equivalent of \$300 dollars (0.34% of a year's consumption) from the oil price increase, particularly in the middle of their life cycle. 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 high-education households lose around \$2,300 dollars (2.6% of a year's 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 have a 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, because they can now acquire equities more cheaply. This force causes oil price shocks to be highly regressive, even though dividend payouts modestly 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 then show that these results are both qualitatively and quantitatively robust to allowing for general idiosyncratic risk, borrowing constraints, and higher-order terms in the welfare effects.

Finally, we compare our results with those implied by a frontier business-cycle model: the two-asset heterogeneous-agent New Keynesian (HANK) model presented in Auclert, Bardóczy, Rognlie, and Straub (2021), following influential work in Kaplan, Moll, and Violante (2018). First, we validate our approach by showing that the feasible set approach, when applied to model-simulated data, replicates the true value changes of all households even when idiosyncratic risk is large, borrowing constraints are tight, and shocks are relatively large. We then show that this workhorse model generates different welfare effects of calibrated monetary and oil supply shocks than what is implied by the feasible set approach applied to the data, both on average and distributionally. We argue that the workhorse model misses the distributional welfare effects of these shocks primarily because it does not account for life-cycle dynamics in savings, housing, or meaningful household portfolio choice. Adding these ingredients to the workhorse model is therefore likely to be a fruitful direction for future work.

Our paper makes three contributions. The first is conceptual: the source of inflationary shocks matters for inflation's distributional consequences, and the effect of such shocks must be evaluated all throughout the budget constraint. The second is methodological: we demonstrate how to measure the distributional welfare consequences of generic macroeconomic shocks in settings with idiosyncratic risk and borrowing constraints. The third is empirical: expansionary monetary policy is progressive, whereas oil supply contractions are regressive.

Literature Review. Our feasible set approach highlights four dimensions of exposure to macroeconomic shocks: consumption expenditures, labor income, portfolios, and transfer income. Several literatures study each of these dimensions in isolation, which we detail below.

A growing literature has examined the distributional effects of trend inflation through the lens of household expenditures and consumption baskets. Hobijn and Lagakos (2005); Kaplan and Schulhofer-Wohl (2017); Jaravel (2019); Argente and Lee (2021) and Jaravel (2021) show substantial heterogeneity in average increases in the cost of living in the US. A consistent finding in this literature is that households at the bottom of the earnings distribution and older households have faced larger average inflation rates than have richer and younger households. A parallel literature studies how particular shocks affect the cost of living for different households. Cravino, Lan, and Levchenko (2020) show that monetary policy has a disproportionate effect on the cost-of-living of low-income households, partly because these households consume goods with less sticky prices. Clayton, Jaravel, and Schaab (2018) show that prices are more rigid in sectors selling to college-educated households, so that monetary shocks have a larger impact on the price of consumption for low-income households (see also Cravino and Levchenko (2017) Faber and Fally (2022), Kuhn, Kehrig, and Ziebarth (2021), Orchard (2022), Lauper and Mangiante (2023)). We find that differences in the consumption channel are small relative to the labor income or portfolio channels, which highlights the value of combining all movements of the budget constraint into one composite number.

There is also a recent literature focusing solely on what we term the labor-income channel. Bartscher, Kuhn, Schularick, and Wachtel (2021) find that accommodative monetary policy increases employment more for black households than for 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. Hubert and Savignac (2023) carry out a similar exercise using French administrative data and finds that the effects of ECB monetary policy shocks on labor income are U-shaped along the labor income distribution (see also Coglianese, Olsson, and Patterson (2022), Amberg, Jansson, Klein, and Rogantini Picco (2021) and Lee, Macaluso, and Schwartzman (2022)).

On the portfolio side, the seminal paper of Doepke and Schneider (2006) considers the redistribution of wealth 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, whereas young, middle-class households with fixed-rate mortgages are the main winners. 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.

Some recent papers study more than one of the channels suggested by our feasible set approach. Coibion, Gorodnichenko, Kueng, and Silvia (2017) show that contractionary monetary policy increases inequality in both consumption and earnings. They find that financial-income inequality responds by more than labor income inequality, which our results echo. Chang and Schorfheide (2024) use functional vector autoregression (fVAR) techniques to document that expansionary monetary policy shocks reduce earnings inequality, weakly increase consumption inequality, and have no effect on financial-income inequality, using data from the Consumer Expenditure Survey (CEX) and Current Population Survey (CPS). Andersen, Huber, Johannesen, Straub, and Vestergaard (2022) use Danish administrative data to show that income, consumption, and wealth increases from softer monetary policy are monotonically increasing in pre-shock income levels, and find an important role for non-labor income. Ferreira, Leiva, Nuño, Ortiz, Rodrigo, and Vazquez (2024) study the distributional consequences of the 2021 inflation in Spain and find that the consumption heterogeneity channel is an order of magnitude smaller than what we term the labor income and portfolio channels in this episode for household balance sheets (see also Pallotti, Paz-Pardo, Slacalek, Tristani, and Violante (2023)).

Our paper makes three contributions to this large reduced-form literature. Empirically, following a long literature in public finance,[1] we study money-metric welfare movements, rather than wealth, consumption, or income inequality. This approach allows us to compare the relative magnitudes of each channel for welfare, and reveals that the portfolio channel appears to be the largest. Methodologically, we show how one can measure welfare effects of macroeconomic shocks accounting for idiosyncratic risk, borrowing constraints, and non-homothetic preferences. Conceptually, we estimate the response to both oil and monetary shocks, and show that these two sources of inflation have radically different distributional consequences.

The large reduced-form literature is complemented by a set of papers that 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). 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.[3] Glover, Heathcote, Krueger, and Ríos-Rull (2020), Gagliardone and Gertler (2023) and Rubbo (2023) study recent shock episodes through the lens of structural macro models. Pugsley and Rubinton (2023) study the distributional welfare effects of the Volcker disinflation in a calibrated HANK model. 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 takes a different tack, not needing to specify the general equilibrium structure to assess welfare changes. It also compares the estimated effects with those implied by a benchmark two-asset HANK model, and finds that the benchmark model does not generate asset accumulation patterns that match the data, and hence a portfolio channel that cannot speak to the empirical evidence.

[1] See e.g. Samuelson (1974), Deaton (1989), Deaton and Zaidi (2002), and Saez and Stantcheva (2016).

[2] Tzamourani (2021) measures the unhedged interest rate exposure channel from Auclert (2019) in the eurozone.

[3] See also Rubbo (2024) and Schaab and Tan (2023), who study the aggregate and distributional implications of monetary policy in multi-sector heterogeneous-agent models.

Our paper is also related to recent papers seeking to measure the welfare effects of price movements. 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. Dávila and Schaab (2022) considers how to make aggregate 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 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 (see also Baqaee and Burstein (2022), Baqaee, Burstein, and Koike-Mori (2022) and Jaravel and Lashkari (2022)) Many of these papers consider flexible ways to measure non-homothetic price indices, but abstract from dynamic decision-making. One exception is Baqaee, Burstein, and Koike-Mori (2024), who account for dynamic decision-making in price index theory. In contrast to their flexible "top-down" approach of measuring welfare changes without specifying the channel driving the change, we employ a Taylor approximation and envelope arguments to build up to welfare from the "bottom-up"; that is, by seeking to measure and sum the different channels through which household welfare can move.

Methodologically, our paper is most closely related to 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, using a money-metric measure of welfare from a first-order approximation to the value function. Like us, they 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.

Our work differs from this important paper in a few central ways. First, we focus our methodology on the effects of identified, exogenous shocks, showing how to aggregate impulse response functions estimated using standard time-series techniques. Fagereng et al. (2022) instead attempt to evaluate the welfare effects of large *observed* changes in asset prices, which is more difficult to conceive of without a theory of where these changes arose from (in particular, the role of preference changes). Second, we focus our evaluation on *all* welfare effects of inflationary shocks, not just those stemming from asset-price changes. Our methodology allows us to size and aggregate all of these pieces in a consistent manner. Third, we develop and estimate our results in the presence of general idiosyncratic risk. Finally, we validate our methodology within a frontier two-asset HANK model, and show that the model misses inflationary shocks' redistributive effects, due to the lack of a realistic age structure and portfolio dynamics.

2. FRAMEWORK

This section presents our framework to analyze the distributional consequences of inflationary shocks. 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 aggregate empirical IRFs using cross-sectional data to estimate the first-order welfare impact of shocks on these different agents.

Setting. Time is discrete and indexed by t. There is a continuum of households indexed by i. There is both aggregate and idiosyncratic uncertainty; let s_t denote a history of realizations of states of the world, including idiosyncratic states, up to period t. We will call this history a state for convenience, but it should be read to include a sequence for previous realizations of stochastic variables before t. There are J consumption goods, indexed by $j \in \{1, \ldots, J\}$, with good j having price $p_{jt}(s_t)$ in period t.

There are $K + 1$ long-lived assets, indexed by $k \in \{0, 1, \ldots, K\}$, available for trading in each period. Asset k pays a nominal dividend $D_{kt}(s_t)$ and may be traded at a price $Q_{kt}(s_t)$ in state s_t. We assume that asset $k = 0$ is a one-period nominal bond which pays one unit in all states s_t. We define the cumulative return on buying a sequence of these bonds from period 0 to t as $R_{0\to t}(s_t) \equiv \prod_0^t Q_\tau(s_\tau)^{-1}$. Lastly, 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. 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 $N_{kt}^{ag}(s_t)$ denote the amount of asset k held by group g of initial age a at time t given a realization of s_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 $\chi_k^{ag}(\Delta N_{kt})$. The one-period bond is assumed to not be subject to adjustment costs.

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

Households have time-separable preferences with subjective discount factor $\beta_t^{ag} \in (0, 1)$. This is read as a household of initial age a, in group g, discounting period t from period zero, where the rate at which they discount is potentially non-constant. The household has preferences over consumption, labor, and asset holdings. We assume that each household type derives utility from consumption via an aggregator of goods

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

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 $\mathcal{C}^{ag}(\cdot)$ is increasing and continuously differentiable in all its arguments.

Households of type ag's preferences may be summarized by the differentiable utility function

$U^{ag}(C_t^{ag}, \{N_{kt}^{ag}\}_{k=1}^K, L_t^{ag})$, where L_t^{ag} is the labor supplied by households of initial age a at time t. We assume that $U^{ag}(\cdot)$ 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.[4] Excluding the quantity of one-period bonds from being in the utility function directly allows us to conveniently characterize an Euler equation for the households in terms of expected marginal utilities of consumption and the return on the bond.

Labor income for individual i is given by the product of three terms: $W_t^{ag}(s_t)e_t^i(s_t)L_t^{ag}(s_t)$. W_t^{ag} is an aggregate component of wages that varies with the aggregate state of the economy, as well as the age and group of the individual. $e_t^i(s_t)$ is an idiosyncratic stochastic process for efficiency units of labor, which has support on $\mathbb{R}_+$ and satisfies $\mathbb{E}_0[e_t^i] = 1$. This captures both general transitory and permanent fluctuations to idiosyncratic labor income. The stochastic process for e_t^i may be different for different household groups g and initial ages a. The state s_t contains realizations of both aggregate and idiosyncratic processes. We assume aggregate and idiosyncratic risks are independent so that we can partition the state into an aggregate component, s_t^A, and an individual component s_t^i: $s_t = \{s_t^A, s_t^i\}$ and $\pi_t(s_t) = \pi_t(s_t^A)\pi_t(s_t^i)$.

A representative type g household of initial age a takes as given its initial stock of asset holdings $\{N_{k,-1}\}_k$ and is a price taker. It solves the following utility-maximization problem:

(2)
$$V^{ag}(\{N_{k,-1}\}_k) = \max_{\{\{c_{jt}^{ag}(s_t)\}_j, L_t^{ag}(s_t), \{N_{kt}^{ag}(s_t)\}_k\}_{t=0,s}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta_t^{ag}\delta_t^{ag} U^{ag}(C_t^{ag}(s_t), \{N_{kt}^{ag}(s_t)\}_{k=1}^K, L_t^{ag}(s_t)),$$

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

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

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

$$\lim_{T\to\infty} \mathbb{E}_0[R_{0\to T}^{-1} N_{kT}^{ag} Q_{kT}] \geq 0, \qquad \forall k \in \{0, \dots, K\}. \tag{3}$$

Stochastic Structure. We suppose that the (general equilibrium) law of motion for the prices of the economy admits a VAR representation. Concretely, we assume that dividends, asset prices, goods prices, wages, and transfers are stochastic, and take the form

$$\begin{aligned} D_{kt} = \bar{D}_{kt}\exp(v_{kt}^D)^\sigma, \quad Q_{kt} = \bar{Q}_{kt}\exp\left(v_{kt}^Q\right)^\sigma, \quad p_{jt} = \bar{p}_{jt}\exp\left(v_{jt}^p\right)^\sigma, \\ W_t^{ag} = \bar{W}_t^{ag}\exp\left(v_t^{W^{ag}}\right)^\sigma, \quad T_t^{ag} = \bar{T}_t^{ag}\exp\left(v_t^{T^{ag}}\right)^\sigma, \end{aligned} \tag{4}$$

[4] Allowing assets to directly impact utility is a common tool in monetary and financial economics to capture the liquidity values of assets like cash (Sidrauski, 1967; Van den Heuvel, 2008). In addition, households may draw utility directly from service flows from assets such as housing, and durable goods such as cars. Lastly, bequests may be thought of separately as drawing utility from the *value* of assets one holds. We consider this specification separately in the Online Appendix and explore the sensitivity of our results to this specification in Section 8.

where $\sigma > 0$ is a parameter that scales the variance of the aggregate 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^D_{kt}). We assume that the shock processes are functions of current and lagged values of a structural shock vector $\boldsymbol{\epsilon}_t$, such that

$$v^D_{kt} = \theta^D_k(L)\boldsymbol{\epsilon}_t, \quad v^Q_{kt} = \theta^Q_k(L)\boldsymbol{\epsilon}_t, \quad v^p_{jt} = \theta^p_j(L)\boldsymbol{\epsilon}_t,$$
$$v^{W^{ag}}_t = \theta^{W^{ag}}(L)\boldsymbol{\epsilon}_t, \quad v^{T^{ag}}_t = \theta^{T^{ag}}(L)\boldsymbol{\epsilon}_t,$$

where each $\theta(L)$ is a lag operator matrix of finite dimension, and the elements of $\boldsymbol{\epsilon}_t$ are mutually uncorrelated. We collect these v^x_t into a vector $\mathbf{v}_t$. We further assume that the structural shocks $\mathbf{v}_t$ have no direct effect on household utility functions; we therefore rule out preference shocks, such as discount rate shocks. Finally, we assume that these structural aggregate shocks are independent from the idiosyncratic income process for each individual.

We leave the production structure of the economy unspecified, as long as equilibrium can be written in this fashion. Importantly, the aggregate economy need not be efficient.

Following Stock and Watson (2018), we define the vector of structural impulse responses of the collection of variables affecting households' budget constraint $\mathbf{v}_t \equiv (\{v^p_{jt}\}_j, \{v^Q_{kt}, v^D_{kt}\}_k, \{v^{W^{ag}}_t, v^{T^{ag}}_t\}_{ag})$ at time t to the n^{th} entry of the structural shock vector $\boldsymbol{\epsilon}$ at time $t = 0$ as

$$\boldsymbol{\Psi}_{n,t} \equiv \mathbb{E}_0[\mathbf{v}_t|\epsilon^n_0 = 1] - \mathbb{E}_0[\mathbf{v}_t|\epsilon^n_0 = 0].$$

Elements of this vector are denoted with superscripts, such as $\Psi^{p,j}_{n,t}$ for the consumption price of good j, and $\Psi^{Q,k}_{n,t}$ for the asset price of asset k.

2.1 Welfare Response to Shocks

Our notion of a change in welfare follows the idea of an impulse response at time 0. Specifically, for an innovation to a fundamental shock ϵ^n_0 at time 0, we define

$$d\tilde{V}^{ag} \equiv V^{ag}(\{N_{k,-1}\}_k; \epsilon^n_0 = 1, \epsilon^{-n}_0 = 0) - V^{ag}(\{N_{k,-1}\}_k; \epsilon^n_0 = 0, \epsilon^{-n}_0 = 0).$$

This is the change in welfare for an agent of age a at time 0 from a unit-sized shock to element n of the structural shock vector, holding all other fundamental shocks constant at time 0. To facilitate comparisons across agents in different groups, we follow a long literature in public finance and welfare economics, and normalize this change in welfare by the marginal utility of a dollar in the absence of a shock. This money-metric utility calculation measures how much different households would be willing to pay in dollars at time 0 to avoid the fundamental shock at time 0. We call this change $dV^{ag} \equiv d\tilde{V}^{ag}/U^{ag}_c(\cdot)$.

We first consider a simple baseline case that will guide our empirical analysis. In our baseline, we characterize money-metric welfare changes in response to shocks as the variance of both aggregate and idiosyncratic risk tends to zero (so that $\sigma \to 0$ and $Var(e^i_t) \to 0$).

Proposition 1. *As both aggregate and idiosyncratic risk become small, the change in money-metric welfare from an impulse to an element n of the fundamental shock vector at $t = 0$ is to a first order*

$$dV^{ag} = \sum_t R_{0\to t}^{-1} \Bigg(\underbrace{-\sum_j p_{j,t} c_{jt}^{ag} \Psi_{n,t}^{p,j}}_{\text{Consumption Price Changes}} + \underbrace{W_t^{ag} L_t^{ag} \Psi_{n,t+h}^{W^{ag}}}_{\text{Labor Income Changes}} + \\ + \sum_k \Big[\underbrace{N_{kt-1}^{ag} D_{kt} \Psi_{n,t}^{D,k}}_{\text{Asset Income Changes}} - \underbrace{Q_{kt} \Delta N_{kt}^{ag} \Psi_{n,t}^{Q,k}}_{\text{Asset Price Changes}} \Big] + \underbrace{T_t^{ag} \Psi_{n,t}^{T^{ag}}}_{\text{Transfer Income Changes}} \Bigg), \tag{5}$$

where choices are evaluated at their deterministic values.

All proofs are contained in Appendix A. Proposition 1 states that the first-order money-metric welfare gain in response to a fundamental shock 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 weights the percentage change in the price of each good induced by the shock by total spending on each good, ignoring 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 shock.

Second, the shock may induce changes in labor income for the household if their wage moves. We term this the "labor-income channel" of the shock.

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, echoing 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*.[5] A rise in the price of the S&P500 at time t is mainly relevant for those at a point in their life cycle in t 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). This logic is also clear in Dávila and Korinek (2018), who show that the product of an agent's net trading position and the induced equilibrium asset price change is crucial for the distributional effects of shocks. We refer to the effect of the shock on asset prices and dividends as its "portfolio channel".[6]

Finally, the shock may shift the present value of taxes owed or transfers paid to the household. We term this the "transfer channel".

Intuitively, a shock to the expected 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, it suffices to simply consider the present discounted value of movements in households' budget constraints. Note this does not mean that we *assume* that households do

[5]As we will soon see, this may not be true if the household is subject to borrowing constraints.

[6]It is worth clarifiying that the appearance of the asset terms in this equation is not caused by assets entering the utility function; whether or not a particular asset quantity has a direct effect on utility has no effect on this formula. Instead, they arise from changes in asset prices and dividends moving the agent's budget set. If asset *values* enter utility, there is an additional term measuring how utility changes directly. We explore this in Section 8.

not substitute, just that the welfare effects of doing so are negligible for small shocks. In addition, if households are on their Euler equation, risk is small, and bonds do not enter utility, they discount future movements in prices by the risk free rate. Lastly, we note that the welfare difference expressed here is a *discrete difference*, expressed as the dollar gain in a world where element n of the fundamental shock is set to 1 at time 0, relative to a world where it is zero (as in the language of time series impulse response functions).[7] It estimates the true welfare response with a "certainty equivalence" welfare measure where aggregate risk is small.

Proposition 1 forms the foundation of our empirical strategy. It provides a method to appropriately aggregate estimated IRFs of macroeconomic inflationary shocks into a welfare metric. It is non-parametric in the sense that it holds without specifying the general equilibrium structure of the economy or the nature of the utility function.[8] In particular, 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. It is, however, rather stark, in that as a measurement tool, it will only perform well against the true dV^{ag} when risk is small, and households face no borrowing constraints. A large literature in macroeconomics suggests that these are important for both describing behavior and evaluating welfare. We now turn to incorporating the effect of these two considerations.

2.2 Idiosyncratic Risk

In the absence of complete markets, marginal utilities of consumption are generally not equated across states. In this case, aggregate price changes that shift consumption in all idiosyncratic states must be weighted by the expected marginal utility of such a price change. To lighten notation, we drop dependence on household type (a, g), but everything henceforth should be understood as holding within a household group g and initial age a.

To state the next result, let $E(C_t, \{p_{jt}\})$ denote the household's static expenditure function given a desired consumption level C_t and price vector $\{p_{jt}\}$. Further, define the marginal cost of consumption C_t as

$$P_t(C_t, \{p_{jt}\}) \equiv \frac{\partial(E(C_t, \{p_{jt}\})}{\partial C_t}.$$

This can be thought of as an ideal price index when preferences are homothetic, but may in general depend on C_t and hence total income. As a slight abuse of notation, let $U_C(s_t)$ denote the marginal utility of the consumption under the household's optimal choices after realization of history s_t. Now define the shifter

$$\Theta_t^W \equiv Cov\left(\frac{U_C(s_t)/P_t(s_t)}{\mathbb{E}_0[U_C/P_t]}, \frac{W_t(s_t)e(s_t)L_t(s_t)}{\mathbb{E}_0[W_te_tL_t]}\right).$$

This captures how deviations of marginal utility from the mean move with (de-meaned) labor

[7] This is discrete despite the fact that we invoke an envelope theorem for the proposition. The envelope theorem is applied to changes in σ: the scale of aggregate risk in the economy.

[8] Indeed, there is no requirement that the economy be efficient or without distortions.

income. When markets are complete and aggregate risk is absent, households perfectly smooth consumption, and so this covariance is zero. With uninsurable labor-income risk and risk-averse preferences, these Θ_t^W will generally be negative.[9] Similarly, define

$$\Theta_t^{p,j} \equiv Cov\left(\frac{U_C(s_t)/P_t(s_t)}{\mathbb{E}_0[U_C/P_t]}, \frac{p_{j,t}(s_t)c_{jt}(s_t)}{\mathbb{E}_0[p_{j,t}c_{jt}]}\right) \qquad \Theta_t^T \equiv Cov\left(\frac{U_C(s_t)/P_t(s_t)}{\mathbb{E}_0[U_C/P_t]}, \frac{T_t(s_t)}{\mathbb{E}_0[T_t]}\right)$$
$$\Theta_t^{D,k} \equiv Cov\left(\frac{U_C(s_t)/P_t(s_t)}{\mathbb{E}_0[U_C/P_t]}, \frac{D_{kt}(s_t)N_{kt-1}(s_{t-1})}{\mathbb{E}_0[D_{kt}N_{kt-1}]}\right) \qquad \Theta_t^{Q,k} \equiv Cov\left(\frac{U_C(s_t)/P_t(s_t)}{\mathbb{E}_0[U_C/P_t]}, \frac{Q_{kt}(s_t)\Delta N_{kt}(s_t)}{\mathbb{E}_0[Q_{kt}\Delta N_{kt}]}\right)$$

as the covariance of (detrended) marginal utility of consumption with (detrended) good expenditures, asset holdings, asset expenditures, and transfers. We can then develop a similar result to Proposition 1, but incorporate general idiosyncratic risk.

Proposition 2. *As only aggregate risk becomes small ($\sigma \to 0$), the first-order change in money-metric welfare from an impulse to an element n of the fundamental shock vector at $t = 0$ is*

$$dV = \sum_t R_{0\to t}^{-1}\Bigg[\Bigg(\underbrace{-\sum_j \mathbb{E}_0[p_{j,t}c_{jt}]\Psi_{n,t}^{p,j}(1+\Theta_t^{p,j})}_{\text{Consumption Price Changes}} + \underbrace{\mathbb{E}_0[W_t e_t^i L_t]\Psi_{n,t}^W(1+\Theta_t^W)}_{\text{Labor Income Changes}}$$
$$+\sum_k\left[\underbrace{\mathbb{E}_0[N_{kt-1}D_{kt}]\Psi_{n,t}^{D,k}(1+\Theta_t^{D,k})}_{\text{Asset Income Changes}} - \underbrace{\mathbb{E}_0[Q_{kt}\Delta N_{kt}]\Psi_{n,t}^{Q,k}(1+\Theta_t^{Q,k})}_{\text{Asset Price Changes}}\right] + \underbrace{\mathbb{E}_0[T_t]\Psi_{n,t}^T(1+\Theta_t^T)}_{\text{Transfer Income Changes}}\Bigg)\Bigg], \tag{6}$$

where choices and risk-adjustment shifters are evaluated at $\sigma = 0$.

The formula in Proposition 2 is very similar to that found in Proposition 1. However, price movements are now weighted by the covariance of marginal utilites of consumption with each element of the budget set (labor income, goods expenditure etc.). Note that when idiosyncratic risk vanishes, households are risk-neutral or there are complete markets, all Θ_t^x are equal to zero and the welfare change collapses to that of Proposition 1.[10]

Intuitively, agents discount more heavily riskier income streams that co-vary negatively with the marginal utility of consumption. If the agent cannot smooth labor-income fluctuations across states, they discount expected wage movements in the future at a higher rate than the risk-free rate. The same is true of expected future consumption price changes or changes in asset prices and dividends.

In principle, the covariances in Proposition 2 are estimable, though doing so requires parametric assumptions on marginal utilities of consumption that are not required in Proposition 1. We discuss our strategy for doing so below.

[9] Appendix G provides intuition for when these Θ are large or small.
[10] One can follow similar steps to show that a similar formula holds in response to unanticipated perturbations to the deterministic components of the stochastic processes, even when aggregate risk does not vanish, only that the structural impulse responses $\Psi_{n,t}^x$ are replaced with the perturbations in the deterministic components.

2.3 Borrowing Constraints

The previous results leverage an Euler equation in the riskless bond, whose interest rate can then be used to discount future price changes (with appropriate adjustments for idiosyncratic risk). However, a large literature has emphasized that borrowing constraints may be important for consumption and savings decisions (Kaplan and Violante, 2024). Poorer households in particular may not be on their Euler equations and wish to borrow to smooth consumption fluctuations, but are unable to do so. Suppose that the household faces a constraint that net worth cannot fall below a certain level in any state, so that

$$\sum_k Q_{kt}(s_t)N_{kt}(s_t) \geq \underline{b}, \tag{7}$$

for some $\underline{b} \leq 0$. We can modify Proposition 1 as follows.

Proposition 3. *When households face the additional constraint in (7), as both aggregate and idiosyncratic risk become small, the first-order change in money-metric welfare from an impulse to an element n of the fundamental shock vector at $t = 0$ is*

$$\begin{aligned} dV = \sum_t R^{-1}_{0\to t} \overbrace{\prod_{s=0}^{t-1}(1+\tau_s)^{-1}}^{\text{Discount Wedge}} \Big(& -\sum_j p_{j,t}c_{jt}\Psi^{p,j}_{n,t} + W_t L_t \Psi^W_{n,t} + T_t \Psi^T_{n,t} \\ & + \sum_k \left[N_{kt-1}D_{kt}\Psi^{D,k}_{n,t} - Q_{kt}\Delta N_{kt}\Psi^{Q,k}_{n,t} \right] + \underbrace{\frac{\tau_t}{1+\tau_t}\sum_k Q_{kt}N_{kt}\Psi^Q_{kt}}_{\text{Value of Relaxed Constraint}} \Big), \end{aligned} \tag{8}$$

where choices are evaluated at their deterministic values and τ_t solves

$$\frac{P_{t+1}U_C(C_t, \{N_{kt}\}, L_t)}{P_t U_C(C_{t+1}, \{N_{kt+1}\}, L_{t+1})} = \left(\frac{\beta_{t+1}}{\beta_t}\right)\left(\frac{\delta_{t+1}}{\delta_t}\right) R_t(1+\tau_t). \tag{9}$$

Equation (8) differs from Proposition 1 in two ways. First, there is a "wedge" in the Euler equation which affects the rate at which future price movements are discounted, here denoted τ_t. If the constraint does not bind, $\tau_t = 0$ and there is no wedge in the Euler equation. The more binding the constraint, the larger is τ_t, and the greater the difference in discount rates.

The second additional term reflects the extra value from rising asset prices after the impulse, which can help in relaxing the borrowing constraint of the consumer. Again, this term is larger the higher is τ_t or, equivalently, the more binding the constraint.[11]

As with idiosyncratic risk and Proposition 2, obtaining values for τ_t from the data requires imposing parametric structure on the utility function of each consumer. We outline our strategy for doing so below.

[11] We show in Appendix A that a similar expression holds for general constraints of the form $G(\mathbf{x}, \mathbf{z}) \leq 0$, where $\mathbf{x}$ is the set of choice variables for the household and $\mathbf{z}$ is the set of objects the household takes as exogenous (such as prices, parameters etc.).

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

$$Q_{kt} N_{kt} \geq 0, \tag{10}$$

where we let μ_{kt} be the Lagrange multiplier on the short-selling constraint for asset k.[12] Because, by assumption, this constraint does not affect the risk-free bond ($k = 0$), such constraints do not distort the Euler equation and so do not affect effective discount rates. One can show that constraints on short-selling leads to one extra term in the expression for the welfare response to shocks, given by $\mu_{kt} Q_{kt} N_{kt} \Psi^{Q,k}_{n,t}$. Note, however, that this term is necessarily zero if the constraint takes the no-short-selling form of equation (10): 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.

2.4 Discussion

The three propositions above guide our empirical analysis of inflationary shocks. They permit the econometrician to aggregate IRFs of very different objects – the price of food and the S&P500, for instance – into a welfare-relevant common unit. Our strategy is to estimate these empirical IRFs for identified shocks, and then use cross-sectional data to aggregate them to first-order welfare effects. We begin by studying the welfare effects of shocks using Proposition 1, which does not rely on any parametric forms for utility. We then consider idiosyncratic risk and borrowing constraints, which come at the cost of requiring some parametric structure on utility.[13]

One major benefit of this approach is that the data required to estimate the formulas in Propositions 1-3 are readily available from aggregate time series and publicly available household surveys. First, we require time-series data on prices to estimate the IRFs $\mathbf{\Psi}_{n,t}$, specifically data on consumption prices, wages, and asset prices, all of which are easily accessible from national statistical agencies. The second piece of data we require is life-cycle data on consumption expenditures by good ($p_{jt} c_{kt}$), labor-income profiles ($W_t e^i_t L_t$), and asset accumulation profiles (ΔN_{kt}), which we draw from various large-scale household surveys. In all three cases, the formula calls for the life-cycle variables to be evaluated at the deterministic steady state. We approximate this by assuming that 2019 represents a stochastic steady state (i.e., is absent aggregate shocks, so that $\mathbf{v}_t = 0$), and further assuming that the stochastic steady state is "close" in the sense of household choice variables to the deterministic steady state. We explore the robustness of our results to using different years of household survey data in the Appendix and find this makes little difference to our conclusions. However, this assumption may be less tenable in other contexts, and for shocks other than the ones we study.

[12] By dividing both sides by Q_{kt}, this formulation also accounts for constraints of the form $N_{kt} \geq 0$.

[13] We additionally account for durable consumption through a utilization approach. As we show in Appendix A.3, this leads fluctuations in durable goods' prices to appear both in the consumption channel – where the price of consumption is tied to depreciation and the price of new durables – and the portfolio channel.

Another 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 much more heterogeneity than is usually tractable in a structural model. Our empirical application below includes dozens of consumer and asset prices, and a rich age and group structure, which would likely be infeasible in a structural setting. In addition, 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 they face. 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.[14]

An alternative way to study the welfare effects of a shock in reduced form would be to estimate the shock's effect on the variables over which households have preferences, such as consumption and labor supply. Our feasible set approach carries three benefits over such an exercise. First, price data are readily available, while consumption data are often measured with error. Second, measuring "consumption" is difficult when people have non-homothetic preferences: without group-specific price indices, comparing changes in nominal spending across groups may not be informative about welfare changes (Oberfield, 2023). In contrast, prices are relatively easy to measure. Third, the portfolio channel could lead to long lags between a shock and its effect on consumption, which could be difficult to convincingly measure.[15] Finally, the feasible set approach does not rely on any assumptions regarding how households trade off, say, consumption and leisure, which would be necessary to aggregate responses of consumption and leisure into one composite "welfare change".

Nevertheless, the feasible set approach is limited in a few key ways. First and most importantly, the feasible set approach approximates the first-order effect of shocks around an economy with no aggregate risk.[16] When aggregate risk or higher-order effects are large, the quality of the linear approximation may suffer.[17] We study the influence of second-order effects in Section 8, and find them to be small in our application, but they may be large elsewhere.

Second, the feasible set approach does not give guidance on welfare aggregation of the sort studied by, for instance, Dávila and Schaab (2022). Rather, the feasible set approach gives a method to characterize an individual household's willingness to pay for a given shock. Further, we do not impose that assets are in fixed net supply, so accumulation by one group of households need not correspond 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

[14] McKay and Wolf (2023) make a similar assumption to show that policy shocks inform policy counterfactuals. In addition, we assume that households are price takers in all markets, and that they face linear price schedules.

[15] For example, changes in asset prices today can change the amount of wealth an agent accumulates today, which may have impacts on consumption for decades hence.

[16] In the zero-risk world, the reason for holding different asset classes is to trade off differential return timings, precautionary savings against idiosyncratic risk, the utility flow from assets and the costs of changing the holdings.

[17] Certainty equivalent in aggregates is commonly employed in many business cycle macro models and techniques (e.g. Boppart, Krusell, and Mitman (2018)).

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 US household sector.

Third, although extending the basic formula of Proposition 1 to handle general idiosyncratic risk and borrowing constraints greatly extends the power of the approach, we want to be careful not to claim that our approach is a panacea for heterogeneous-agent welfare calculations. A key assumption necessary to derive Propositions 2 and 3 is that aggregate and idiosyncratic risks are independent. This seems like a reasonable assumption for many risks, such as health shocks, but for unemployment, where job-losing and job-finding rates vary with the business cycle, it may be less reasonable.

Finally, the feasible set approach is not well suited to studying taste and uncertainty shocks.[18] This is partly why we apply the approach to study the response to identified monetary and oil shocks, which are likely orthogonal to preference shocks. In addition, it relies on the assumption that households are price-takers. This is probably reasonable when it comes to consumption and asset prices, but may be less so in labor markets. In a wider sense, our setting will not apply to any form of non-linear pricing. In addition, parametric assumptions on utility are required to handle idiosyncratic risk and borrowing constraints. It is also limited by the extent to which high-quality time-series data are available to credibly estimate IRFs. It is worth keeping these caveats in mind.

3. DATA

This section describes our data on household consumption, income, and portfolios, as well as time series of prices, dividends, and wages. Appendix B provides further information.

Throughout, we focus on household groups g defined by the educational attainment of the household head. We distinguish households by education for three reasons. First, education is a readily available statistic in many datasets. Second, education may often be a better proxy for a household's permanent income than their income in any given year. Third, education is a fixed characteristic of the household that maps cleanly to our organizing framework and, as we describe below, allows us to project forward household decisions using a synthetic cohort approach. As such, our distributional effects will consider average welfare changes between educational groups. In Appendix F, we repeat our analysis by income quintile in 2019.

We compute life-cycle profiles of consumption, wages, asset holdings, and transfer income within each education group. We consider households whose head is at least 25 years old, and combine all households over 75 into one group. Our baseline approach measures the cross-sectional consumption, portfolio, and income variables using 2019 data and assumes 2019 represents a steady state. Thus, a represents the household head's age as of 2019. A

[18] Baqaee and Burstein (2023) show taste shocks only affect welfare starting at second order.

period t denotes a quarter, in keeping with our available consumption data.

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 separately tracking the price of "Medical Equipment and Supplies" in 2006. Since we need long time series to estimate our regressions, we only track categories that satisfy three criteria: they must (1) be available at least back to 1998, (2) represent a sufficiently large share of the aggregate consumption bundle, and (3) add up to 100% of consumption. This leaves us with 25 consumption goods, roughly at the level of the BLS' Consumer Price Index (CPI) categories.[19]

We use data from the interview component of the 2019 CEX to measure group-specific life-cycle consumption of goods. The CEX is a nationally representative quarterly survey run by the BLS that provides data on expenditures of US consumers at the household level. Its broad coverage of all components of household expenditure makes it uniquely well-suited to our exercise.[20]

We group the expenditure categories into 25 groups for which we have a CPI price series.[21] We then calculate the average expenditure of households of age a in group g on each good j. Next, we use locally-weighted scatterplot smoothing (LOWESS) over the life cycle within each group g to minimize large swings in consumption caused by measurement error. These (smoothed) average expenditures form our estimate of $p_{j0}c_{j0}^{ag}$. In our baseline scenario, we assume a constant life-cycle profile of consumption of each good, so that $p_{jt}c_{jt}^{a,g} = p_{j0}c_{j0}^{a+t,g}$ absent any shocks. For example, 25-year-old households at $t = 0$ will have the same baseline expenditure on each good j in period $t = 4$ as did 26-year-old households in period $t = 0$.

Labor Income Data. For labor income, we use monthly earnings information from the CPS. We use the "Outgoing Rotations Group" (ORG) component to construct quarterly age $\times$ education group average weekly earnings profiles in the 2019 CPS. As above, we smooth these averages over the life cycle within each group g using a LOWESS procedure. We include those with zero income in this exercise, and use this as our measure of $\mathbb{E}[W_t e_t^i L_{it}]$. As with consumption data, our baseline scenario assumes a constant life-cycle profile of earnings absent any shocks, so that $\mathbb{E}[W_t^{a,g} L_t^{a,g}] = \mathbb{E}[W_0^{a+t,g} L_0^{a+t,g}]$ for all t, g, and a. We multiply weekly earnings by 13 to get quarterly earnings.[22]

[19] Prior 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 are only recently available. However, there is no *a priori* reason for inflation rates of finer product categories to be differentially *responsive* to short-run shocks.

[20] Whereas 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 representative US dataset that accounts for all of household's expenditure.

[21] Note that households do not, in general, report healthcare expenditures covered by Medicare or Medicaid.

[22] We do not calculate time series of earnings separately by age and education because we could not find significant differences in the response of earnings by age conditional on education, but including multiple age groups in our estimation substantially increased the noise of our estimates.

Portfolio Data. We use the Survey of Consumer Finances (SCF) to measure household balance sheets by age and education level. The SCF is a triennial nationally representative survey in which respondents are asked 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. 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 all ages. Next, we perform a LOWESS smoothing over the life cycle within groups. Finally, we approximate $Q_{k0}\Delta N_{k0}^{a,g}$ with the estimated change in asset holdings between adjacent ages implied by the LOWESS: $Q_{k0}N_{k0}^{a,g} - Q_{k0}N_{k0}^{a-1,g}$.[23] We again assume a constant life cycle so that $Q_{kt}\Delta N_{kt}^{a,g} = Q_{k0}\Delta N_{k0}^{a+t,g}$.[24]

The SCF data directly give us the value of asset holdings $Q_{k0}N_{k0}^{ag}$. To recover the no-shock dividend income of each asset class, we use data on dividend yields in 2019, which report D_{k0}/Q_{k0}. Multiplying the value of the asset holding from the SCF by the dividend yield returns $D_{k0}N_{k0}^{ag}$ 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.[25] Thus, the dividend component is only relevant for equities and mortgages. We proxy the dividend yield for equities using the dividend yield of the S&P500. Data on effective mortgage interest rates come from the National Income and Product Accounts.

We use a variety of asset-price indices for our analysis. Equity price returns, estimated dividend yields, and dividend growth are computed from the value-weighted indices 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. Effective mortgage rates are obtained from the Bureau of Economic Analysis (BEA).

Transfer Income. We measure transfer income, which we define as the sum of means-tested transfer income and social insurance payments, using the Survey of Income and Program Participation (SIPP).[26] It is difficult to estimate IRFs for transfer income by group, because the

[23] This approach has the benefit of filtering out movements in the value of assets that arise from short-run 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.

[24] 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 $t-1$. Relaxing this strong assumption would require panel data on asset holdings and trades for various household groups, which is seldom available in the US.

[25] The assumption that nominal bond coupon payments are fixed and do not respond to shocks is a reasonable assumption for the US: adjustable-rate mortgages 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 US treasury bonds and corporate bonds usually pay a fixed coupon.

[26] 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

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, since Social Security payments, which form the bulk of transfer income, are explicitly indexed to the CPI. Because 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.[27]

4. Estimating Impulse Response Functions (IRFs)

This section describes our approach to estimating the requisite impulse responses to two inflationary macroeconomic shocks: a shock to world oil prices, and US monetary expansions.

Our first application considers responses to an oil-supply-news shock. We use the Känzig (2021) oil-surprise series as our source of identifying variation, which uses high-frequency movements in oil futures prices in short windows around OPEC production announcements to address the endogeneity of oil prices. In a sufficiently tight window, global economic conditions are unlikely to change, isolating the impact of news about future oil supply.

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 CPI using monthly data from 1974:M1 to 2017:M12. 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). We describe this in more detail in Appendix C.

Our identifying assumption is that unexpected OPEC supply announcements are exogenous to other fundamental drivers of our outcomes, conditional on the SVAR controls. Following Känzig (2021), the SVAR-IV approach is used in the shock-identification step for its additional precision in finite samples (Li, Plagborg-Møller, and Wolf, 2022). Nonetheless, 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 let z_t ber the estimated oil supply news shock. For constant $\mathbf{c}$ and coefficients $\mathbf{A}_j$, 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:

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

Security, unemployment insurance and the G.I. Bill.

[27] Transfer income is small for the majority of the population. Appendix Figure B1 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. Therefore, transfers received by young households have only a small effect on the total welfare effect of inflationary shocks.

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 from the "internal instrument" recursive causal ordering. We store the element of $\mathbf{H}_{\cdot 1}$ corresponding to the response of outcome variable y_t^n as $\mathbf{\Psi}_{n,0}$. The impulse responses for subsequent horizons $\mathbf{\Psi}_{n,h}$ can be computed by propagating the oil-supply-news shock through the VAR model.

The "identified shock" view is a product of our setting's differing outcome variable sample lengths. While most outcome variables have long samples, some are shorter. 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-surprise series. In contrast, a procedure that combines shock estimation and IRF computation is constrained by the outcome variable's available sample.

We follow an analogous 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 CPI. The instrument for the one-year government bond rate—the three-month-ahead monthly Fed futures surprises—is updated using data from Gürkaynak, Karasoy-Can, and Lee (2022). Mirroring the Gertler and Karadi (2015) baseline specification, the reduced-form VAR is estimated using data from 1979:M7-2019: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 estimate IRFs for four years in all of our applications.[28] Estimation of effects over longer horizons is challenging in time series contexts. Our exercise is thus best-equipped to study the short-run effects of inflationary shocks. We explore longer-run effects in Section 8. Standard errors for the impulse responses are computed using a moving block bootstrap (Jentsch and Lunsford, 2019). A description of our approach to estimating standard errors for our welfare calculations is provided in Appendix C.2.

5. Estimated IRFs

This section reports the IRFs to monetary and oil price shocks.

[28]We do not estimate IRFs for business wealth as because scant data on these assets' returns or prices are available.

Shocks and Aggregate CPI. Figure 1 plots the IRF of our shock series and aggregate CPI to our oil-supply and monetary shocks. 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. We scale the size of the shock to represent a 10% increase in the West Texas Intermediates (WTI) crude oil price, because the standard deviation of monthly oil price growth is around 10%. 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 42 basis points after two quarters before converging back to the pre-shock path for the aggregate price index. This finding is consistent with Känzig (2021)'s findings for the aggregate economy.[29]

Panels C and D of Figure 1 plots the estimated response of the one-year treasury yield (Panel C) and aggregate CPI-U (Panel D) in response to the monetary shock. We scale the shock to represent a 25 basis point decline in the one-year treasury yield, which is a common adjustment in the Federal Funds Rate. The initial 25 basis point decline in treasury yields gradually dissipates over the subsequent two years. We have scaled the oil and monetary shocks such that they generate a similar impact response of aggregate inflation: on impact, the 25 basis point decline in nominal interest rates generates an increase in the aggregate CPI-U of 15.6 basis points, which rises to 55 basis points after two quarters.

Consumption Prices. The path of aggregate CPI masks rich heterogeneity in the price responses of different goods. To visualize the effect of our inflationary shocks on disaggregated goods prices, Figure 2 presents coefficient plots of impulse responses for all of our disaggregated CPI subcategories measured at a 24-month horizon. Panel A plots the response to an oil price shock, and Panel B plots the response to monetary shocks.

Consider the responses to the oil shock first. There is substantial heterogeneity in the response of consumer prices to the oil shock. Panel A shows that, intuitively, motor fuel experiences by far the largest increase in response to the crude oil price shock, with a price increase of around 4%, 10 times larger than the response of aggregate CPI. The next largest price movements come from "fuel and utilities", "information technology, hardware and services", and "public transportation". All of these goods rely heavily on energy in production. By contrast, the price of goods such as medical care, recreation or education—which do not have a large energy cost share in production—show no response to the oil shock.

Panel B shows a similar pattern for monetary shocks. We estimate that monetary expansions most affect the price of motor fuel and fuel and utilities. This finding could reflect a relatively small degree of nominal price stickiness or inelastic short-run supply curves in these sectors, or US nominal demand increases affecting world oil prices. Thus, the extent to which inflationary oil supply and monetary shocks affect household well-being through the consumption channel will be primarily determined by household expenditures on motor fuel.[30]

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

[30] We plot the full impulse response for motor fuel and fuel and utilities in Appendix Figure F1.

FIGURE 1: Impulse Responses for Shocks and Aggregate CPI-U

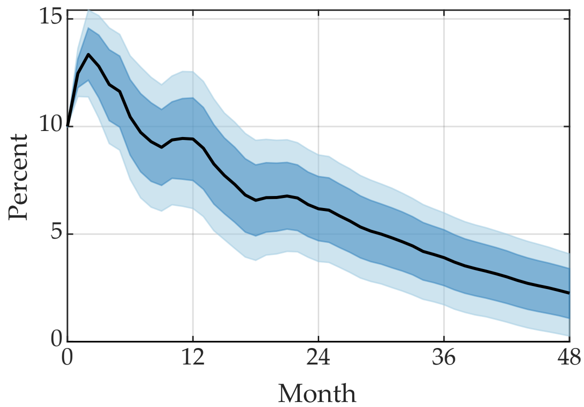

PANEL A: OIL SHOCK, REAL WTI CRUDE OIL PRICE

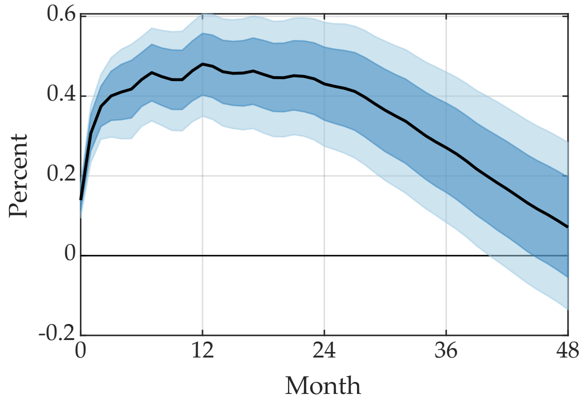

PANEL B: OIL SHOCK, AGGREGATE CPI-U

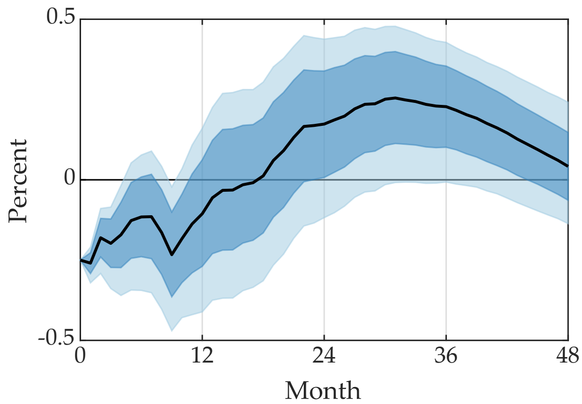

PANEL C: MONETARY SHOCK, 1-YEAR TREASURY BOND YIELD

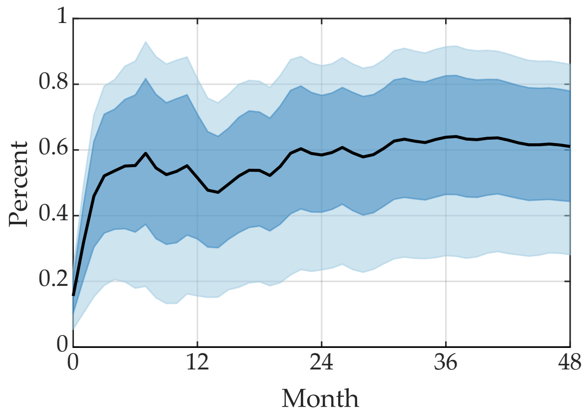

PANEL D: MONETARY SHOCK, AGGREGATE CPI-U

Notes: Figure plots cumulative impulse response functions (IRFs) to our shocks. Panels A and B plot responses to inflationary oil supply news shocks constructed by Känzig (2021). Oil shocks are normalized to represent a 10% increase in the Real West Texas Intermediates (WTI) crude oil price in high-frequency windows around OPEC supply announcements. Panels C and D plot responses 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. Dark blue regions specify the 68% confidence interval, and light blue regions specify the 90% confidence intervals.

Labor Income. Figure 3 plots the estimated response of log labor income to the oil supply shock (Panels A, C, and E) and monetary shock (Panels B, D, and F) for our three education groups. We aggregate all age groups together for these plots. The figure shows that the oil supply contraction leads to a reduction in labor earnings. Note, however, that the decline in earnings from the oil shock does not exhibit large differences between our three education groups. Indeed, two years after the shock, high school or less households' income declines by 0.31 log points, while some college households decline by 0.41 log points and college-educated households see a decline of 0.24 log points.

While the consumption channel was similar across inflationary oil price shocks and monetary shocks, the labor-income channel is quite different. While oil price increases lead to declines in earnings, inflationary monetary shocks have the opposite effect. A 25 basis point cut in interest

FIGURE 2: 24-Month Response of Disaggregated CPI Prices to Oil Price and Monetary Shocks

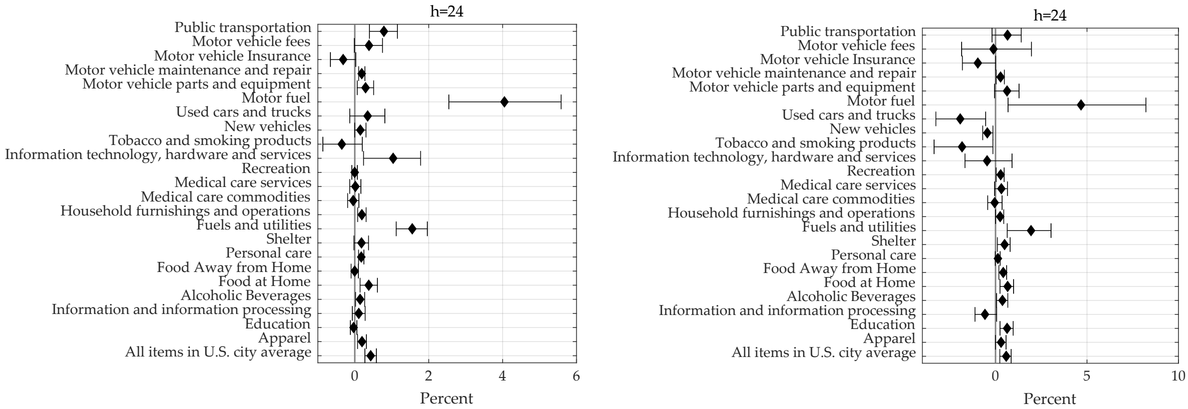

PANEL A: OIL SHOCK PANEL B: MONETARY SHOCK

Notes: Figure plots cumulative impulse response functions (IRFs) of CPI consumption goods 24 months after inflationary oil-supply-news shocks (Panel A) and monetary shocks (Panel B). Shocks normalized to represent a 10% increase in the Real West Texas Intermediates (WTI) crude oil price or a 25 basis point reduction in the one-year treasury yield. IRFs estimated using the "internal instrument" 12-lag SVAR procedure explained in Section 4. Error bars represent 90% confidence intervals.

rates leads to a 1.11 log point increase in earnings for those with at most a high school education after two years. There is a somewhat smaller response for those with some college (0.84 log points) and those with a college degree or more (0.73 log points). These patterns indicate that the labor-income channel will push toward a welfare gain from expansionary monetary shocks that is smallest for those with some college education, opposite to the oil-price shock.

Portfolios. Figure 4 shows the response of asset prices and dividends 24 months after the shock impact for both the oil supply (Panel A) and monetary shocks (Panel B).[31] Focusing first on Panel A, we estimate that most asset prices and dividend series do not respond to the oil shock. However, the value of the S&P500 declines by almost 3% two years after the initial oil price shock. This decline is partially accounted for by dividend payments, which fall by around 1%. Meanwhile, we find essentially no effect on any other asset price. 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.

Equity holders lose from the oil price shock because they receive lower dividend income. Importantly, however, the oil supply contraction is beneficial to those who were planning to accumulate equity, because doing so is now cheaper. Thus, the strength of the portfolio channel for oil price shocks depends critically on who holds and who is accumulating equities.

Panel B shows that many of these patterns are flipped for monetary shocks. Interest rate reductions cause increases in dividends and prices for the S&P500. These effects are large, mirroring previous work: a 25 basis point decline in interest rates leads to a 3.5 percentage point increase in stock prices on impact and nearly a 5 percentage point increase after 24 months, though

[31] We plot the full impulse response of key asset prices in Appendix Figures F2 and F3.

FIGURE 3: Impulse Responses of Labor Income to Shocks

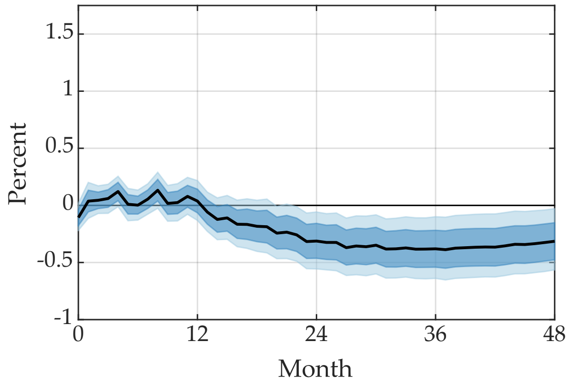

PANEL A: OIL SHOCKS
HIGH SCHOOL OR LESS

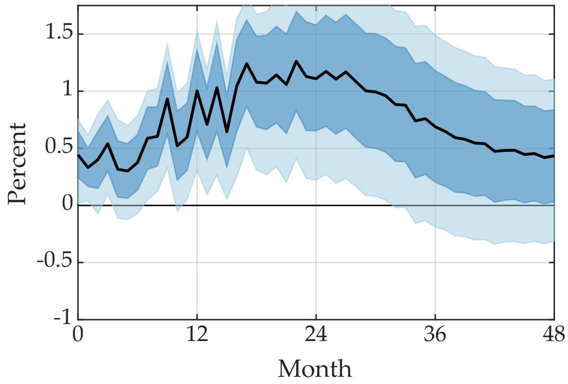

PANEL B: MONETARY SHOCKS
HIGH SCHOOL OR LESS

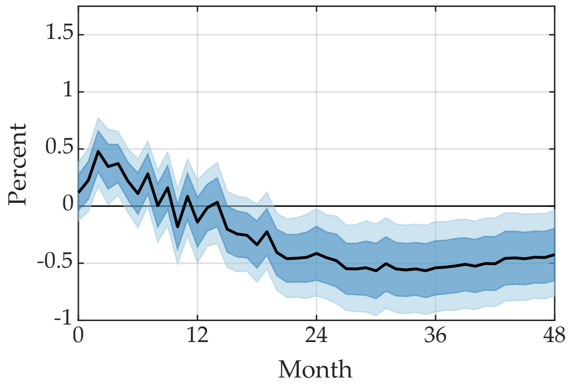

PANEL C: OIL SHOCKS
SOME COLLEGE

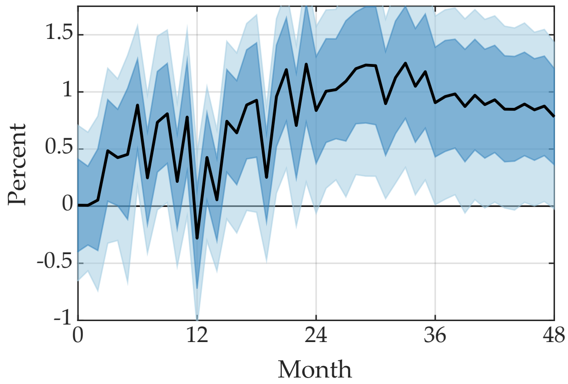

PANEL D: MONETARY SHOCKS
SOME COLLEGE

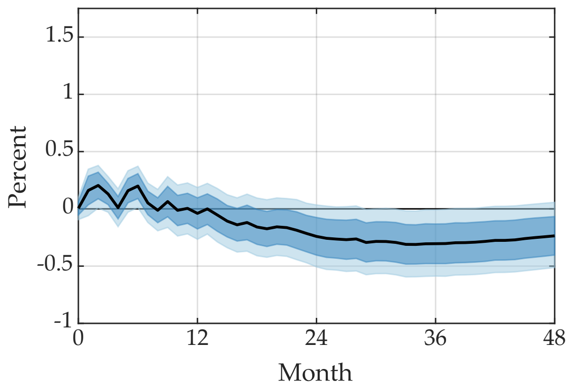

PANEL E: OIL SHOCKS
BACHELOR'S +

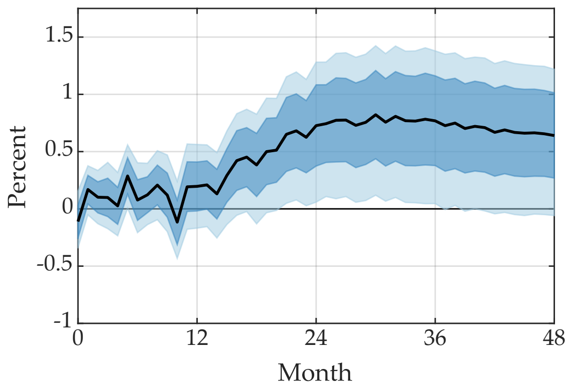

PANEL F: MONETARY SHOCKS
BACHELOR'S +

Notes: Figure plots cumulative impulse response functions (IRFs) of log average weekly earnings for households in our three education groups to inflationary oil supply news shocks (Panels A, C and E) and monetary shocks (Panels B, D and F). Earnings data constructed from the Outgoing Rotations Group (ORG) of the Current Population Survey (CPS). Dark and light blue regions represent 68% and 90% confidence intervals, respectively.

confidence bands are wide. In addition, declines in interest rates lead to gradual increases in house prices, which peak at around a 2.5 percentage point increase after 3 years.

The portfolio channel for monetary shocks principally benefits those who are selling equities

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

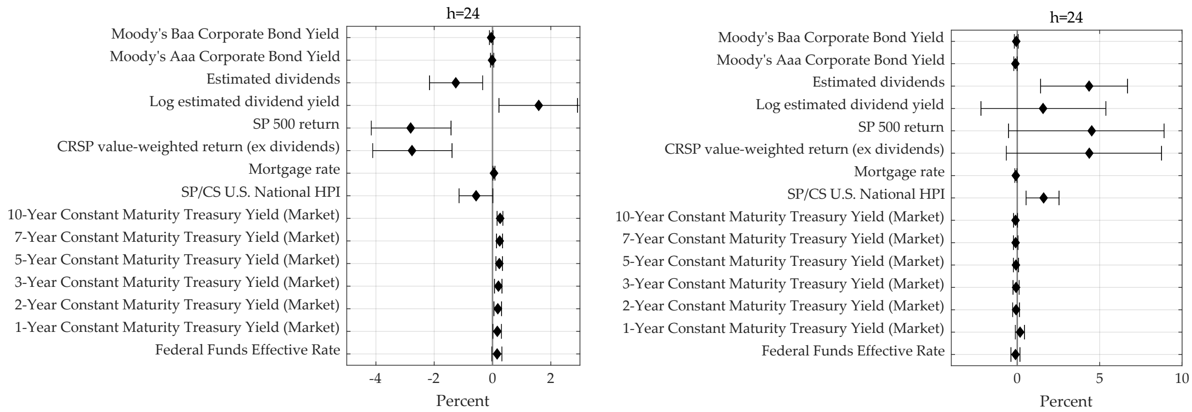

PANEL A: OIL SHOCK

PANEL B: MONETARY SHOCK

Notes: Figure plots cumulative impulse response functions (IRFs) of asset prices and dividends 24 months after inflationary oil supply news shocks (Panel A) and monetary shocks (Panel B). Shocks normalized to represent a 10% increase in the Real West Texas Intermediates (WTI) crude oil price or a 25 basis point reduction in the one-year treasury yield. Error bars represent 90% confidence intervals.

or their home. This contrasts with oil shocks, which benefit those who would *buy* equities. The different behavior of these asset prices will turn out to drive the quantitative differences in welfare responses of these two shocks.

6. RELEVANT FEATURES OF CROSS-SECTIONAL DATA

Figure 5 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 the literature (Browning and Lusardi, 1996; Attanasio, 1999).

Panel B shows the share of household consumption expenditures that are accounted for by the two categories most responsive to oil and monetary shocks. The figure shows that the least educated households 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.

Turning to the portfolio channel, Figure 6 plots asset holdings and accumulation patterns over the life cycle by education. The largest share of assets for most households is housing. Equities constitute a larger share of assets for older households and those with at least a college education. The dominance of housing in portfolios illustrates the importance of accounting for its

FIGURE 5: Life-cycle Consumption Expenditures and Shares

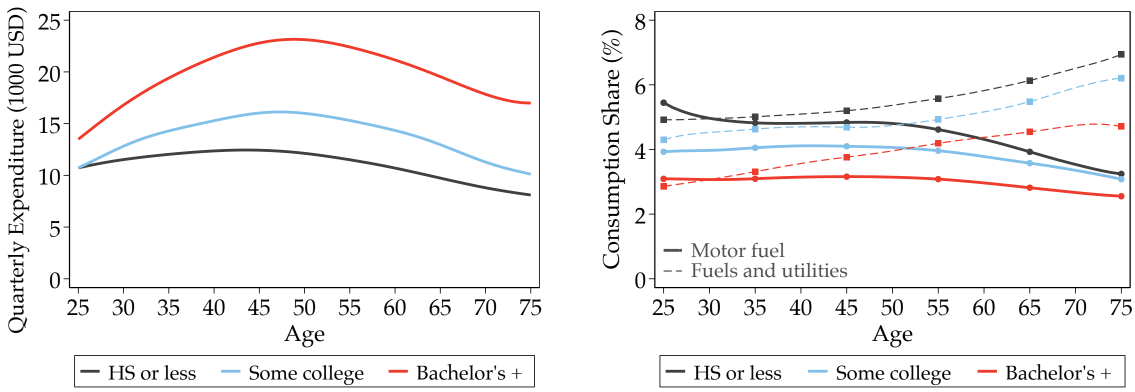

PANEL A: TOTAL CONSUMPTION EXPENDITURES

PANEL B: CONSUMPTION SHARE OF MOTOR FUEL AND UTILITIES

Notes: Panel A plots total quarterly expenditure by group, and Panel B plots shares of expenditure on motor fuel and fuel and utilities 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 by age.

durable nature as both a consumption good and store of value (see Appendix A.3).

Panel B plots the accumulation of equity over the life cycle. All education groups accumulate equity during middle age, before slowing accumulation or decumulating after the retirement age.[32] This hump-shaped accumulation pattern is especially strong for those with a college education. 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. Thus, reductions in housing prices will benefit younger college-educated households on average. Older households, however, decumulate housing, and as such house price increases are beneficial for them.

7. MONEY-METRIC WELFARE CALCULATIONS

This section aggregates our estimated IRFs into money-metric welfare effects of inflationary oil price and monetary shocks following Section 2. We do this first in the baseline case which does not rely on parametric utility functions but ignores idiosyncratic risk and borrowing constraints. We then add these features to assess their importance.

[32]The 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 6: Life-Cycle Asset Portfolios and Accumulation

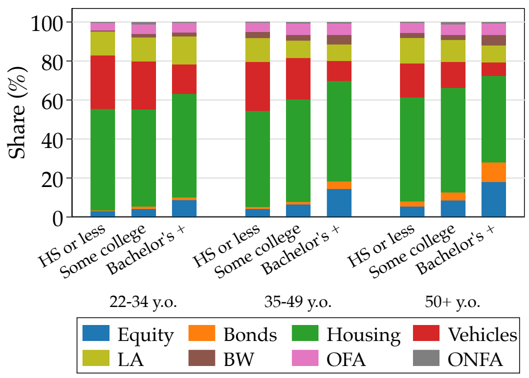

PANEL A: PORTFOLIO SHARES

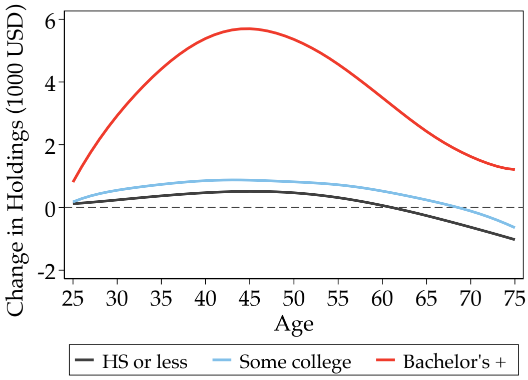

PANEL B: EQUITY ACCUMULATION

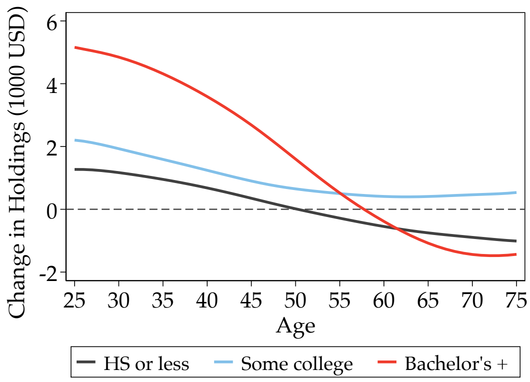

PANEL C: HOUSING ACCUMULATION

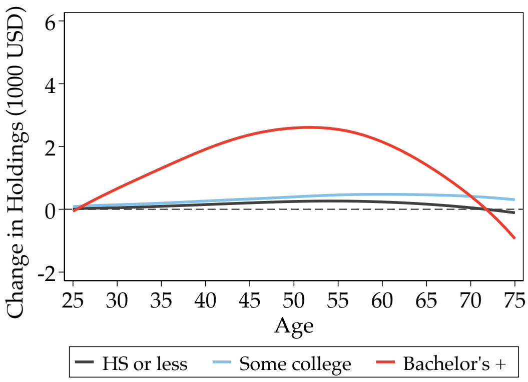

PANEL D: NON-CORP. BOND ACCUMULATION

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. ONFA 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 are from the Survey of Consumer Finances for 2019.

7.1 Baseline Results Without Idiosyncratic Risk or Borrowing Constraints

Welfare Effects of Oil Price Shocks. We estimate IRFs out to a horizon of 16 quarters. Thus, in calculating the welfare formula in equation (5), we limit ourselves to the cumulative effects over a truncated, four-year horizon. In practice, a tradeoff exists between the precision of the estimates and the length of period studied. We extend the period over which welfare changes are calculated in Section 8 below.

Figure 7 plots money-metric welfare losses from a 10% increase in the price of oil brought about by announced oil supply contractions. Panels A through C separately plot welfare losses from the consumption, labor-income and portfolio channels, respectively, while Panel D reports the total welfare loss. We normalize four-year money-metric welfare losses by four-year consumption expenditures to facilitate comparison across groups.[33]

[33]In Appendix Table F1, we report the raw money-metric welfare change in response to an oil shock and a

Panel A shows that oil supply contractions harm all households through the consumption channel. In addition, the least-educated households lose the most as a share of consumption. However, these differences are small: were consumption prices the only thing to respond to oil supply shocks, the least-educated households must be paid around 0.23% of four-year consumption to be made whole, whereas college-educated households must be paid around 0.1%. There is also little difference over the life cycle in the welfare losses from the consumption channel. These patterns reflect those presented above: while less educated households spend a larger share of their consumption bundle on motor fuel and fuel and utilities, these differences are quantitatively small.[34]

A similar pattern arises when one focuses solely on labor income (Panel B). The labor-income channel reduces less educated households' welfare by around 0.12%, and that of college-educated households by around 0.13%, a negligible difference. Note, however, that the labor-income channel does have important differential effects over the life-cycle. Retired households, for example, are unaffected by the labor-income channel.

The greatest difference among groups is in the portfolio channel. Panel C shows that this channel is negligible for young low-education workers, but large and *positive* for middle-aged high-education workers. That is, middle-aged high-education workers see fairly large welfare gains as a result of the portfolio channel. This is because high-education households accumulate equity in the middle of their life, so that temporarily falling equity prices are beneficial to them. 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.[35]

Finally, Panel D adds all the channels together to a composite welfare change following Proposition 1.[36] We find that young and middle-aged households with no more than a high school degree must be paid 0.4% of pre-shock consumption expenditure to achieve the same utility, reflecting primarily the consumption and labor-income channels. This is relatively flat across the life cycle. In stark contrast, younger college-educated households gain 0.15% of consumption relative to the no-shock baseline. Younger college-educated households particularly gain because they are equity-accumulators on net, but older high-education households see small losses reflecting their lower dividend receipts.[37]

Inflationary oil supply contractions therefore appear highly regressive: the least educated households suffer sizable welfare losses, while the most educated households see large wel-

monetary policy shock for our education groups across three age groups.

[34] Cravino and Levchenko (2017) find larger distributional consequences of consumption price changes around the devaluation of the Mexican peso. However, Ferreira et al. (2024) also find small differences in the "consumption channel" during the COVID inflation in Spain.

[35] The importance of the portfolio channel reflects the finding of Lanteri and Rampini (2023) that distributive externalities – which reflect transfers between net buyers and sellers of assets – are a large driver of the pecuniary externalities stemming from asset prices.

[36] We present estimates of the transfer channel in Appendix Figure F6 and find it is relatively small.

[37] Note 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. We explore this possibility in Section 8.

FIGURE 7: Welfare Losses From Inflationary Oil Price Shocks

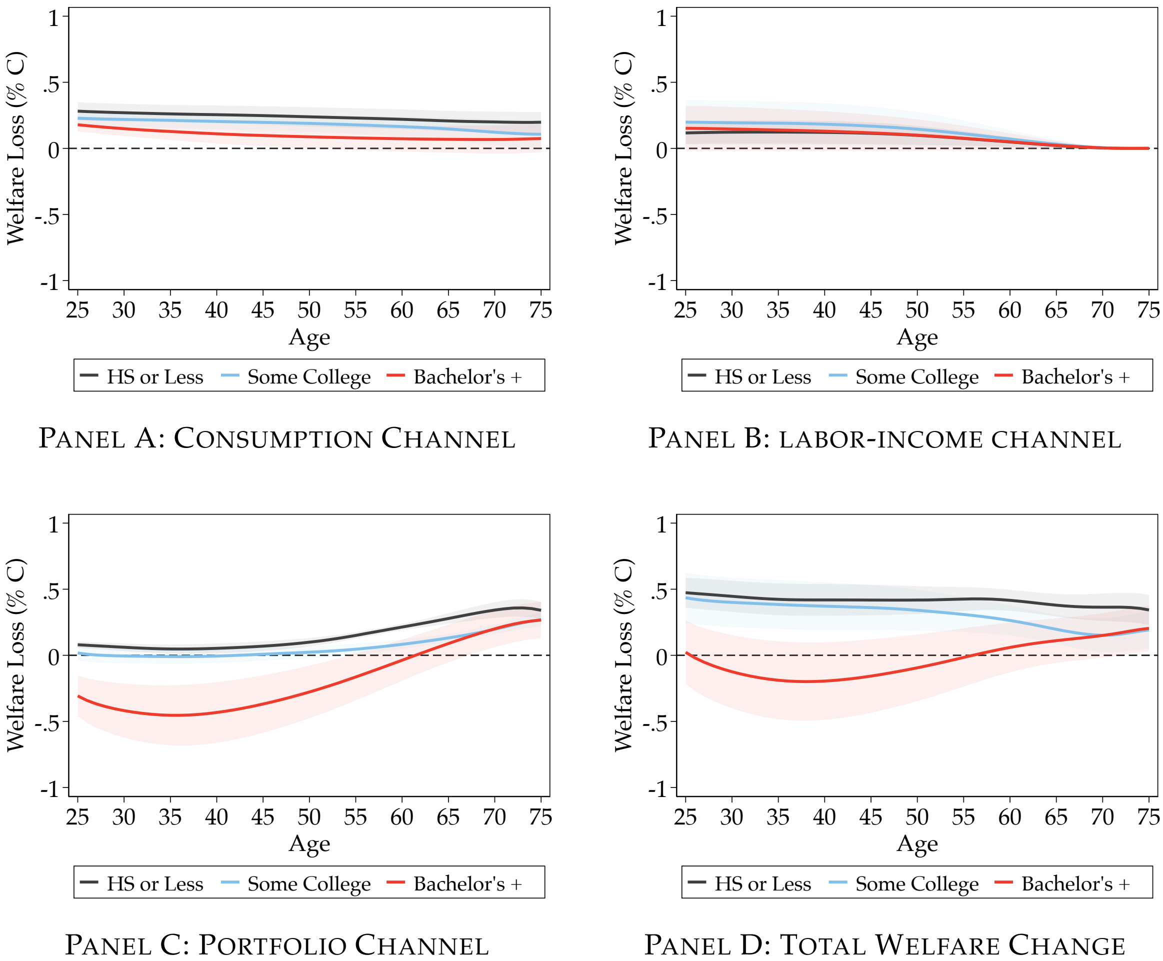

Notes: Figure shows the estimated welfare loss from a 10% positive oil price shock. We normalize the figures by total four-year consumption by age and group, using projected life-cycle consumption patterns in 2019. A negative number represents a welfare gain. Shaded regions represent 90% confidence bands.

fare gains for much of their life cycle.[38] This is mostly due to the portfolio channel: declines in equity prices allow middle-aged college-educated households to save at a lower cost.

Welfare Effects of Monetary Shocks. Figure 8 plots the four-year money-metric welfare losses from a surprise 25 basis point reduction in interest rates caused by monetary policy announcements, scaled by four-year consumption. As with the oil shock, Panel A shows that, in proportional terms, households of all ages and education groups have a similar loss from monetary shocks through the consumption channel. However, households now gain from the labor-income channel. These two channels roughly offset each other.

Contrary to the oil shock, the portfolio channel strongly pushes towards monetary policy shocks being progressive. Panel C shows that the portfolio channel leads to a loss of 0.8% of consumption for middle-aged college-educated households. Again, we see an important life-

[38]The differences between high school or less and bachelor's plus welfare changes are statistically significant until around 55 years of age (see Appendix Figure F9).

FIGURE 8: Welfare Losses From of Inflationary Monetary Policy Shocks

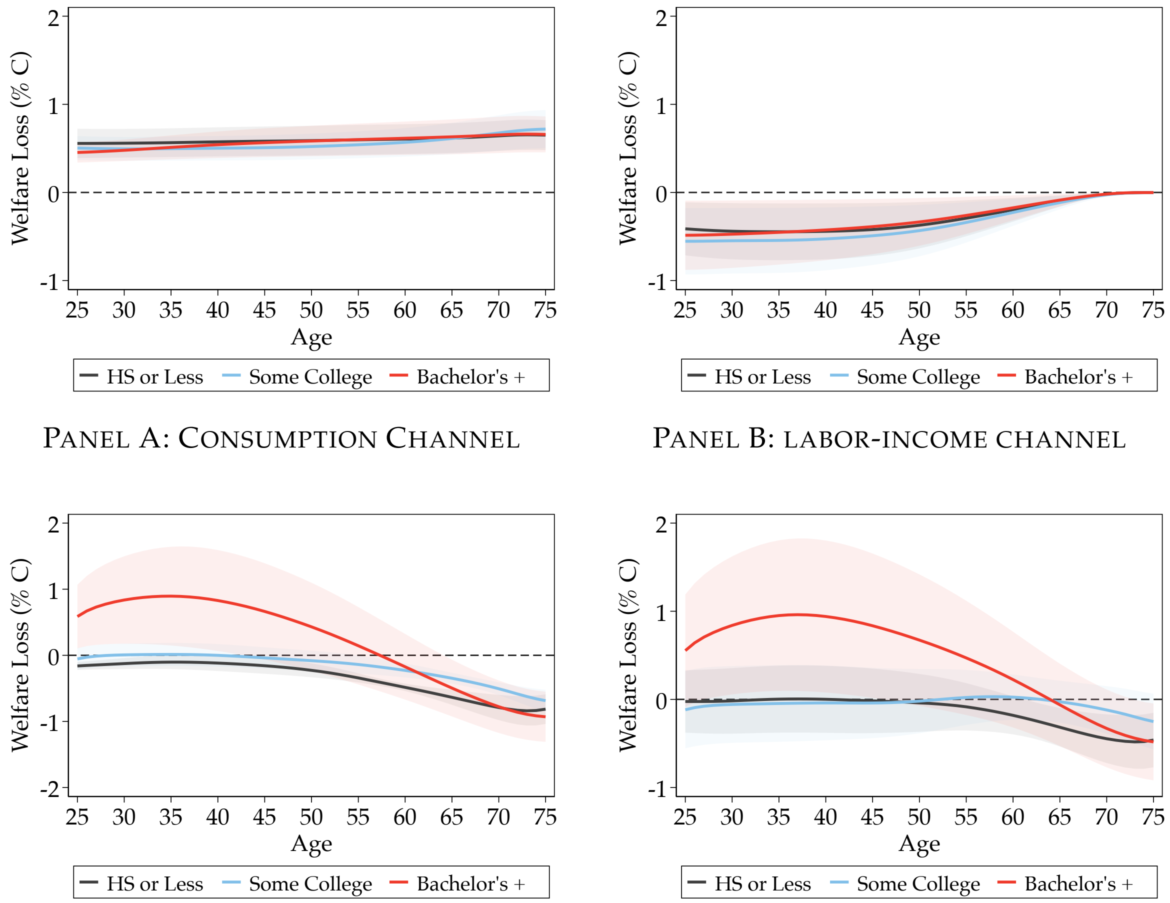

PANEL A: CONSUMPTION CHANNEL

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 normalize the figures by total four-year consumption by age and group, using projected life-cycle consumption patterns in 2019. A negative number represents a welfare gain. Shaded regions represent 90% confidence bands.

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, although 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. Finally, a countervailing force arises through asset income: older college-educated households own more equities and benefit from the increased dividend payouts as a result of expansionary monetary policy.

Combining these effects together reveals that inflationary shocks to monetary policy are on balance progressive. The movements in all prices net to approximately zero for households with at most a high-school education, whereas young and middle-aged high-education households see losses of 0.8% 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 slightly benefit low-education households and hurt high-education households. Rate increases would do the opposite. 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 shock.

7.2 Incorporating Idiosyncratic Risk

To incorporate idiosyncratic risk, one must estimate the adjustment factors Θ_t^x from the data. To do so, we make parametric assumptions on utility functions. We assume that household utility is $U(C, L, \{N_k\}) = \log C + h(L, \{N_k\})$, so that household preferences are log-separable in consumption. We additionally assume that household consumption aggregators are homothetic within period, but may still differ across groups and ages. Under these assumptions, the adjustment factor for a variable x becomes

$$\Theta_t^x = Cov\left(\frac{1/P_tC_t(s_t^i)}{\mathbb{E}[1/P_tC_t(s_t^i)]}, \frac{x_t(s_t^i)}{\mathbb{E}[x_t(s_t^i)]}\right).$$

Therefore, given data on consumption expenditures and the variable x, one can estimate the covariance between marginal utilities of consumption and the x variable across realizations of the idiosyncratic state s_t^i. To do so, we simply compute the cross-sectional covariance between the reciprocal expenditures and x_t. For instance, in the case of Θ_t^W, we estimate the covariance of the reciprocal of consumption expenditure with labor income in the consumer expenditure survey, within age $\times$ education bins.[39] This exercise likely overestimates the magnitude of the covariance between marginal utilities and the variable x, because the cross-sectional covariance additionally includes permanent unobserved heterogeneity within our education x age groups. Therefore, our approach gives the largest chance for risk to affect our results.

We present and discuss our estimated values for Θ^x in Appendix F. Nearly all of the estimated values of Θ^x range between -0.2 and -0.4, except for transfer income. To assess the magnitudes of these Θ's, one can consider two extreme cases. First, if there is perfect consumption insurance against idiosyncratic risk, then there is no variance in consumption and all Θ's are equal to 0. Second, suppose households are fully hand-to-mouth, in which case income shocks pass through to consumption one-for-one. In this case, Θ^x is equal to the covariance of x with $1/x$. This ranges from around -0.5 to -1, depending on the variable. Generally, the estimated Θ's are around one half of this hand-to-mouth benchmark, indicating a moderate degree of consumption insurance. These Θ's are also approximately the same magnitude as what would be implied by a workhorse two-asset HANK model, as we discuss in Section 9 below.[40]

[39]To maximize power, we group households into 5-year age bins for this exercise. We additionally assume a constant value for Θ^x over time. This likely overstates the importance of risk, because risk-adjustment factors are zero in periods for which the state is known, such as period 0. In the HANK model considered below, this use of a constant Θ^x over time did not meaningfully affect the performance of our feasible set approach. Ideally, one would use long-run individual panel data to estimate these Θ, but such data are not available in our context.

[40]Our estimates also align well with similar objects estimated in recent work on Italian household data in Krueger, Malkov, and Perri (2023).

FIGURE 9: Welfare Losses Accounting for Risk and Borrowing Constraints

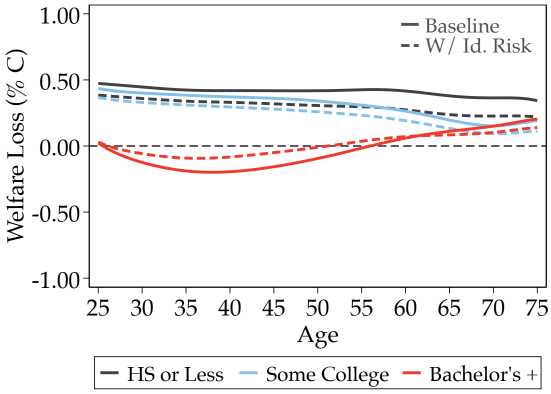

PANEL A: OIL SHOCKS WITH RISK

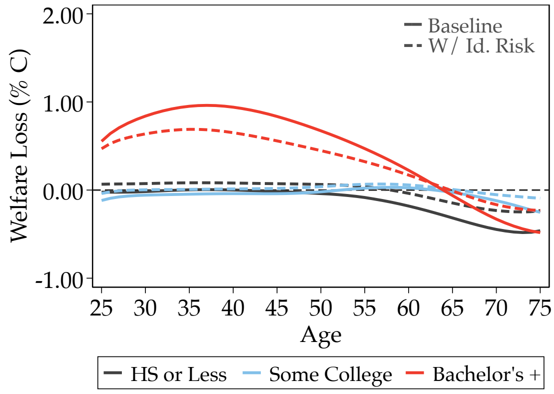

PANEL B: MONETARY SHOCKS WITH RISK

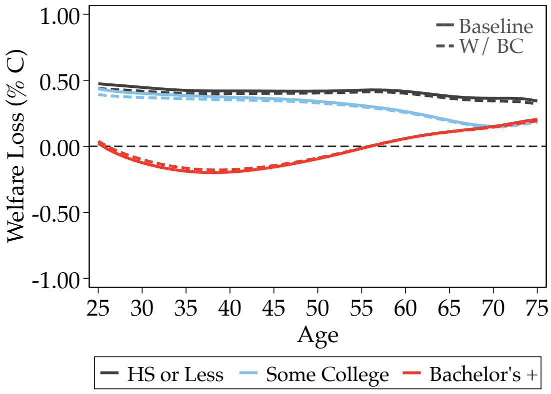

PANEL C: OIL SHOCKS WITH BC

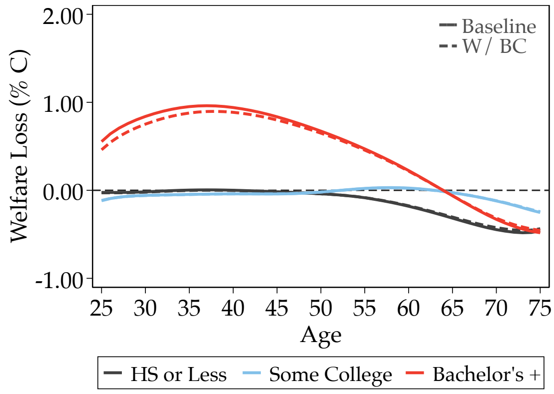

PANEL D: MONETARY SHOCKS WITH BC

Notes: Figure shows the estimated welfare loss from a 10 percentage point increase in the price of oil (Panels A and C) or a 25 basis point cut to the federal funds rate (Panels B and D) after accounting for idiosyncratic risk (Panels A and B) or borrowing constraints (Panels C and D). We divide the figures by total four-year consumption by age and group, using projected life-cycle consumption patterns in 2019.

The top row of Figure 9 reports our estimated welfare effects accounting for idiosyncratic risk. The shape and magnitude of the welfare effects is extremely similar for both oil shocks and monetary shocks across the age distribution. However, the magnitude of the welfare effects are generally a little smaller, due to the additional discounting implied by the riskiness of future income streams and price changes.

Clearly, the degree of risk aversion will matter for the quantitative effects estimated here. In Appendix G we study the forces that affect the risk-adjustment shifters Θ's. In addition, we provide a comparison of the estimated shifters when we move from log separability in consumption to CRRA-separable with a relative risk aversion coefficient of 2. This effectively doubles the estimated shifters, reducing the estimated welfare losses due to greater discounting of future labor income steams. However, the qualitative patterns of regressivity and progressivity presented in Figure 9 and further above are unaffected.

7.3 Incorporating Borrowing Constraints

We follow Proposition 3 to compute the welfare effects including borrowing constraints. To implement this formula, we require an estimate of the borrowing constraint wedges τ. To do so, we maintain the assumption that utility is log-separable in consumption and that the consumption aggregator is homothetic, so that equation (9) can be written

$$1 + \tau_t = \frac{P_{t+1}C_{t+1}}{P_t C_t} \left(\frac{\beta_t}{\beta_{t+1}} \right) \left(\frac{\delta_t}{\delta_{t+1}} \right) \frac{1}{R_t}.$$

We calculate the average τ_t^a by age and type by using annual consumption data from the CEX to compute yearly consumption growth by age for the years 1990 to 2019. We assume a constant discount rate $\beta = (0.98)^{\frac{1}{4}}$, and use empirical death rates taken from the Social Security Administration. The value of τ_t over the life cycle is plotted by education group in Appendix Figure F11. These wedges decline with age, reflecting slower growth in consumption.

If all households of type (a, g) are constrained, Proposition 3 can be applied directly. However, the data show that only a fraction of households have negative net worth; thus, it is unlikely that most households within a group are constrained. As depicted in Appendix Figure F11, younger households are more likely to have negative net worth than are older households.

Because many datasets do not have information about household net worth, computing labor income and consumption patterns for those who appear constrained is not possible. We therefore consider the following approach. From the SCF, we identify households with negative net worth and compute their average holdings ($Q_{kt}N_{kt}$) for each of the asset classes we consider. Then we compute the welfare change according to Proposition 3 using these averages for computing the value of a relaxed constraint. To account for the fact that not all households are constrained, we weight this term by the share of households constrained for that type. We include the discount wedge for all households for this exercise.

Panels C and D of Figure 9 plots the results of this exercise. The dashed lines show the welfare losses from shocks incorporating borrowing constraints, whereas the solid lines are our baseline results ignoring borrowing constraints. Changes in asset prices leads to marginally smaller welfare losses due to this tightened constraint. However, accounting for borrowing constraints makes little quantitative difference. This is because, by definition, people who are constrained have small asset holdings. Therefore, movements in asset prices have a small effect on the value of their holdings relative to the effect that asset-price movements have on the cost of asset accumulation for these households.

We stress that this result does not imply that borrowing constraints are unimportant. As has been emphasized elsewhere, borrowing constraints and the extent to which households can self-insure through savings is extremely important for the *level* of welfare. They are also crucial determinants of the inputs to our formula – both individual consumption behavior and the impulse response of aggregate variables to exogenous shocks (Aiyagari, 1994; Kaplan and

Violante, 2024). Rather, our result suggests that borrowing constraints are not quantitatively important for the effect of these aggregate shocks on household well-being *conditional* on the consumption behavior of households and impulse responses implied by that shock. Note that this may not be true if the shocks in question were extremely persistent, as the discount wedge introduced by constraints may be more quantitatively important for such long-lived shocks. We study these point further in Section 9.

8. 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 the robustness of the paper's conclusions to alternative model specifications and estimation strategies. We probe the size of second-order effects on welfare – such as the importance of substitution patterns – and test robustness to alternative estimation strategies, alternative measurement assumptions, incorporating asset values (rather than quantities) in utility, and the longer-run welfare effects of oil supply and monetary shocks. Table 1 collects the results of these robustness exercises. Additional details of our robustness exercises are contained in Appendix D.

Second-Order Effects. The envelope theorem states that behavior changes in response to the shock are not first-order welfare-relevant. As shocks become larger, this approximation may become poor. We now study second-order welfare effects incorporating behavioral change.

To do so, we need some additional notation. First, define the vector of "square impulses" as

$$\mathbf{\Psi}^2_{n,t} \equiv \mathbb{E}_t[\mathbf{v}_t^2 | \epsilon_0^n = 1] - \mathbb{E}[\mathbf{v}_t^2 | \epsilon_0^n = 0].$$

which denotes the impulse response of the second moment of the error processes to a fundamental shock.[41] These are the additional terms that appear in a second-order welfare expansion. They enter similarly to the first-order terms, and take the form

$$\Xi_2 \equiv \frac{1}{2} \sum_t R^{-1}_{0\to t} \bigg(- \sum_j p_{j,t} c_{jt} \mathbf{\Psi}^{2p,j}_{n,t} + W_t L_t \mathbf{\Psi}^{2W}_{n,t} + T_t \mathbf{\Psi}^{2T}_{n,t} + \sum_k \Big[N_{kt-1} D_{kt} \mathbf{\Psi}^{2D,k}_{n,t} - Q_{kt} \Delta N_{kt} \mathbf{\Psi}^{2Q,k}_{n,t} \Big] \bigg).$$

They represent the welfare loss from additional volatility induced by the shock. In practice, these square impulse response terms tend to be very small. A second, more important set of additional terms contain behavioral elasticities, such as the extent to which consumption of good j responds to changes in wages, consumption prices, or asset prices in different periods. To capture these forces, we first define an additional "cross-impulse" response matrix as

$$\mathbf{\Psi}^X_{n,t,s} \equiv \mathbb{E}_t[\mathbf{v}_t \mathbf{v}_s' | \epsilon_0^n = 1] - \mathbb{E}[\mathbf{v}_t \mathbf{v}_s' | \epsilon_0^n = 0].$$

Lastly, let $\omega_t \in \Omega_t \equiv \{\{p_{jt}\}_j, \{Q_{kt}\}_k, \{D_{kt}\}_k, W_t, T_t\}$ denote an element of the set of all external

[41] See Appendix C.3 for details on how we estimate these.

variables in the agent's budget set. We then have an additional term given by:

$$\Xi_3 \equiv \frac{1}{2}\sum_t R_{0\to t}^{-1}\Bigg(-\sum_j p_{j,t}c_{jt}\sum_s\sum_{\omega_s\in\Omega_s}\frac{dlnp_{jt}c_{jt}}{dln\omega_s}\Psi_{n,t,s}^{p,j\omega_s} + \sum_s\sum_{\omega_s\in\Omega_s} W_tL_t\frac{dlnL_t}{dln\omega_s}\Psi_{n,t,s}^{W\omega_s}$$
$$+\sum_k\left[N_{kt-1}D_{kt}\sum_s\sum_{\omega_s\in\Omega_s}\frac{dlnN_{kt-1}}{dln\omega_s}\Psi_{n,t,s}^{D,k\omega_s} - |Q_{kt}\Delta N_{kt}|\sum_s\sum_{\omega_s\in\Omega_s}\frac{dQ_{kt}\Delta N_{kt}/|Q_{kt}\Delta N_{kt}|}{dln\omega_s}\Psi_{n,t,s}^{Q,k\omega_s}\right]\Bigg).$$

This term captures the response of household choices to price changes, multiplied by the size of the price changes induced by the shock. If prices change for goods with highly elastic demand, this term might be sizable, even if the cross-impulse response functions are small.

A final term modifies the original first-order expansion, which indexes how the marginal utility of a dollar at time 0 changes as the setting becomes more uncertain. We collect all these pieces in the following proposition. Let dV_2 denote the second-order welfare response to a shock, and dV as above is the first-order expansion explored in Proposition 1.

Proposition 4. *As $\sigma \to 0$, the second-order change in money-metric welfare dV_2 from an impulse to an element n of the fundamental shock vector at $t = 0$ is*

$$dV_2 = dV(1 + \frac{1}{2}\frac{d\lambda_0}{d\sigma}) + \Xi_2 + \Xi_3, \tag{11}$$

where choices are evaluated at $\sigma = 0$.

Computing Ξ_2 is straightforward from the cross-sectional data and impulse responses used in the calculation of Proposition 1. However, a full estimation of all the behavioral elasticities in Ξ_3 is daunting. This would require estimating elasticities for each consumption category, each asset-class holding, labor supply, and accumulation patterns with respect to all prices, and at all possible time horizons. This is effort is beyond the scope of this paper.[42]

Instead, to understand the potential size of second-order effects, we proceed as follows. We start by setting non-contemporaneous elasticities to zero, in order to size the magnitude of contemporaneous terms. Generally, we expect cross-time impacts (consumption today in response to future price changes, or consumption tomorrow in response to today's price changes) to be much smaller than contemporaneous responses. We then take values for other elasticities either from the literature, or to maximize the size of the second-order effects. Our exact approach is detailed in Appendix D.1. Briefly, we assume a demand for consumption goods follows a constant elasticity of substitution (EOS) demand system, with an EOS of 4 (as in Hottman, Redding, and Weinstein, 2016). We alternately assume households spend or save all of their additional income in order to maximize the size of the behavioral elasticities. We take a labor supply elasticity of 2 to be in line with the larger values seen in the macro literature, and assume that labor supply responds to income shocks with an elasticity of -0.2 following lottery-based evidence of Cesarini, Lindqvist, Notowidigdo, and Östling (2017). The response

[42]Many of these elasticities have yet to be estimated in the literature. For example, Auclert, Rognlie, and Straub (Forthcoming) consider the matrix of "intertemporal marginal propensities to consume", which here would be the terms denoting how consumption responds to shocks to transfers in different periods. They note that only the first column of this matrix has been estimated empirically, and use a model to discipline the other columns.

TABLE 1: Robustness of Estimated Welfare Effects of Oil and Monetary Shocks

	Oil Supply News Shock			Monetary Policy Shock		
Specification	$\leq$ HS (1)	Some College (2)	College+ (3)	$\leq$ HS (4)	Some College (5)	College+ (6)
Baseline	-0.422% (-$819)	-0.336% (-$830)	+0.084% (+$304)	+0.080% (+$141)	+0.024% (+$58)	-0.654% (-$2308)
Second Order	-0.423% (-$821)	-0.336% (-$830)	+0.080% (+$291)	+0.071% (+$124)	+0.020% (+$49)	-0.678% (-$2395)
Project No-Shock Choices	-0.420% (-$845)	-0.322% (-$826)	+0.138% (+$513)	+0.047% (+$79)	-0.022% (-$59)	-0.744% (-$2734)
Asset Values in Utility Function	-0.378% (-$736)	-0.279% (-$691)	+0.264% (+$955)	-0.024% (-$56)	-0.108% (-$269)	-1.077% (-$3835)
Long-Run Welfare Effects (% Consumption up to 80 y.o.)	-0.133% (-$1509)	-0.076% (-$1360)	+0.050% (+$1656)	+0.003% (+$286)	-0.039% (-$535)	-0.183% (-$5506)

Notes: This Table reports the robustness of our baseline welfare results to a variety of specifications. All columns are life-cycle averages. Welfare effects are reported as a share of four-year consumption for the first four exercises and as a share of lifetime consumption for the final exercise. Money-metric gains are reported in parentheses.

of labor supply to asset prices is taken to be 0.035 based on the evidence of Chodorow-Reich, Nenov, and Simsek (2021). The magnitude of the elasticity of asset demand to prices and dividends is taken to be 0.02 based on the estimates of Gabaix, Koijen, Mainardi, Oh, and Yogo (2023). Finally, we assume that $d\lambda_0(\sigma = 0)/d\sigma \approx 0$.

The second row of Table 1 reports the results of computing Proposition 4 under these assumptions. Quantitatively, the welfare estimates are little changed. Including non-contemporaneous effects in the calculation would increase the difference from the baseline, but if (as we expect) they are smaller than contemporaneous effects calculated above, this should not change the qualitative conclusion. That, we stress that the feasible set approach is best for measuring first-order effects.

Project No Shock Choices. In our baseline, we assume that the expected consumption expenditures, wages and assets holdings by age and group are stable absent shocks. Of course, in the background there is real growth in these variables as the aggregate economy expands. For robustness, we examine the effects on our results of projecting out into the future log-linear time trends in these variables, so that, for example, the expected expenditures on good j at time $t+h$, $p_{jt+h}c^{ag}_{jt+h}$ is not simply evaluated at $p_{j,t}c^{a+h,g}_{j,t}$, but grows. We explain our procedure for doing so in Appendix D.2. Our results, presented in the third row of Table 1, are similar to the baseline, but somewhat larger for highly educated households' response to the oil shock, reflecting growth over time in asset accumulation.

Asset Values in Utility. Our baseline formulation supposes households may have preferences over the *quantity* of assets, held as a stand-in for households' liquidity preferences, or preferences over the service flows of durable goods such as housing and cars. One might equally suppose that households derive utility from the *value* of assets that they hold. For instance, households may hold valuable assets due to a bequest motive. We show in Appendix

D.3 that, in the small-noise limit, the welfare effects of shocks is the same as in Proposition 1, with one additional term reflecting the utility gain from higher asset values, which is given by:

$$\sum_t R_{0\to t}^{-1}\left(1 - R_{0\to t}^{-1}\mathbb{E}_0\left[\frac{Q_{kt+1}}{Q_{kt}}\right]\right) Q_{kt} N_{kt} \Psi_{n,t}^{Q,k}.$$

If households enjoy owning high-value assets, an increase in the price of assets they hold increases their utility by the product of their marginal utility of the asset and the change in the value of the asset. Optimal portfolio choice implies that the marginal utility of the asset must be tied to the excess return on the asset: if a household holds a relatively low-return asset, they must derive utility from that asset. Thus, one can write the marginal utility of asset values as a function of returns – which are measurable – in the limit as aggregate risk becomes small. Note, however, that this derivation assumes that households do not value the riskless bond.[43]

We calculate this term assuming that households expect asset-price appreciation to be the average appreciation over our sample. As shown in the fourth row of Table 1, doing so has little impact on our results for less educated households, but has some effect on the more educated households because they have larger asset holdings.

Long-Run Welfare Effects. Our baseline results above present discounted welfare effects over a four-year forward horizon. The primary reason for this was the precision of our IRF estimates; the standard errors tend to become large past a horizon of four years. To project beyond this horizon, we must make an assumption about what happens to the values of $\mathbf{v}_t$ after four years. Our baseline results can be read as assuming these values immediately return to zero. We now consider the case where they stay forever at the values recorded after 48 months. Here, the price movements induced by the fundamental shock are permanent. We then recalculate our welfare formula up until 80 years of age for all of our household types.

Our key qualitative findings are unchanged in this exercise, but the magnitudes of the effects are larger in money-metric space and smaller as a share of lifetime consumption. Those with at most a high school education would be willing to pay 0.133% of lifetime consumption to avoid a 10% oil price shock, whereas college-plus households gain around 0.05% of lifetime consumption. For the monetary shock, the least educated households are roughly unaffected by the shock, while college-plus households would be willing to pay 0.183% of consumption to avoid the shock on average – this is around one-quarter the size of the four-year consumption effect. The change in magnitudes arises from life-cycle savings behavior: households who accumulate assets today will eventually sell those assets (though at a discounted point in the future) so are less harmed by permanent asset-price appreciation.

[43] Away from the zero-risk world, high-return assets need not be associated with such a disutility, because holdings depend on the return structure of the asset and its covariance with marginal utilities. Whether our framework accounts for such motives depends on the quality of the linearization around the economy with $\sigma = 0$.

9. Validation of Approach and Lessons For Models

Our final exercise is to validate the accuracy of our approach in a setting where we know the true welfare responses to shocks: a state-of-the-art heterogeneous-agent dynamic stochastic general equilibrium model. Within the model, one can calculate the true value change — not just the first-order approximation — in a setting with borrowing constraints, large shocks, and large idiosyncratic risk. Appendix E provides further details of these exercises.

9.1 Validation within a Two-Asset HANK Model

The model is identical to that in Auclert et al. (2021) in both setting and calibration, and is described in detail in Appendix E. In the model, households are subject to idiosyncratic earnings risk e_i, which follows an AR(1) process. They consume a single consumption good c and supply labor N as dictated by a labor union. Preferences are separable between consumption and labor supply, and feature constant relative risk-aversion (CRRA) over consumption. Households have access to a liquid bond b and illiquid account a, and are subject to borrowing and short-selling constraints as in Kaplan et al. (2018). The household's recursive problem is

$$V_t(e_t^i, b_{it-1}, a_{it-1}) = \max_{c_{it}, b_{it}, a_{it}} \left\{ u(c_{it}, N_t) + \beta \mathbb{E}_t \left[V_{t+1}(e_{it+1}, b_{it}, a_{it}) | e_t^i \right] \right\} \tag{12}$$

$$\text{s.t. } c_{it} + b_{it} + a_{it} = (1 - \tau_t) W_t N_t e_t^i + (1 + r_t^a) a_{it-1} + (1 + r_t^b) b_{it-1} - \chi(a_{it}, a_{it-1})$$

$$a_{it} \geq 0, \quad b_{it} \geq \underline{b}.$$

Here, τ_t is a linear income tax, w_t is the real wage, r_t^a and r_t^b are returns on the illiquid and liquid accounts, respectively, and $\chi(a_{it}, a_{it-1})$ is an adjustment cost of the illiquid account. The model's production side is largely standard and features sticky prices and wages so that monetary policy has real effects.

The model provides a useful laboratory to validate our feasible set approach to computing welfare effects, because one can calculate true welfare changes exactly within the model. To calculate these changes, we use the following algorithm. We first calculate the steady-state value function V^{SS} and policy functions $c^*(\cdot), a^*(\cdot), b^*(\cdot)$ following the endogenous grid-point method of Carroll (2006). We then calculate the impulse response to a shock using the Sequence-Space Jacobian (SSJ) methodology of Auclert et al. (2021). We assume that the economy returns to steady state T periods after the shock. We then solve backwards for the value function in period t given the value function in period $t+1$ using the household problem (12), and plugging in for the path of wages W_t, union-imposed labor supply N_t, and rental rates r_t^a, r_t^b. We repeat this process to form an estimate of $V_0(\cdot)$ at every point on the state space. The change in value from the shock is then given by $d\tilde{V}^{FULL} \equiv V_0 - V^{SS}$.

To calculate welfare changes using our feasible set approach, we first simulate a large number of households from the steady state of this model, and record the path of the consumption, asset choices, and idiosyncratic labor supply. We then compute money-metric welfare gains

FIGURE 10: Comparison of Value Change in a Two-Asset HANK Model

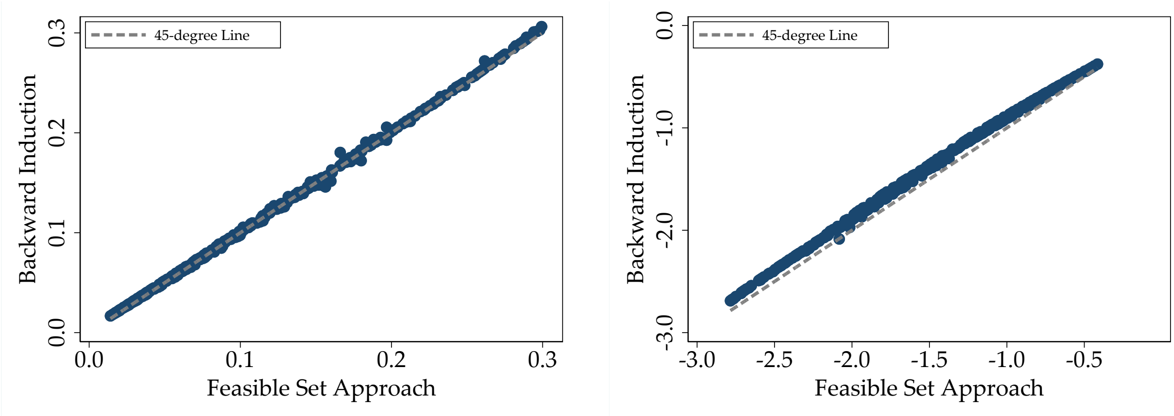

PANEL A: MONETARY SHOCK

PANEL B: TFP SHOCK

Notes: Figure reports the estimated welfare effects computed using our feasible set approach (horizontal axis) and the true welfare change calculated using backwards induction (vertical axis) within a two-asset HANK model (see Appendix E). Panel A considers a 25 basis point reduction in the policy rate, while Panel B considers a 1% decline in TFP. We assume log utility, idiosyncratic risk is calibrated to match a cross-sectional standard deviation of log earnings of 0.5, and households are not allowed to borrow in either the liquid or illiquid asset. Each dot represents the estimated value change for a point on the state space. Axis scales are percentages of steady-state consumption.

dV^{FS} arising from the impulse responses using the formula of Proposition 1, just as we do in the data. We do this following our baseline methodology, which ignores idiosyncratic risk and the constraint effect, in order to assess the performance of the approach requiring the fewest assumptions on preferences.[44]

We perform this routine for both monetary shocks and shocks to aggregate total factor productivity (TFP) (which is the closest thing to an oil shock in this class of models). For our baseline exercises, we consider a 25 basis point cut in interest rates to mirror the size of the monetary shock in our reduced-form IRFs, and a reduction in TFP of 1%. We assume both shocks decay according to an AR(1) process with persistence 0.4 for monetary shocks and 0.9 for TFP shocks. In our baseline scenario, we assume log utility over consumption, an idiosyncratic shock process calibrated to have a cross-sectional standard deviation of log earnings of 0.5 and persistence of 0.966 following Auclert et al. (2021), and a tight borrowing constraint of the form $b_{it} \geq 0$.

Figure 10 compares the money-metric welfare gains from an inflationary monetary shock (Panel A) or TFP shock (Panel B) as calculated from the full model dV^{FULL} (plotted on the vertical axis) and our feasible set approach dV^{FS} (plotted on the horizontal axis). Each dot is a different point in the discretized state space of earnings, liquid assets, and illiquid assets and plots money-metric welfare changes, scaled by households' steady-state consumption (i.e., $d\tilde{V}_{ia}/(u'(c_{ia}^{SS}) * c_{ia}^{SS})$). The figure shows that our feasible set approach generates welfare effects very similar to the full model: all dots lie on or near the 45 degree line.

[44]Because the household does not choose its own labor supply, movements in N appear as taste shocks in this setting. We remove these taste shocks in our baseline analyses, because our feasible set approach returns the value change net of taste shocks.

Table 2 presents several statistics comparing the feasible set approach with the backward induction method for the true welfare change. Column (1) shows our baseline exercise. For monetary shocks, the root-mean-squared error (RMSE) between our approach and the full model is 0.04% of quarterly consumption, and the error is never larger than 0.23%. For the TFP shock, the RMSE rises to 1.76% with maximum absolute error of 2.38%, which is small relative to an average value change of 22.6% of quarterly consumption. That is, the feasible set approach calculates the correct change in value in a model with idiosyncratic risk and tight borrowing constraints.

Furthermore, the feasible set approach successfully characterizes the distribution of welfare shocks in this model. To make this point, we regress the true value change on the value change from the feasible set approach. Both the slope and R^2 from this regression are consistently close to 1. These results suggest that the feasible set approach is well-suited to studying welfare effects of shocks in normal times.

Discrepancies between the feasible set approach and backward-induction approach can arise for three reasons. First, the feasible set approach is valid to a first order, but the full model contains higher-order effects, such as behavioral responses to the shock. Second, the feasible set approach averages over observed histories of shocks; thus, idiosyncratic risk may be imperfectly captured in finite samples. Finally, the feasible set approach is not well suited to jointly studying borrowing constraints and idiosyncratic risk. We study the importance of these considerations in columns (2) through (5) of Table 2.

Column (2) supposes households' preferences for consumption are summarized by a constant relative risk aversion utility function, with a coefficient of relative risk aversion of 2. Unsurprisingly, the feasible set approach continues to perform well under this different specification for utility, reflecting the robustness of the approach to utility function parameterization.[45]

Column (3) doubles the standard deviation of idiosyncratic risk in the model so that the cross-sectional standard deviation of log earnings is 1. The feasible set approach again performs well under the monetary shock. The approach does perform slightly worse for the TFP shock, but the root-mean-squared error remains around one-third of a standard deviation of value changes. The poorer performance arises because the TFP shock generates a larger movement in wages, which scales the uninsurable risk that households face (see Appendix Figure E1).

Column (4) loosens the borrowing constraint so that liquid assets may fall as low as -0.15, which is about 25% of steady-state aggregate consumption. In this case, the feasible set approach continues to perform extremely well. In this case, the welfare gains from expansionary monetary shocks are larger, reflecting the gains accrued by borrowers.

Finally, column (5) considers a much larger shock for which higher-order effects may matter. The monetary shock considers a 1 percentage-point reduction in the policy rate, while the TFP shock considers a 5% decline in TFP. Both shocks decay according to an AR(1) with persistence

[45]Note the mean value changes are different because equilibrium IRFs are different under different utility functions.

TABLE 2: Comparing Performance of Feasible Set Approach to Full Model Value Changes

	Baseline (1)	CRRA 2 (2)	High Risk (3)	Loose Borrow (4)	Big Shock (5)
PANEL A: Monetary Shock					
Root-mean-squared error	0.04%	0.04%	0.12%	0.46%	1.97%
Max absolute error	0.23%	0.26%	0.62%	0.77%	3.41%
Mean Value Change: Full	2.0%	1.2%	1.9%	8.8%	35.6%
Mean Value Change: Feasible Set	2.0%	1.2%	2.0%	9.2%	37.5%
S.D. Value Change: Full	1.3%	1.0%	2.3%	4.4%	18.0%
S.D. Value Change: Feasible Set	1.3%	1.0%	2.2%	4.5%	18.2%
Regression Slope	1.004	0.998	1.008	0.989	0.988
R^2 from Feasible Set	0.999	0.998	0.999	0.999	0.999
PANEL B: TFP Shock					
Root-mean-squared error	1.76%	1.13%	4.85%	1.71%	9.05%
Max absolute error	2.38%	3.28%	8.91%	2.50%	11.9%
Mean Value Change: Full	-22.6%	-11.4%	-18.6%	-22.2%	-112.9%
Mean Value Change: Feasible Set	-24.3%	-12.4%	-22.9%	-23.9%	-121.7%
S.D. Value Change: Full	9.7%	5.7%	14.5%	9.6%	48.8%
S.D. Value Change: Feasible Set	9.9%	6.1%	15.6%	9.7%	49.5%
Regression Slope	0.982	0.938	0.923	0.984	0.984
R^2 from Feasible Set	0.999	0.995	0.986	0.999	0.998

Notes: Table reports the error in estimated welfare effects computed using our feasible set approach when compared with an exact welfare change calculated by backwards induction within a two-asset HANK model (see Appendix E). Column (1) reports our baseline exercise. Column (2) assumes a CRRA utility function with a risk coefficient of 2. Column (3) assumes twice as much idiosyncratic risk as the baseline model, so that the cross-sectional standard deviation of log earnings is 1. Column (4) allows households to borrow in the liquid asset up to 25% of aggregate consumption. Column (5) considers a 1 percentage point reduction in the policy rate (Panel A) or a 5% reduction in TFP (Panel B), both with AR(1) persistence 0.9. The table reports errors and value changes in units of steady-state quarterly consumption.

0.9. In this case, the root mean squred error of the feasible set approach is larger: it grows to almost 2% of consumption for the monetary shock and 9% for the TFP shock. However, these errors are small compared with the average value change from such a shock: the large monetary shock generates an average welfare gain of 35.6% of quarterly consumption, whereas the large TFP shock generates welfare losses of around 113% of quarterly consumption. Furthermore, the approach continues to identify which households benefit the most from the shock: the regression slope and R^2 values remain close to 1. Still, this highlights that our feasible set approach is best used to study relatively small shocks.

9.2 Does the Model Reproduce the Welfare Changes in the Data?

In this section, we compare the distributional consequences of oil supply and monetary shocks implied by the two-asset HANK model studied above with those found in the data in Section 7. To do so, we feed through the model a path for interest rates and TFP that is consistent with our estimated monetary and oil supply shocks. Specifically, we feed in a sequence for the policy rate given by Panel A of Figure 1 to simulate our monetary shock, and a path of TFP

TABLE 3: Money-Metric Value Changes, Model versus Data

	Oil Shock		Monetary Shock	
	Model (1)	Data (2)	Model (3)	Data (4)
RMSE of feasible set approach	0.24%	–	0.04%	–
Mean Value Change	-2.73%	-0.17%	0.52%	-0.36%
SD Total Value Change	0.74%	0.23%	0.22%	0.44%
Mean Labor Income Effect	-2.73%	-0.16%	0.53%	0.43%
Mean Portfolio Effect	0.002%	0.09%	-0.002%	-0.29%
Mean Consumption + Transfer Channel	0%	-0.10%	0%	-0.58%
Regressivity Coefficient γ	0.542	0.492	-0.027	-0.983
	(0.034)	(0.075)	(0.013)	(0.084)
Regressivity of labor-income channel	0.540	-0.036	-0.026	0.217
	(0.034)	(0.034)	(0.013)	(0.076)
Regressivity of Portfolio Channel	0.002	0.526	-0.002	-1.067
	(0.0002)	(0.039)	(0.0001)	(0.104)

Notes: Table compares the estimated welfare effects of our estimated oil price and monetary shocks in both the data (columns 2 and 4) and our baseline two-asset HANK model (columns 1 and 2). The model is detailed in Appendix E, see text for details on how to construct the model shocks. The regressivity coefficient γ is the coefficient from a regression of value change as a percent of four-year consumption on log initial consumption, with heteroskedasticity robust standard errors reported in parentheses.

given by the implied impulse response of aggregate TFP to the Känzig (2021) oil shocks.[46] We then compare the model-implied value changes with the value changes we find in our data. Throughout, we calculate model-implied changes using our feasible set approach and weight according to the stationary distribution.

Table 3 reports statistics of the model-implied (columns (1) and (3)) and data-implied (columns (2) and (4)) value changes from the estimated monetary shock (first two columns) and oil shock (second two columns). The workhorse two-asset HANK model finds much larger value changes on average as a result of these shocks than the data suggests.

Furthermore, the model suggests different regressivity to what we uncover in the data. One summary measure of the regressivity of a shock is the extent to which high-consumption households benefit more from it. To measure this, we estimate regressions of the form

$$dV_i/(u'(c_i) \times c_i) = \alpha + \gamma \ln c_i + \epsilon_i. \tag{13}$$

The dependent variable is the money-metric change in utility from a shock, scaled by consumption. The independent variable is a household type's log consumption. The coefficient γ reports the semi-elasticity of welfare changes to consumption levels. We run this regression in both the model and the data, where c_i is the consumption of a household i either in the model's steady state or in the data in 2019.

[46]Our VAR estimates give paths for interest rates and TFP out to four years. After four year, we assume these paths decay to 0 following an AR(1) process with persistence 0.4 (for monetary) and 0.7 (for TPF).

The model estimates a somewhat similar regressivity coefficient γ (0.542) to what is in the data for oil shocks (0.492). However, the standard deviation of value changes in the model is approximately three times that in the data. In this context, the data suggests that oil shocks are about thrice as regressive as the model suggests, whereas monetary shocks in the data are an order of magnitude more progressive than the model suggests.

Why does this model generate welfare effects of these shocks that are different from the data? The primary reason is that it does not generate realistic household portfolios. Because this standard model lacks both a realistic portfolio choice problem and a life cycle – which give households a reason to accumulate savings at some points in their lives and borrow or spend down savings at other times – the model-implied portfolio effects are orders of magnitude smaller than what we find in the data. Rather, the model attributes nearly all of the welfare effects to the labor-income channel. What's more, the labor-income channel accounts for nearly all of the regressivity implied by the model, which is in stark contrast to the prominent role for the portfolio channel that we find in the data.

Overall, our results suggest three lessons for structural analysis. First, although idiosyncratic risk and borrowing constraints are undoubtedly important for steady-state welfare and the impulse response of the economy to structural shocks, they may not be quantitatively important for the welfare effects *conditional* on impulse responses and policy functions. Note that this may not be true for extremely persistent shocks or shocks that take a long time to manifest, because changes in discounting may matter more for such shocks. Second, standard calibrations of the current workhorse business cycle model – a two-asset HANK model – do not capture the welfare effects of monetary and oil supply shocks inferred from empirical impulse responses. Third, models seeking to better match the welfare effects should focus most on incorporating life cycles and/or realistic portfolio choice problems to give scope for important portfolio effects as in the data (Kekre and Lenel, 2022; Bardóczy and Velásquez-Giraldo, 2024).

10. Conclusion

This paper estimates the incidence of inflationary macroeconomic shocks by proposing a new methodology accounting for movements of all aspects of the budget constraint. In our framework, the money-metric welfare effect of a shock on a given household may be estimated by aggregating empirical IRFs 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.

Our framework is a price theoretic approach to valuing 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 US 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, to study the impact of oil and monetary shocks in different countries and contexts, or to use high-quality administrative data to study distributional effects among more granular household definitions. Our results also urge structural models to introduce life cycles and portfolio choice to better capture the welfare effects of macroeconomic shocks. Finally, our framework does not easily incorporate uncertainty shocks, large aggregate risk, or preference shocks; doing so would be fruitful ground for future work.

REFERENCES

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

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.

ANDERSEN, A. L., K. HUBER, N. JOHANNESEN, L. STRAUB, AND E. T. VESTERGAARD (2022): "Disaggregated Economic Accounts," Working Paper 30630, National Bureau of Economic Research.

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.

AUCLERT, A., B. BARDÓCZY, M. ROGNLIE, AND L. STRAUB (2021): "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 89.

AUCLERT, A., M. ROGNLIE, AND L. STRAUB (Forthcoming): "The Intertemporal Keynesian Cross," *Journal of Political Economy*.

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.

——— (2024): "Sufficient Statistics for Measuring Forward-Looking Welfare," *Quarterly Journal of Economics*.

BAQAEE, D. R. AND A. BURSTEIN (2022): "Aggregate Welfare and Output with Heterogeneous Agents," Working paper.

——— (2023): "Welfare and Output with Income Effects and Taste Shocks," *The Quarterly Journal of Economics*, 138, 769–834.

BARDÓCZY, B. AND M. VELÁSQUEZ-GIRALDO (2024): "HANK Comes of Age," Tech. rep., Finance and Economics Discussion Series 2024-052.

BARTSCHER, A. K., M. KUHN, M. SCHULARICK, AND P. WACHTEL (2021): "Monetary Policy and Racial Inequality," *FRB of New York Staff Report*.

BASU, S., J. G. FERNALD, AND M. S. KIMBALL (2006): "Are Technology Improvements Contractionary?" *American Economic Review*, 96, 1418–1448.

BAUMEISTER, C. AND J. D. HAMILTON (2019): "Structural interpretation of vector autoregressions with incomplete identifcation: Revisiting the role of oil supply and demand shocks," *American Economic Review*, 109, 1873–1910.

BERAJA, M. AND N. ZORZI (2024): "Durables and Size-Dependence in the Marginal Propensity to Spend," *NBER Working Paper Series*.

BOPPART, T., P. KRUSELL, AND K. MITMAN (2018): "Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative," *Journal of Economic Dynamics and Control*, 89, 68–92.

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.

CARROLL, C. D. (2006): "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," *Economics Letters*, 91, 312–320.

CESARINI, D., E. LINDQVIST, M. J. NOTOWIDIGDO, AND R. ÖSTLING (2017): "The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries," *American Economic Review*, 107, 3917–46.

CHANG, M. AND F. SCHORFHEIDE (2024): "On the Effects of Monetary Policy Shocks on Income and Consumption Heterogeneity," Working Paper 32166, National Bureau of Economic Research.

CHETTY, R., A. GUREN, D. MANOLI, AND A. WEBER (2013): "Does Indivisible Labor Explain the Difference Between Micro and Macro Elasticities? A Meta-analysis of Extensive Margin Elasticities," *NBER macroeconomics Annual*, 27, 1–56.

CHODOROW-REICH, G., P. T. NENOV, AND A. SIMSEK (2021): "Stock Market Wealth and the Real Economy: A Local Labor Market Approach," *American Economic Review*, 111, 1613–57.

CLAYTON, C., X. JARAVEL, AND A. SCHAAB (2018): "Heterogeneous Price Rigidities and Monetary Policy," *SSRN Electronic Journal*.

COGLIANESE, J., M. OLSSON, AND C. PATTERSON (2022): "Monetary Policy and the Labor Market: A Quasi-Experiment in Sweden," *Working Paper*.

COIBION, O., Y. GORODNICHENKO, L. KUENG, AND J. SILVIA (2017): "Innocent Bystanders? Monetary policy and inequality," *Journal of Monetary Economics*, 88, 70–89.

CRAVINO, J., T. LAN, AND A. A. LEVCHENKO (2020): "Price stickiness along the income distribution and the effects of monetary policy," *Journal of Monetary Economics*, 110, 19–32.

CRAVINO, J. AND N. A. A. LEVCHENKO (2017): "The Distributional Consequences of Large Devaluations," *American Economic Review*, 107.

DÁVILA, E. AND A. KORINEK (2018): "Pecuniary Externalities in Economies with Financial Frictions," *The Review of Economic Studies*, 85, 352–395.

DÁVILA, 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.

DEATON, A. (1989): "Rice Prices and Income Distribution in Thailand: a Non-parametric Analysis," *The Economic Journal*, 99, 1–37.

DEATON, A. AND S. ZAIDI (2002): *Guidelines for Constructing Consumption Aggregates for Welfare Analysis*, vol. 135, World Bank Publications.

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.

FABER, B. AND T. FALLY (2022): "Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data," *Review of Economic Studies*, 89, 1420–1459.

FAGERENG, A., M. GOMEZ, E. GOUIN-BONENFANT, M. HOLM, B. MOLL, AND G. NATVIK (2022): "Asset-Price Redistribution," Tech. rep., Working paper, LSE.

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.

FERREIRA, C., J. M. LEIVA, G. NUÑO, Á. ORTIZ, T. RODRIGO, AND S. VAZQUEZ (2024): "The Heterogeneous Impact of Inflation on Households' Balance Sheets," Working paper.

GABAIX, X., R. S. J. KOIJEN, F. MAINARDI, S. OH, AND M. YOGO (2023): "Asset Demand of U.S. Households," Working Paper 32001, National Bureau of Economic Research.

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.

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.

HOBIJN, B. AND D. LAGAKOS (2005): "Inflation inequality in the United States," *Review of Income and Wealth*, 51, 581–606.

HOTTMAN, C. J., S. J. REDDING, AND D. E. WEINSTEIN (2016): "Quantifying the Sources of Firm Heterogeneity," *The Quarterly Journal of Economics*, 131, 1291–1364.

HUBERT, P. AND F. SAVIGNAC (2023): "Monetary Policy and Labor Income Inequality: the Role of Extensive and Intensive Margins," Tech. rep., CEPR Discussion Paper No. 18130. CEPR Press, Paris & London.

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., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

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.

KAPLAN, G. AND G. L. VIOLANTE (2024): "The Marginal Propensity to Consume in Heterogeneous Agent Models," *Annual Review of Economics*.

KEKRE, R. AND M. LENEL (2022): "Monetary Policy, Redistribution, and Risk Premia," *Econometrica*, 90, 2249–2282.

KRUEGER, D., E. MALKOV, AND F. PERRI (2023): "How Do Households Respond to Income Shocks?" Tech. rep., National Bureau of Economic Research.

KUHN, F., M. KEHRIG, AND N. L. ZIEBARTH (2021): "Welfare effects of gas price fluctuations," .

LANTERI, A. AND A. A. RAMPINI (2023): "Constrained-Efficient Capital Reallocation," *American Economic Review*, 113, 354–395.

LAUPER, C. AND G. MANGIANTE (2023): "Monetary Policy Shocks and Inflation Inequality," *SSRN Electronic Journal*.

LEE, M., C. MACALUSO, AND F. SCHWARTZMAN (2022): "Minority Unemployment, Inflation, and Monetary Policy," Working Paper.

LI, D., M. PLAGBORG-MØLLER, AND C. K. WOLF (2022): "Local Projections vs. VARs: Lessons From Thousands of DGPs," Working Paper 30207, National Bureau of Economic Research.

MCKAY, A. AND C. WOLF (2023): "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?" .

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.

PALLOTTI, F., G. PAZ-PARDO, J. SLACALEK, O. TRISTANI, AND G. L. VIOLANTE (2023): "Who Bears the Costs of Inflation? Euro Area Households and the 2021–2022 Shock," Working Paper 31896, National Bureau of Economic Research.

PLAGBORG-MØLLER, M. AND C. K. WOLF (2021): "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 89, 955–980.

PUGSLEY, B. AND H. RUBINTON (2023): "The Uneven Welfare Costs of the Volcker Disinflation," Tech. rep., Federal Reserve Bank of St Louis.

RUBBO, E. (2023): "What Drives Inflation? Lessons From Disaggregated Price Data," *Working Paper*.

——— (2024): "Monetary Non-Neutrality in the Cross-Section," .

SAEZ, E. AND S. STANTCHEVA (2016): "Generalized Social Marginal Welfare Weights for Optimal Tax Theory," *American Economic Review*, 106, 24–45.

SAMUELSON, P. A. (1974): "Complementarity: An Essay on the 40th Anniversary of the Hicks-Allen Revolution in Demand Theory," *Journal of Economic literature*, 12, 1255–1289.

SCHAAB, A. AND S. Y. TAN (2023): "Monetary and Fiscal Policy According to HANK-IO," .

SIDRAUSKI, M. (1967): "Rational Choice and Patterns of Growth in a Monetary Economy," *American Economic Review*, 57, 534–544.

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 Propositions 1-3

This subsection proves the main propositions of Section 2. It first proves a general result for the first-order welfare effect of structural shocks with general convex constraints and idiosyncratic risk using a small noise approximation in aggregate risk. It then specializes this general result to give Propositions 1, 2, and 3. We suppress dependence on household group g and initial age a for notational expedience.

Preliminaries. Recall that prices, wages, dividends, and transfers follow stochastic processes given by

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

$$W_t = \bar{W}_t \exp\left(v_t^W\right)^\sigma, \quad T_t = \bar{T}_t \exp\left(v_t^T\right)^\sigma, \tag{A1}$$

where

$$v_{kt}^D = \theta_k^D(L)\boldsymbol{\epsilon}_t, \quad v_{kt}^Q = \theta_k^Q(L)\boldsymbol{\epsilon}_t, \quad v_{jt}^p = \theta_j^p(L)\boldsymbol{\epsilon}_t,$$
$$v_t^W = \theta^W(L)\boldsymbol{\epsilon}_t, \quad v_t^T = \theta^T(L)\boldsymbol{\epsilon}_t.$$

Letting s_t denote the history of the world in period t, summarizing the history of realizations of $\boldsymbol{\epsilon}$ and idiosyncratic histories s_t^i up until t, write the household problem in sequence form (omitting ag superscripts for notational clarity):

$$V(\sigma, \boldsymbol{\epsilon}_0) = \max_{\{\{c_{jt}(s_t)\}_j, \{N_{kt}(s_t)\}_k, L_t(s_t)\}_{t=0}^\infty} \sum_{t=0}^\infty \beta_t \delta_t \sum_{s_t} \pi_t(s_t) U(C_t(s_t), \{N_{kt}(s_t)\}_k, L_t(s_t))) \tag{A2}$$

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

$$\sum_j p_{jt}(s_t) c_{jt}(s_t) = \sum_k [N_{kt-1}(s_{t-1})) D_{kt}(s_t) - Q_{kt}(s_t) \Delta N_{kt}(s_t) - \chi_k(\Delta N_{k,t}(s_t))]$$
$$+ W_t(s_t) e_t^i(s_t) L_t(s_t) + T_t(s_t),$$

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

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

and a set of $M_t \in \mathbb{N}^+$ additional constraints in each period t of the form

$$G_t^m(\mathbf{x}(s_t); \mathbf{z}(s_t)) \leq 0.$$

where $\mathbf{x} = \{\{c_{jt}\}_j, \{N_{kt}\}_k, L_t\}_t$ is the set of household choice variables and $\mathbf{z}$ is a set of objects the household takes as given, such as prices, wages, dividends, transfers, or parameters (e.g. parameters governing a borrowing constraint). Two natural special cases are short-selling constraints, in which $G_t^m(\mathbf{x}(s_t); \mathbf{z}(s_t)) = -N_{kt}(s_t)$, and net-worth constraints in which $G_t^m(\mathbf{x}(s_t); \mathbf{z}(s_t)) = -\sum Q_{kt}(s_t) N_{kt}(s_t)$. We assume that the space of feasible choices that satisfy these constraints is convex and that a unique optimum exists. We let $\lambda_t(s_t)$ be the Lagrange multiplier on the period t budget constraint after realization s_t and $\mu_t^m(s_t)$ be the Lagrange multiplier on the m^{th} general constraint in period t. Write the associated Lagrangian as

$$\begin{aligned}\mathcal{L} = \sum_t \beta_t \delta_t \sum_{s_t} \pi_t(s_t) \Bigg(& U(C_t(s_t), \{N_{kt}(s_t)\}_k, L_t(s_t)) - \lambda_t(s_t) \Bigg[\sum_j p_{jt}(s_t) c_{jt}(s_t) \\ & - \sum_k [N_{kt-1}(s_{t-1}) D_{kt}(s_t) - Q_{kt}(s_t) \Delta N_{kt}(s_t) - \chi_k(\Delta N_{k,t}(s_t))] \\ & - W_t(s_t) e_t^i(s_t^i) L_t(s_t) - T_t(s_t) \Bigg] + \sum_{s_t} \pi_t(s_t) \sum_{m=1}^{M_t} \mu_t^m(s_t) G_t^m(\mathbf{x}(s_t), \mathbf{z}(s_t)) \Bigg).\end{aligned}$$

We have written the household's value function such that it depends on the initial realization of the aggregate state $\boldsymbol{\epsilon}_0$. This problem is parameterized by $\sigma \in [0, 1]$, which indexes a perturbation from an economy whose aggregate variables are deterministic.

Lastly, define the risk-adjustment shifters Θ as

$$\Theta_t^W \equiv Cov\left(\frac{\lambda_t(s_t)}{\mathbb{E}_0[\lambda_t(s_t^i)]}, \frac{W_t(s_t) e_t^i(s_t) L_t(s_t)}{\mathbb{E}_0[W_t(s_t) e(s_t) L_t(s_t)]} \right) \tag{A3}$$

$$\Theta_t^{p,j} \equiv Cov\left(\frac{\lambda_t(s_t)}{\mathbb{E}_0[\lambda_t(s_t)]}, \frac{p_{j,t}(s_t)(s_t) c_{jt}(s_t)}{\mathbb{E}_0[p_{j,t}(s_t) c_{jt}]} \right) \qquad \Theta_t^T \equiv Cov\left(\frac{\lambda_t(s_t)}{\mathbb{E}_0[\lambda_t(s_t)]}, \frac{T_t(s_t)}{\mathbb{E}_0[T_t(s_t)]} \right)$$

$$\Theta_t^{D,k} \equiv Cov\left(\frac{\lambda_t(s_t)}{\mathbb{E}_0[\lambda_t(s_t)]}, \frac{D_{kt}(s_t) N_{kt-1}(s_{t-1})}{\mathbb{E}_0[D_{kt}(s_t) N_{kt-1}(s_{t-1})]} \right) \qquad \Theta_t^{Q,k} \equiv Cov\left(\frac{\lambda_t(s_t)}{\mathbb{E}_0[\lambda_t(s_t)]}, \frac{Q_{kt(s_t)} \Delta N_{kt}(s_t)}{\mathbb{E}_0[Q_{kt}(s_t) \Delta N_{kt}(s_t)]} \right),$$

which will be equal to their counterparts in the text under the assumptions maintained in Proposition 2. We now state a preliminary result.

Lemma 1. *In the limit as $\sigma \to 0$, the first-order change in money-metric welfare from an impulse response to an element n of the fundamental shock vector at $t = 0$ approaches*

(A4)

$$dV = \sum_t R_{0\to t}^{-1} \prod_{\tau=0}^{t-1} \overbrace{\mathbb{E}_0[1+\tilde{\mu}_\tau]}^{\text{Discount Wedge}} \Bigg(-\sum_j \overbrace{\mathbb{E}_0[p_{j,t}c_{jt}(s_t^i)](1+\Theta_t^{p,j})\Psi_{n,t}^{p,j}}^{\text{Consumption Price}} + \overbrace{W_t\mathbb{E}_0[e_t^i(s_t^i)L_t(s_t^i)](1+\Theta_t^W)\Psi_{n,t}^W}^{\text{Labor Income}}$$

$$\underbrace{+\sum_k \Big[\mathbb{E}_0[N_{kt-1}(s_{t-1}^i)]D_{kt}(1+\Theta_t^{D,k})\Psi_{n,t}^{D,k} - Q_{kt}\mathbb{E}_0[\Delta N_{kt}(s_t^i)](1+\Theta_t^{Q,k})\Psi_{n,t}^{Q,k}\Big]}_{\text{Portfolio}} + \underbrace{\mathbb{E}_0[T_t(s_t^i)](1+\Theta_t^T)\Psi_{n,t}^T\Bigg)}_{\text{Transfer}}$$

$$\underbrace{+\sum_t \beta_t\delta_t \sum_{m=1}^{M_t}\sum_{z\in\mathbf{z}} \mathbb{E}_0\left[\mu_t^m(s_t)\frac{\partial G_t^m(\mathbf{x}(s_t),\mathbf{z}(s_t))}{\partial z}\frac{dz(s_t)}{d\sigma}\bigg|_{\epsilon_0^n=1,\epsilon_0^{-n}=0}\right] - \mathbb{E}_0\left[\mu_t^m(s_t)\frac{\partial G_t^m(\mathbf{x}(s_t),\mathbf{z}(s_t))}{\partial z}\frac{dz(s_t)}{d\sigma}\bigg|_{\boldsymbol{\epsilon}_0=0}\right]}_{\text{Constraints}},$$

where choices and risk-adjustment shifters are evaluated at $\sigma = 0$.

Proof of Lemma 1: We approximate $V(\sigma, \boldsymbol{\epsilon}_0)$ around $\sigma = 0$ with the following Taylor approximation:

(A5)
$$V(\sigma,\boldsymbol{\epsilon}_0) \approx V(0) + \frac{dV(\sigma,\boldsymbol{\epsilon}_0)}{d\sigma}\bigg|_{\sigma=0}\sigma.$$

When $\sigma = 0$, $V(0)$ does not depend on $\boldsymbol{\epsilon}$ since the stochastic processes have zero variance. Thus, the $V(0)$ term will cancel when considering the effect of a shock to $\boldsymbol{\epsilon}$. To compute $dV(\cdot)/d\sigma$, we differentiate the Lagrangian with respect to σ. Substituting in the definitions of the stochastic processes for prices, wages, dividends, and transfers, and applying the Envelope Theorem of Oyama and Takenawa (2018) yields

(A6)
$$\begin{aligned}\frac{dV}{d\sigma} = &\sum_t \beta_t\delta_t \sum_{s_t}\pi_t(s_t)\lambda_t(s_t)\Bigg(-\sum_j p_{j,t}(s_t)c_{jt}(s_t)v_{jt}^p + W_t(s_t)e_t^i(s_t)L_t(s_t)v_t^W \\ &+ T_t(s_t)v_t^T(s_t) + \sum_k\Big[N_{kt-1}(s_{t-1}))D_{kt}(s_t)v_{kt}^D(s_t) - Q_{kt}(s_t)\Delta N_{kt}(s_t)v_{kt}^Q(s_t)\Big]\Bigg) \\ &+ \sum_t\beta_t\delta_t\sum_{s_t}\pi_t(s_t)\sum_{m=1}^{M_t}\mu_t^m(s_t)\sum_{z\in\mathbf{z}(s_t)}\frac{\partial G_t^m(\mathbf{x}(s_t),\mathbf{z}(s_t))}{\partial z}\frac{dz(s_t)}{d\sigma}\bigg|_{\boldsymbol{\epsilon}_0},\end{aligned}$$

where choice variables are understood to be evaluated at the optimum. Partition the state vector into aggregate and individual components so that $s_t = \{s_t^A, s_t^i\}$ and $\pi_t(s_t) = \pi_t(s_t^A)\pi_t(s_t^i)$. Multiply and divide each period's $\lambda_t(s_t^i)$ by $\mathbb{E}_0[\lambda_t(s_t^i)|s_t^A]$ to write equation (A6) as

(A7)
$$\begin{aligned}\frac{dV}{d\sigma} = &\sum_t\beta_t\delta_t\sum_{s_t^A}\mathbb{E}_0[\lambda_t(s_t^i)|s_t^A]\pi_t(s_t^A)\sum_{s_t^i}\pi_t(s_t^i)\frac{\lambda_t(s_t^i)}{\mathbb{E}_0[\lambda_t(s_t^i)|s_t^A]}\Bigg(-\sum_j p_{j,t}c_{jt}(s_t^i)v_{jt}^p + W_te_{it}^i(s_t)L_t(s_t^i)v_t^W \\ &+ T_tv_t^T + \sum_k\Big[N_{kt-1}(s_{t-1}^i))D_{kt}v_{kt}^D - Q_{kt}\Delta N_{kt}(s_t^i)v_{kt}^Q\Big]\Bigg) \\ &+\sum_t\beta_t\delta_t\sum_{s_t^A}\sum_{s_t^i}\pi_t(s_t^A)\pi_t(s_t^i)\sum_{m=1}^{M_t}\mu_t^m(s_t)\sum_{z\in\mathbf{z}}\frac{\partial G_t^m(\mathbf{x}(s_t),\mathbf{z}(s_t^A))}{\partial z}\frac{dz(s_t^A)}{d\sigma}.\end{aligned}$$

Next, observe that the s_t first-order condition for the riskless bond $N_{0t}(s_t)$ implies

$$\text{(A8)} \qquad \beta_t \delta_t \lambda_t(s_t) Q_{0t}(s_t) + \beta_t \delta_t \sum_{m=1}^{M_t} \mu_t^m(s_t) \frac{\partial G_t^m(\mathbf{x}(s_t), \mathbf{z}(s_t))}{\partial N_{0t}} = \delta_{t+1}\beta_{t+1}\mathbb{E}_t\left[\lambda_{t+1}(s_{t+1})|s_t\right].$$

Next define

$$\tilde{\mu}_t(s_t) \equiv \frac{1}{Q_{0t}\mathbb{E}_t[\lambda_t(s_t)]}\left(\sum_{m=1}^{M_t} \mu_t^m(s_t)\frac{\partial G_t^m(\mathbf{x}(s_t), \mathbf{z}(s_t))}{\partial N_{0t}}\right),$$

so that equation (A8) can be written

$$\text{(A9)} \qquad \delta_t \beta_t Q_{0t}(s_t)\lambda_t(s_t) + \tilde{\mu}_t(s_t)\mathbb{E}_t[\lambda_t(s_t)]Q_{0t}(s_t)\beta_t\delta_t = \delta_{t+1}\beta_{t+1}\mathbb{E}_t[\lambda_t(s_{t+1})|s_t].$$

As $\sigma \to 0$, this becomes

$$\delta_t \beta_t Q_{0t}\lambda_t(s_t) + \tilde{\mu}_t(s_t)\mathbb{E}_t[\lambda_t(s_t)]Q_{0t}\beta_t\delta_t = \delta_{t+1}\beta_{t+1}\mathbb{E}_t[\lambda_t(s_{t+1})|s_t].$$

Taking expectations of both sides, and iterating this equation back to period 0 and employing the law of iterated expectations yields

$$\text{(A10)} \qquad \delta_t\beta_t\mathbb{E}_t[\lambda_t(s_t)] = \lambda_0 R_{0\to t}^{-1}\prod_{\tau=0}^{t-1}\mathbb{E}_0[1+\tilde{\mu}_\tau].$$

Plugging this into equation (A7) gives

(A11)

$$\begin{aligned}
\frac{dV}{d\sigma} &= \lambda_0 \sum_t R_{0\to t}\prod_{\tau=0}^{t-1}\mathbb{E}_0[1+\tilde{\mu}_\tau]\sum_{s_t^i}\pi_t(s_t^i)\frac{\lambda_t(s_t^i)}{\mathbb{E}_0[\lambda_t(s_t^i)]}\Bigg(-\sum_j p_{j,t}c_{jt}(s_t^i)\mathbb{E}_0[v_{jt}^p] + W_t e_{it}^i(s_t)L_t(s_t^i)\mathbb{E}_0[v_t^W] \\
&\quad + T_t\mathbb{E}_0[v_t^T] + \sum_k\left[N_{kt-1}(s_{t-1}^i))D_{kt}\mathbb{E}_0[v_{kt}^D] - Q_{kt}\Delta N_{kt}(s_t^i)\mathbb{E}_0[v_{kt}^Q]\right]\Bigg) \\
&\quad + \sum_t \beta_t\delta_t\sum_{s_t^i}\pi_t(s_t^i)\sum_{m=1}^{M_t}\mu_t^m(s_t^i)\sum_{z\in\mathbf{z}}\mathbb{E}_0\left[\frac{\partial G_t^m(\mathbf{x}(s_t^i),\mathbf{z})}{\partial z}\frac{dz}{d\sigma}\,|\,s_t^i\right].
\end{aligned}$$

Using the fact that $\mathbb{E}[XY] = Cov(X,Y) + \mathbb{E}[X]\mathbb{E}[Y]$, write the above in covariance form as

$$\begin{aligned}
\text{(A12)}\quad \frac{dV}{d\sigma} &= \lambda_0 \sum_t R_{0\to t} \prod_{\tau=0}^{t-1} \mathbb{E}_0[1+\tilde{\mu}_\tau] \Bigg(-\sum_j \mathbb{E}_0[p_{j,t} c_{jt}(s_t^i)]\mathbb{E}_0[v_{jt}^p] + W_t \mathbb{E}_0[e_t^i(s_t^i) L_t(s_t^i)]\mathbb{E}_0[v_t^W] \\
&+ \sum_k \Big[\mathbb{E}_0[N_{kt-1}(s_{t-1}^i)] D_{kt} \mathbb{E}_0[v_{kt}^D] - Q_{kt}\mathbb{E}_0[\Delta N_{kt}(s_t^i)]\mathbb{E}_0[v_{kt}^Q]\Big] + \mathbb{E}_0[T_t(s_t^i)]\mathbb{E}_0[v_t^T] \\
&- \sum_j Cov\left(\frac{\lambda_t(s_t^i)}{\mathbb{E}_0[\lambda_t(s_t^i)]}, \frac{p_{j,t} c_{jt}(s_t^i)}{\mathbb{E}_0[p_{j,t} c_{jt}]}\right) \mathbb{E}_0[p_{j,t} c_{jt}(s_t^i)]\mathbb{E}_0[v_{jt}^p] \\
&+ Cov\left(\frac{\lambda_t(s_t^i)}{\mathbb{E}_0[\lambda_t(s_t^i)]}, \frac{L_t(s_t^i) e_t^i(s_t^i)}{\mathbb{E}_0[L_t w_t]}\right) W_t \mathbb{E}_0[L_t(s_t^i) e_t^i(s_t^i)]\mathbb{E}_0[v_t^W] \\
&+ \sum_k Cov\left(\frac{\lambda_t(s_t^i)}{\mathbb{E}_0[\lambda_t(s_t^i)]}, \frac{Q_{kt-1} N_{kt-1}(s_{t-1}^i)}{Q_{kt-1}\mathbb{E}_0[N_{kt-1}]}\right) \mathbb{E}_0[N_{kt-1}(s_{t-1}^i)] D_{kt} \mathbb{E}_0[v_{kt}^D] \\
&- \sum_k Cov\left(\frac{\lambda_t(s_t^i)}{\mathbb{E}_0[\lambda_t(s_t^i)]}, \frac{Q_{kt}\Delta N_{kt}(s_t^i)}{Q_{kt}\mathbb{E}_0[\Delta N_{kt}]}\right) \mathbb{E}_0[\Delta N_{kt}(s_t^i)] Q_{kt} \mathbb{E}_0[v_{kt}^Q] \\
&+ Cov\left(\frac{\lambda_t(s_t^i)}{\mathbb{E}_0[\lambda_t(s_t^i)]}, \frac{T_t(s_t^i)}{\mathbb{E}_0[T_t]}\right) \mathbb{E}_0[T_t(s_t^i)]\mathbb{E}_0[v_t^T] \Bigg) \\
&+ \sum_t \beta_t \delta_t \sum_{m=1}^{M_t} \sum_{z\in\mathbf{z}} \mathbb{E}_0\left[\mu_t^m(s_t^i) \frac{\partial G_t^m(\mathbf{x}(s_t^i), \mathbf{z})}{\partial z} \frac{dz}{d\sigma}\right].
\end{aligned}$$

Observe that we can then write equation (A12) as

(A13)

$$\begin{aligned}
\frac{dV}{d\sigma} &= \lambda_0 \sum_t R_{0\to t} \prod_{\tau=0}^{t-1} \overbrace{\mathbb{E}_0[1+\tilde{\mu}_\tau]}^{\text{Discount Wedge}} \Bigg(-\sum_j \overbrace{\mathbb{E}_0[p_{j,t} c_{jt}(s_t^i)](1+\Theta_t^{p,j}) v_{jt}^p}^{\text{Consumption Price}} + \overbrace{W_t \mathbb{E}_0[e_t^i(s_t^i) L_t(s_t^i)](1+\Theta_t^W) v_t^W}^{\text{Labor Income}} \\
&+ \underbrace{\sum_k \Big[\mathbb{E}_0[N_{kt-1}(s_{t-1}^i)] D_{kt}(1+\Theta_t^{D,k}) v_{kt}^D - Q_{kt}\mathbb{E}_0[\Delta N_{kt}(s_t^i)](1+\Theta_t^{Q,k}) v_{kt}^Q\Big]}_{\text{Portfolio}} + \underbrace{\mathbb{E}_0[T_t(s_t^i)](1+\Theta_t^T) v_t^T}_{\text{Transfer}} \Bigg) \\
&+ \underbrace{\sum_t \beta_t \delta_t \sum_{m=1}^{M_t} \sum_{z\in\mathbf{z}} \mathbb{E}_0\left[\mu_t^m(s_t^i) \frac{\partial G_t^m(\mathbf{x}(s_t^i), \mathbf{z})}{\partial z} \frac{dz}{d\sigma}\right]}_{\text{Constraints}},
\end{aligned}$$

using the definitions in (A3).

Define the change in welfare $d\tilde{V}$ from an impulse to element n of the structural shock vector, for any value of σ, as

$$\text{(A14)}\qquad d\tilde{V} \equiv V(\sigma, \epsilon_0^n = 1, \epsilon_0^{-n} = 0) - V(\sigma, \epsilon_0^n = 0, \epsilon_0^{-n} = 0).$$

Using (A5), note that the $V(0)$ term drops out of this expression, so that:

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

Evaluating at $\sigma = 1$ and plugging in using (A13) and the definitions of the impulse response functions $\mathbf{\Psi}$ yields equation (A4) as desired:

(A15)

$$\begin{aligned} dV = \sum_t R_{0\to t} \prod_{\tau=0}^{t-1} \overbrace{\mathbb{E}_0[1+\tilde{\mu}_\tau]}^{\text{Discount Wedge}} \Bigg(& - \sum_j \overbrace{\mathbb{E}_0[p_{j,t} c_{jt}(s_t^i)](1+\Theta_t^{p,j})\Psi_{n,t}^{p,j}}^{\text{Consumption Price}} + \overbrace{W_t \mathbb{E}_0[e_t^i(s_t^i) L_t(s_t^i)](1+\Theta_t^W)\Psi_{n,t}^W}^{\text{Labor Income}} \\ & + \underbrace{\sum_k \left[\mathbb{E}_0[N_{kt-1}(s_{t-1}^i)] D_{kt}(1+\Theta_t^{D,k})\Psi_{n,t}^{D,k} - Q_{kt}\mathbb{E}_0[\Delta N_{kt}(s_t^i)](1+\Theta_t^{Q,k})\Psi_{n,t}^{Q,k} \right]}_{\text{Portfolio}} + \underbrace{\mathbb{E}_0[T_t(s_t^i)](1+\Theta_t^T)\Psi_{n,t}^T}_{\text{Transfer}} \Bigg) \\ & + \underbrace{\sum_t \beta_t \delta_t \sum_{m=1}^{M_t} \sum_{z\in\mathbf{z}} \mathbb{E}_0 \left[\mu_t^m(s_t^i) \frac{\partial G_t^m(\mathbf{x}(s_t^i), \mathbf{z}(s_t))}{\partial z} \left.\frac{dz(s_t)}{d\sigma}\right|_{\epsilon_0^n=1, \epsilon_0^{-n}=0} \right] - \mathbb{E}_0 \left[\mu_t^m(s_t^i) \frac{\partial G_t^m(\mathbf{x}(s_t^i), \mathbf{z}(s_t))}{\partial z} \left.\frac{dz(s_t)}{d\sigma}\right|_{\epsilon_0=0} \right]}_{\text{Constraints}}, \end{aligned}$$

where recall that $dV = d\tilde{V}/\lambda_0$.

Proof of Proposition 1. As the variance of idiosyncratic risk tends to zero, so too do the risk-adjustment shifters. Therefore, all $\Theta_t^x = 0$. Without constraints besides the budget constraint, all $\mu_t^m = 0$, so that $\tilde{\mu}_t = 0$, and there is neither a discount wedge nor constraint effect. It is then easy to see that (A15) becomes the expression in Proposition 1.

Proof of Proposition 2. When there are no constraints other than the budget constraint, households remain on their expected Euler equation. Therefore, $\tilde{\mu}_t(s_t) = 0$, and there is no discount wedge. In addition, the constraint effects are zero. Thus, we need only show that the adjustment factors Θ_t^x defined above are the same as those defined in the main text. To do so, note that the first-order condition for consumption good j in period t in the general problem is

$$\lambda_t(s_t^i) p_{jt}(s_t) - \sum_{m=1}^{M_t} \mu_t^m(s_t) \frac{\partial G_t^m(\mathbf{x}(s_t), \mathbf{z}(s_t))}{\partial c_{jt}(s_t)} = U_C(C_t(s_t), \{N_{kt}(s_t)\}_k, L_t(s_t)) \frac{\partial \mathcal{C}(C_t(s_t), \{c_{jt}(s_t)\}_j)}{\partial c_{jt}(s_t)}.$$

Given the household's expenditure minimization problem, one can show that

$$p_{jt}(s_t) = P_t(s_t) \frac{\partial \mathcal{C}(C_t(s_t), \{c_{jt}(s_t)\}_j)}{\partial c_{jt}(s_t)}.$$

Since there are no constraints besides the budget constraints, the first-order condition can therefore be written

$$\lambda_t(s_t^i) = \frac{U_C(C_t(s_t), \{N_{kt}(s_t)\}_k, L_t(s_t))}{P_t(s_t)}. \tag{A16}$$

Substituting this into the definition for Θ_t^x (A3) yields the result.

Proof of Proposition 3. Here we specialize the set of constraints as $G_t(\mathbf{x}, \mathbf{z}) = -\sum_k Q_{kt} N_{kt} \leq 0$. In this case,

$$\frac{\partial G_t(\mathbf{x}, \mathbf{z})}{\partial N_{0t}} = -Q_{0t}, \qquad \frac{\partial G_t(\mathbf{x}, \mathbf{z})}{\partial z} = -N_{0t}, \qquad \frac{dz}{d\sigma} = Q_{kt} v^Q_{kt}.$$

Therefore, the discount wedge becomes

$$\tilde{\mu}_t(s_t) = -\frac{\mu_t(s_t)}{\mathbb{E}_0[\lambda_t(s_t)]} = -\frac{\mu_t}{\lambda_t},$$

where the second equality follows as idiosyncratic risk becomes small. The constraint effect becomes

$$\beta_t \delta_t \mu_t Q_{kt} N_{kt} \Psi^{Q,k}_{n,t}.$$

Finally, plugging (A16) into a the first-order condition for risk-free bonds (A8) when idiosyncratic and aggregate risk are both small yields

$$\beta_t \delta_t (1 + \tilde{\mu}_t) \frac{U_C(C_t, \{N_{kt}\}_k, L_t)}{P_t} = R_t \beta_{t+1} \delta_{t+1} \frac{U_C(C_{t+1}, \{N_{kt+1}\}_k, L_{t+1})}{P_{t+1}}. \tag{A17}$$

Defining $1 + \tau_t \equiv (1 + \tilde{\mu}_t)^{-1}$ yields equation (9) in the main text. Finally, note that $\mu_t = \lambda_t \tilde{\mu}_t$, and perform the same steps as above to substitute out λ_t for λ_0. This yields the formula for Proposition 3 with no idiosyncratic risk and borrowing constraints.

A.2 Proof of Proposition 4

We now consider a second-order expansion of the value function as

$$V_2(\sigma) = V(0) + \frac{dV(0)}{d\sigma}\sigma + \frac{1}{2}\frac{d^2 V(0)}{d\sigma^2}\sigma^2.$$

Then, to calculate the change in welfare from an impulse to the fundamental shock ϵ_1, we difference the welfare expansions between a world that receives an impulse to the fundamental shock vector at time 0, and one that does not, giving

$$\begin{aligned} dV_2(\sigma) &= V(\sigma, \epsilon_1 = 1) - V(\sigma, \epsilon_1 = 0) \\ &= \left(\frac{dV(0, \epsilon_1 = 1)}{d\sigma} - \frac{dV(0, \epsilon_1 = 0)}{d\sigma} \right)\sigma + \frac{1}{2}\left(\frac{d^2 V(0, \epsilon_1 = 1)}{d\sigma^2} - \frac{d^2 V(0, \epsilon_1 = 0)}{d\sigma^2} \right)\sigma^2. \end{aligned} \tag{A18}$$

We derive the following in an environment of no idiosyncratic risk (i.e., $Var(e^i_t) = 0$), but applying the methodology developed above to the case with general idiosyncratic risk is not difficult (though it is tedious and notationally dense). Now, as before, we have

$$\frac{dV(\sigma)}{d\sigma} = \sum_t \beta_t \delta_t \sum_{s_t} \pi_t(s_t) \Bigg(\lambda_t(s_t) \Bigg[-\sum_j p_{j,t}(s_t) c_{jt}(s_t) v_{jt}^p(s_t) + W_t(s_t) L_t(s_t) v_t^W(s_t) + T_t(s_t) v_t^T(s_t)$$
$$+ \sum_k \Big[N_{kt-1}(s_{t-1}) D_{kt}(s_t) v_{kt}^D(s_t) - Q_{kt}(s_t) \Delta N_{kt}(s_t) v_{kt}^Q(s_t) \Big] \Bigg] \Bigg).$$

Differentiating this again yields

$$\frac{d^2V(\sigma)}{d\sigma^2} = \sum_t \beta_t \delta_t \sum_{s_t} \pi_t(s_t) \frac{d\lambda_t(s_t)}{d\sigma} \Bigg(-\sum_j p_{j,t}(s_t) c_{jt}(s_t) v_{jt}^p(s_t) + W_t(s_t) L_t(s_t) v_t^W(s_t) + T_t(s_t) v_t^T(s_t)$$
$$+ \sum_k \Big[N_{kt-1}(s_{t-1}) D_{kt}(s_t) v_{kt}^D(s_t) - Q_{kt}(s_t) \Delta N_{kt}(s_t) v_{kt}^Q(s_t) \Big] \Bigg)$$
$$+ \sum_t \beta_t \delta_t \sum_{s_t} \pi_t(s_t) \lambda_t(s_t) \Bigg(-\sum_j p_{j,t}(s_t) c_{jt}(s_t) v_{jt}^p(s_t)^2 + W_t(s_t) L_t(s_t) v_t^W(s_t)^2 + T_t(s_t) v_t^T(s_t)^2$$
$$+ \sum_k \Big[N_{kt-1}(s_{t-1}) D_{kt}(s_t) v_{kt}^D(s_t)^2 - Q_{kt}(s_t) \Delta N_{kt}(s_t) v_{kt}^Q(s_t)^2 \Big] \Bigg)$$
$$+ \sum_t \beta_t \delta_t \sum_{s_t} \pi_t(s_t) \lambda_t(s_t) \Bigg(-\sum_j \frac{dp_{j,t}(s_t) c_{jt}(s_t)}{d\sigma} v_{jt}^p(s_t) + \frac{dW_t(s_t) L_t(s_t)}{d\sigma} v_t^W(s_t) + \frac{dT_t(s_t)}{d\sigma} v_t^T(s_t)$$
$$+ \sum_k \Bigg[\frac{N_{kt-1}(s_{t-1}) D_{kt}(s_t)}{d\sigma} v_{kt}^D(s_t) - \frac{dQ_{kt}(s_t) \Delta N_{kt}(s_t)}{d\sigma} v_{kt}^Q(s_t) \Bigg] \Bigg).$$

Using equations (A1) and the chain rule, we can rewrite the last two lines as

$$
\begin{aligned}
\frac{d^2V(\sigma)}{d\sigma^2} &= \sum_t \beta_t \delta_t \sum_{s_t} \pi_t(s_t) \frac{d\lambda_t(s_t)}{d\sigma}\Bigg(-\sum_j p_{j,t}(s_t) c_{jt}(s_t) v^p_{jt}(s_t) + W_t(s_t) L_t(s_t) v^W_t(s_t) + T_t(s_t) v^T_t(s_t) \\
&+ \sum_k \Big[N_{kt-1}(s_{t-1}) D_{kt}(s_t) v^D_{kt}(s_t) - Q_{kt}(s_t) \Delta N_{kt}(s_t) v^Q_{kt}(s_t) \Big] \Bigg) \\
&+ \sum_t \beta_t \delta_t \sum_{s_t} \pi_t(s_t) \lambda_t(s_t) \Bigg(-\sum_j p_{j,t}(s_t) c_{jt}(s_t) v^p_{jt}(s_t)^2 + W_t(s_t) L_t(s_t) v^W_t(s_t)^2 + T_t(s_t) v^T_t(s_t)^2 \\
&+ \sum_k \Big[N_{kt-1}(s_{t-1}) D_{kt}(s_t) v^D_{kt}(s_t)^2 - Q_{kt}(s_t) \Delta N_{kt}(s_t) v^Q_{kt}(s_t)^2 \Big] \Bigg) \\
&+ \sum_t \beta_t \delta_t \sum_{s_t} \pi_t(s_t) \lambda_t(s_t) \Bigg(-\sum_j p_{j,t}(s_t) c_{jt}(s_t) \sum_s \sum_{\omega_s \in \Omega_s} \frac{d \ln c_{jt}}{d \ln \omega_s} v^{\omega_s}_s v^p_{jt}(s_t) \\
&+ W_t(s_t) L_t(s_t) \sum_s \sum_{\omega_s \in \Omega_s} \frac{d \ln W_t(s_t) L_t(s_t)}{d \ln \omega_s} v^{\omega_s}_s v^W_t(s_t) + T_t(s_t) \sum_s \sum_{\omega_s \in \Omega_s} \frac{d \ln T_t(s_t)}{d \ln \omega_s} v^{\omega_s}_s v^T_t(s_t) \\
&+ \sum_k \Bigg[N_{kt-1}(s_{t-1}) D_{kt}(s_t) \sum_s \sum_{\omega_s \in \Omega_s} \frac{d \ln N_{kt-1}(s_{t-1}) D_{kt}(s_t)}{d \ln \omega_s} v^{\omega_s}_s v^D_{kt}(s_t) \\
&- |\Delta N_{kt-1}(s_{t-1}) Q_{kt}(s_t)| \sum_s \sum_{\omega_s \in \Omega_s} \frac{d \Delta N_{kt-1}(s_{t-1}) Q_{kt}(s_t) / |\Delta N_{kt-1}(s_{t-1}) Q_{kt}(s_t)|}{d \ln \omega_s} v^{\omega_s}_s v^Q_{kt}(s_t) \Bigg] \Bigg),
\end{aligned}
$$

for $\omega_s \in \Omega_s \equiv \{\{p_{js}\}_j, \{Q_{ks}\}_k, \{D_{ks}\}_k, W_s, T_s\}$. Taking $\sigma \to 0$ and forming expectations, we get

$$
\begin{aligned}
\frac{d^2V(\sigma)}{d\sigma^2} &= \sum_t \beta_t \delta_t \frac{d\lambda_t}{d\sigma}\Bigg(-\sum_j p_{j,t} c_{jt} \mathbb{E}_0[v^p_{jt}(s_t)] + W_t L_t \mathbb{E}_0[v^W_t] + T_t(s_t) v^T_t \\
&+ \sum_k \Big[N_{kt-1}) D_{kt} \mathbb{E}_0[v^D_{kt}(s_t)] - Q_{kt} \Delta N_{kt} \mathbb{E}_0[v^Q_{kt}] \Big] \Bigg) \\
&+ \sum_t \beta_t \delta_t \lambda_t \Bigg(-\sum_j p_{j,t} c_{jt} \mathbb{E}_0[(v^p_{jt})^2] + W_t L_t \mathbb{E}_0[(v^W_t)^2] + T_t \mathbb{E}_0[(v^T_t)^2] \\
&+ \sum_k \Big[N_{kt-1} D_{kt} \mathbb{E}_0[(v^D_{kt})^2] - Q_{kt} \Delta N_{kt} \mathbb{E}_0[(v^Q_{kt})^2] \Big] \Bigg) \\
&+ \sum_t \beta_t \delta_t \lambda_t \Bigg(-\sum_j p_{j,t} c_{jt} \sum_s \sum_{\omega_s \in \Omega_s} \frac{d \ln c_{jt}}{d \ln \omega_s} \mathbb{E}_0[v^{\omega_s}_s v^p_{jt}] \\
&+ W_t L_t \sum_s \sum_{\omega_s \in \Omega_s} \frac{d \ln W_t L_t}{d \ln \omega_s} \mathbb{E}_0[v^{\omega_s}_s v^W_t] + T_t \sum_s \sum_{\omega_s \in \Omega_s} \frac{d \ln T_t}{d \ln \omega_s} \mathbb{E}_0[v^{\omega_s}_s v^T_t] \\
&+ \sum_k \Bigg[N_{kt-1} D_{kt} \sum_s \sum_{\omega_s \in \Omega_s} \frac{d \ln N_{kt-1} D_{kt}}{d \ln \omega_s} \mathbb{E}_0[v^{\omega_s}_s v^D_{kt}] \\
&- |\Delta N_{kt-1} Q_{kt}| \sum_s \sum_{\omega_s \in \Omega_s} \frac{d \Delta N_{kt-1} Q_{kt} / |\Delta N_{kt-1} Q_{kt}|}{d \ln \omega_s} \mathbb{E}_0[v^{\omega_s}_s v^Q_{kt}] \Bigg] \Bigg).
\end{aligned}
$$

Plugging this expression into (A18) and using the definitions for Ξ_2 and Ξ_3 yields the result.

A.3 Durable Goods

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 $c_{jt} \equiv \varrho_{jt} d_{jt}$, where d_{jt} is the household's *stock* of the durable good at the beginning of period t and $\varrho_{jt} \in [0,1]$ 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} would be the quantity of vehicles owned, whereas ϱ_{jt} 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, where $\delta'(\varrho) \geq 0, \delta''(\varrho) \geq 0$. Under this assumption, one can write the law of motion for the durable as

$$d_{jt+1} = (1 - \delta(\varrho_{jt}))d_{jt} + I_{jt}, \tag{A19}$$

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 $\chi_j(\Delta d_{jt}, \varrho_{jt})$.[47] 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 $j \in \{1, 2 \ldots, \hat{J}\}$ are non-durable, whereas goods $j \in \{\hat{J}+1, \ldots, J\}$ are durable and may be carried across periods. The household's period t budget constraint can be written as

$$\begin{aligned}\underbrace{\sum_{j=1}^{\hat{J}} p_{jt} c_{jt}}_{\substack{\text{Non-Durable}\\\text{Consumption}}} + \underbrace{\sum_{j=\hat{J}+1}^{J} Q_{jt}\delta_j(\varrho_{jt}) d_{jt}}_{\text{Durable Consumption}} &= W_t L_t + T_t + \underbrace{\sum_k \left[N_{k,t-1}D_{k,t} - Q_{k,t}\Delta N_{k,t} - \chi_k(\Delta N_{k,t})\right]}_{\text{Financial Assets}} \\ &\quad - \sum_{j=\hat{J}+1}^{J} \underbrace{\left(Q_{jt}\Delta d_{jt} + \chi_j(\Delta d_{jt}, \varrho_{jt})\right)}_{\substack{\text{Cost of Increasing}\\\text{Durables Stock}}},\end{aligned} \tag{A20}$$

where Q_{jt} is the price of durable good j's purchases. 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 ϱ_{jt}, one can see the price of utility-relevant durable consumption is $p_{jt} \equiv Q_{jt}\delta_j(\varrho_{jt})/\varrho_{jt}$: the dollar value lost to depreciation as a result of usage. On the income side, durables behave as a financial asset with zero dividend. Proposition 1 therefore directly applies to the case with durable goods, as long as one can appropriately measure the depreciation rate of the durables in question.

[47] Our framework requires the value function remain differentiable. We are therefore unable to account for fixed costs of durable adjustment of the sort seen in Beraja and Zorzi (2024).

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 increases in house prices are only positive for welfare. This result shows 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.[48] 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.

We include vehicles and housing as durable goods in our welfare calculation. We assume a constant depreciation rate for housing of 3.64% per year following the IRS. We calculate depreciation for vehicles using data from the National Household Travel Survey (NHTS) and the hedonic price regressions of Dexheimer (2003), as detailed in Appendix B.5.

Given the above treatment of durable goods, when computing their contribution to welfare changes, the IRFs for housing and vehicle prices are used twice. First, changes in vehicle/housing prices impact the costs of consuming the durable good. Consequently, an increase in price induces a welfare loss whenever consumption is positive (i.e., if the household owns a car or a house). This is the same mechanism that works for any other consumption category. The second impact comes from changes in durable holdings. When households increase their durable stock (e.g., by buying a new car), a price increase reduces welfare by making stock adjustments more costly. This mechanism is identical to the one for other portfolio adjustments.

B. Data Appendix

This section describes the data used in our analysis in more detail. It describes the Consumer Expenditure Survey (CEX), Survey of Consumer Finances (SCF), Current Population Survey (CPS), Survey of Income and Program Participation (SIPP), and National Household Travel Survey (NHTS) and outlines our approach to cleaning them. In addition, Tables B1 and B2 present summary statistics from each dataset.

B.1 Consumer Expenditure Survey (CEX)

Data on households' consumption are obtained from the Public Use Micro Data (PUMD) from the Interview section of the CEX. These data are 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

[48] Adjustments to mortgage interest payments are included as a negative dividend (D_{kt}) for mortgages.

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 calculate total quarterly expenditures using the variable recording total expenditure of the household in each quarter. Then, we compute the average expenditure across all quarters in 2019 to remove the seasonality expected in quarterly expenditures.

We focus only on the sample of households whose reference person is between 25 and 80 years old, and top-code age at 75.[49] Our three educational groups – high school (HS) or less, some college, and bachelor's plus – are defined via the education of the survey's reference person.

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.[50]

Finally, we extract mortgage payments from the CEX. Households indicate how much they pay on their mortgage 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 mortgage 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 are 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 website.

For equity holdings, we include directly held stocks and indirect 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 information on the percentage of the financial instrument invested in the corresponding asset class. For example, combination mutual fund holdings are evenly split between bonds and stocks. We further split bonds into corporate and non-corporate bonds. The former are obtained from the variable "Corporate and Foreign bonds,"

[49] We do this to reduce noise among our eldest households. The implicit assumption is that consumption patterns are similar between people aged between 75 and 80.

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

and the latter are all other bonds.[51] 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 obtain average asset holdings for each of our asset classes at the age-education attainment level.

We are interested in accumulation patterns and previous holdings by quarter. We start by linearly interpolating asset holdings between years of age. Then we define the accumulation of each of our asset classes as the difference between the holdings at age a (measured in years) and the holdings at age $a - 1/4$. Previous holdings at age a are defined as the holdings at age $a - 1/4$. To see when this approach is valid, note that we want to measure

$$Q_{kt}(N^a_{kt} - N^a_{kt-1}),$$

We observe $Q_{kt}N^a_{kt}$ and $Q_{kt}N^{a-1}_{kt}$. Therefore, the assumption underpinning our approach is that accumulation profiles in quantities of assets is constant over time so that $N^a_{kt-1} = N^{a-1}_{kt}$. Using a cross-section of asset holdings holds the asset prices fixed. Finally, we run a LOWESS for each asset class over the life cycle, to reduce jumps due to measurement error.

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 has 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 (e.g., depreciation costs) are multiplied by this rate, whereas welfare effects from CPI 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 3.64%, which we convert to quarterly by dividing by 4.

B.3 Current Population Survey (CPS)

Our data on labor income are from the Current Population Survey (CPS) provided by IPUMS. The CPS is a survey jointly sponsored by the US Census Bureau and 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

[51] Some portion of indirectly held bonds may be corporate. However, the SCF does not allow us to know this.

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.

In addition to the basic CPS, households are asked an additional set of questions, known as either the "Outgoing Rotations Group" (ORG) or "Earner Study," in the final month of each rotation. 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. For our approach, we compute average weekly earnings in 2019 for all households of type a, g from the variable EARNWEEK. These averages constitute the wage portion of the budget constraint for our welfare calculations. For the estimation of IRFs, we compute this same average but only at the g level, and then transform it taking logs. We do this for the longest window that the CPS provides. See section B.7 for more details on the treatment of time-series data for 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), which is a monthly household panel survey. 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.[52] 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.[53] After computing this average, we estimate quarterly income from transfers and interpolate transfer income between quarters. Lastly, we smooth these transfer income patterns over the life cycle with a LOWESS smoother.[54] The result is shown in Figure B1.

[52] Explicitly, 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.

[53] As 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.

[54] We use a LOWESS bandwidth of 0.4 for transfer income to better capture the jump at 65.

FIGURE B1: Transfer Income over the Life Cycle

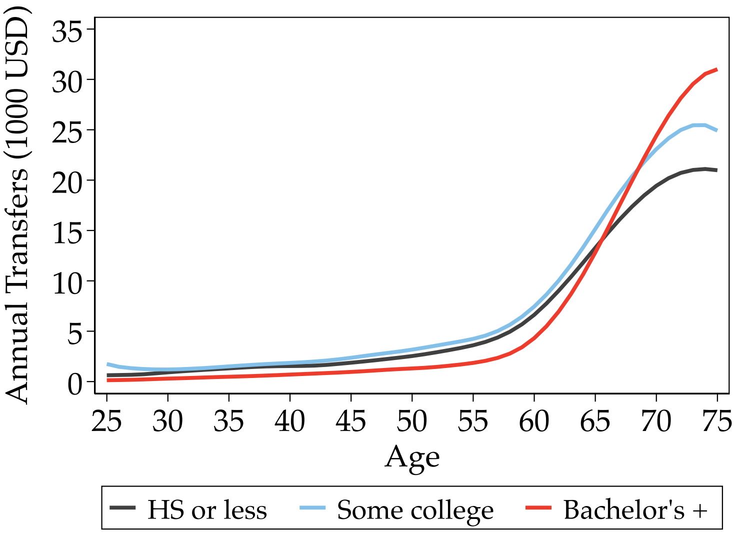

Notes: Figure shows annual income from transfers by group. Data are from the second wave of the 2018 panel of the SIPP. Income is averaged within group and age, and then a LOWESS smoother is applied across age.

B.5 National Household Travel Survey (NHTS)

Data on usage and characteristics of vehicles are 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.[55] 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, top-code age at 75, and construct our educational attainment groups.[56] 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.[57] 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.[58] Finally, we divide the depreciation parameter by 4 to obtain a quarterly estimate.

B.6 Summary Statistics

Table B1 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

[55]More details can be found in the NHTS website.

[56]We 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.

[57]The `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.

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

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 Table B2, 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 ORG of the CPS has 138,270, the SCF has 26,750, and the SIPP has 14,429 observations.

The average annual consumption expenditures in the CEX is $56,138. However, large variance exists 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 $554 for high-school household heads and $1,388 for college-educated households. Likewise, unemployment rates for high school educated households (4.4%) are much higher that of college-educated households (2.8%). 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 $260,000 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).

B.7 Time-Series Data

This section describes the time-series data and processing. When available from the original data source, the seasonally-adjusted series is used. Otherwise, data are seasonally adjusted using the US Census Bureau's X-13ARIMA-SEATS seasonal adjustment program.

Table B3 lists the prices series for which impulse responses are computed. Because impulse responses are estimated using a VAR in log-levels, the CPI price indices are transformed using the 100 $\times$ log transformation. For most categories, the sample coverage is long (available since January 1973).

Table B4 describes the wage data for which impulse responses are computed. For WAGES, we average weekly earnings by demographic group (i.e., education and income) over all households in the CPS. Then, this average is transformed using the 100 $\times$ log transformation. See section B.3 for more details on the treatment of the CPS wage data.

Table B5 displays the asset price variables, mortgage interest payment variables, monetary VAR variables, and oil VAR variables. See below for details:

- ASSETS:
  - Returns (series 42, 43, 45): To match the log-level VAR specification, returns are

TABLE B1: Descriptive Statistics: Cross-Sectional Survey Data in 2019

	HS or Less	Some College	College+	Full Sample
Panel (a): Consumer Expenditure Survey (CEX)				
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
Panel (b): Current Population Survey (CPS)				
Share of sample (%)	32.5	28.4	39.1	100.0
Average age	54.0	51.9	50.1	51.9
Unemployment Rate (%)	4.4	3.7	2.8	3.6
Av. Weekly Earnings	$554	$801	$1388	$950
Observations	45486	39896	52888	138270
Panel (c): Survey of Consumer Finances (SCF)				
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
Share constrained (%)	7.67	12.95	10.82	10.34
Observations	7803	6457	12490	26750
Panel (d): Survey of Income Program Participation (SIPP)				
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. Parentheses below the housing holdings presents 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 ORG sample of the CPS. Only households whose head is at least 25 years old are included.

TABLE B2: Detailed Demographic Statistics: Cross-Sectional Survey Data

	HS or Less	Some College	College+	Full Sample
Panel (a): Consumer Expenditure Survey (CEX)				
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
Panel (b): Current Population Survey (CPS)				
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
Panel (c): Survey of Consumer Finances (SCF)				
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
Panel (d): Survey of Income Program Participation (SIPP)				
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 ORG sample of the CPS. Only households whose head is at least 25 years old are included.

TABLE B3: Description of prices data.

No.	Label	Source	Sample	Trans
1	CPI: Apparel	BLS	1973:01-2019:12	$100 \times \log$
2	CPI: Education	BLS	1994:01-2019:12	$100 \times \log$
3	CPI: Information and information processing	BLS	1994:01-2019:12	$100 \times \log$
4	CPI: Food at Home	BLS	1973:01-2019:12	$100 \times \log$
5	CPI: Alcoholic Beverages	BLS	1973:01-2019:12	$100 \times \log$
6	CPI: Personal care	BLS	1973:01-2019:12	$100 \times \log$
7	CPI: Shelter	BLS	1973:01-2019:12	$100 \times \log$
8	CPI: Fuels and utilities	BLS	1973:01-2019:12	$100 \times \log$
9	CPI: Household furnishings and operations	BLS	1973:01-2019:12	$100 \times \log$
10	CPI: Medical care commodities	BLS	1973:01-2019:12	$100 \times \log$
11	CPI: Medical care services	BLS	1973:01-2019:12	$100 \times \log$
12	CPI: Recreation	BLS	1994:01-2019:12	$100 \times \log$
13	CPI: Postage and delivery services	BLS	1998:12-2019:12	$100 \times \log$
14	CPI: Information technology, hardware and services	BLS	1989:12-2019:12	$100 \times \log$
15	CPI: Food Away from Home	BLS	1973:01-2019:12	$100 \times \log$
16	CPI: Tobacco and smoking products	BLS	1973:01-2019:12	$100 \times \log$
17	CPI: New vehicles	BLS	1973:01-2019:12	$100 \times \log$
18	CPI: Used cars and trucks	BLS	1973:01-2019:12	$100 \times \log$
19	CPI: Leased cars and trucks; Car and truck rental	BLS	1998:12-2019:12	$100 \times \log$
20	CPI: Motor fuel	BLS	1973:01-2019:12	$100 \times \log$
21	CPI: Motor vehicle parts and equipment	BLS	1978:12-2019:12	$100 \times \log$
22	CPI: Motor vehicle maintenance and repair	BLS	1973:01-2019:12	$100 \times \log$
23	CPI: Motor vehicle Insurance	BLS	1973:01-2019:12	$100 \times \log$
24	CPI: Motor vehicle fees	BLS	1998:12-2019:12	$100 \times \log$
25	CPI: Public transportation	BLS	1973:01-2019:12	$100 \times \log$
26	CPI: All items in U.S. city average	BLS	1947:03-2022:11	$100 \times \log$

TABLE B4: Description of Earnings data.

No.	Label	Source	Sample	Trans
27	Weekly earnings (HS or Less)	CPS	1982:01-2022:05	$100 \times \log$
28	Weekly earnings (Some college)	CPS	1982:01-2022:05	$100 \times \log$
29	Weekly earnings (Bachelor's +)	CPS	1982:01-2022:05	$100 \times \log$
30	Weekly earnings (first income quintile)	CPS	1982:01-2022:05	$100 \times \log$
31	Weekly earnings (second income quintile)	CPS	1982:01-2022:05	$100 \times \log$
32	Weekly earnings (third income quintile)	CPS	1982:01-2022:05	$100 \times \log$
33	Weekly earnings (fourth income quintile)	CPS	1982:01-2022:05	$100 \times \log$
34	Weekly earnings (fifth income quintile)	CPS	1982:01-2022:05	$100 \times \log$

cumulated to form an index and logged using the $100 \times \log$ transformation.

  – Log estimated dividend yield (series 44): Let r_t^d give the CRSP value-weighted monthly return including dividends, and let r_t^{ex} give the CRSP value-weighted monthly return excluding dividend. Then, the log estimated dividend yield is computed as $100 \times \log(dp_t) = 100 \times \log\left((r_t^d - r_t^{ex}) \cdot \frac{q_{t-1}}{q_t}\right)$, where q_t gives the ex-dividend return in price space.
  – Estimated dividends (series 49): The estimated dividends series is formed as the product of the estimated dividend yield series and returns ex dividends (in price space) $div_t = dp_t \cdot q_t$. The $100 \times \log$ transformation is then applied to construct series 86.

- The mortgage rate (series 51) comes from the "Mortgage Interest Paid, Owner- and Tenant-Occupied Residential Housing" table from the National Income and Product Accounts. The quarterly series is linearly interpolated after shifting the quarterly value to the final month of the quarter (e.g. a first quarter value is set to March).

- The Basu, Fernald, and Kimball (2006) TFP series (Series 52 and 53) come from the Federal Reserve Bank of San Francisco.[59]

- Data for the Gilchrist and Zakrajšek (2012) Excess Bond Premium series (series 54) come from the Federal Reserve Board.[60]

- Control variables included in the oil VAR (series 55-58) were originally compiled in Baumeister and Hamilton (2019). We use an updated version of this dataset available on Christiane Baumeister's website.[61]

C. Estimation Appendix

C.1 Estimating Impulse Response Functions using SVAR-IV

In this section, we detail the steps for estimating the impulse responses for the oil and monetary applications featured in the main text. For what follows, let $\mathbf{M}_{ij}$ give row i and column j of a matrix $\mathbf{M}$. $\mathbf{M}_{\cdot i}$ gives the i^{th} column.

With the goal of using all available data, the impulse response function (IRF) of an outcome variable of interest x_{nt+h} (e.g., the response of the price of motor fuel to a monetary policy shock after h periods) to some shock ε_t is estimated using the following two-step procedure:

1. Shock estimation: We replicate the SVAR-IV specifications for our shock of interest—like the baseline specifications of Känzig (2021) or Gertler and Karadi (2015). Both papers

[59]The data are downloaded from https://www.frbsf.org/research-and-insights/data-and-indicators/total-factor-productivity-tfp/.

[60]The data are downloaded from https://www.federalreserve.gov/econres/notes/feds-notes/updating-the-recession-risk-and-the-excess-bond-premium-20161006.html.

[61]The data are downloaded from https://sites.google.com/site/cjsbaumeister/research.

TABLE B5: Description of Asset Price, Housing, and Oil Data.

No.	Label	Source	Sample	Trans
35	Federal Funds Effective Rate	FRED	1965:M1-2019:M12	level
36	1-Year Constant Maturity Treasury Yield (Market)	FRED	1965:M1-2019:M12	level
37	2-Year Constant Maturity Treasury Yield (Market)	FRED	1976:M6-2019:M12	level
38	3-Year Constant Maturity Treasury Yield (Market)	FRED	1965:M1-2019:M12	level
39	5-Year Constant Maturity Treasury Yield (Market)	FRED	1965:M1-2019:M12	level
40	7-Year Constant Maturity Treasury Yield (Market)	FRED	1969:M7-2019:M12	level
41	10-Year Constant Maturity Treasury Yield (Market)	FRED	1965:M1-2019:M12	level
42	CRSP value-weighted return (ex dividends)	CRSP	1965:M1-2019:M12	ASSETS
43	CRSP value-weighted return (incl dividends)	CRSP	1965:M1-2019:M12	ASSETS
44	Log estimated dividend yield	CRSP	1965:M1-2019:M12	ASSETS
45	S&P 500 return	CRSP	1965:M1-2019:M12	ASSETS
46	Moody's Aaa Corporate Bond Yield	FRED	1983:M1-2019:M12	level
47	Moody's Baa Corporate Bond Yield	FRED	1986:M1-2019:M12	level
48	Real Oil Price	FRED	1965:M1-2019:M12	100 x log
49	Estimated dividends	CRSP	1965:M1-2019:M12	ASSETS
50	SP/CS U.S. National HPI	FRED	1987:M1-2022:M10	100 x log
51	Mortgage rate	BEA	1977:03-2023:12	100 x log
52	TFP	BFK (2006), updated	1947:06-2023:12	100 x log
53	TFP Utilization-adj	BFK (2006), updated	1947:06-2023:12	100 x log
54	GZ2012 Excess Bond Premium	GZ (2012), updated	1973:M1-2022:M7	level
55	Oil production	BH (2019), updated	1974:M1-2017:M12	100 x log
56	Oil stocks	BH (2019), updated	1974:M1-2017:M12	100 x log
57	World IP	BH (2019), updated	1974:M1-2017:M12	100 x log
58	US IP	BH (2019), updated	1974:M1-2017:M12	100 x log

use SVAR-IV as their baseline specifications for computing their IRFs. In reduced form, the VAR(p) is

$$\mathbf{y}_t = \mathbf{c} + \mathbf{A}_1\mathbf{y}_{t-1} + \mathbf{A}_2\mathbf{y}_{t-2} + \cdots + \mathbf{A}_p\mathbf{y}_{t-p} + \mathbf{v}_t$$

for data $\mathbf{y}_t$ ($k \times 1$), VAR coefficients $\mathbf{A}_1, \ldots, \mathbf{A}_p$, and VAR residuals $\mathbf{v}_t$. The VAR residuals are further decomposed into a linear combination of k orthogonal shocks $\boldsymbol{\varepsilon}_t$

$$\mathbf{v}_t = \mathbf{H}\boldsymbol{\varepsilon}_t,$$

where the impact matrix $\mathbf{H}$ ($k \times k$) is unknown. The vector $\mathbf{y}_t$ contains the variables in the baseline VAR specifications of Känzig (2021) and Gertler and Karadi (2015).

To construct an estimate of the shock of interest (ordered first in the shock vector $\boldsymbol{\varepsilon}_t$ and noted as ε_{1t}), we use SVAR-IV. Both applications provide an instrument z_t, which gives monthly high-frequency surprises around monetary or oil announcements. There are two conditions for the validity of SVAR-IV. First, the instrument z_t is relevant in that it is correlated with the shock of interest ($\mathbb{E}[\varepsilon_{1t}z_t] \neq 0$). Second, instrument z_t is uncorrelated with respect to other current shocks ($\mathbb{E}[\varepsilon_{2:kt}z_t] = 0$).

Ordering the endogenous variable (1Y Treasury for the monetary application and oil prices for the oil application) first, the column of the impact matrix corresponding to the target shock (e.g. the impact effect of oil prices on outcome j, like the price of motor fuel) is identified from

$$\mathbf{H}_{j1} = \frac{\mathrm{Cov}(v_{jt}, z_t)}{\mathrm{Cov}(v_{1t}, z_t)}, \quad j = 1, \ldots, k.$$

The above relationship is an implication of (1) the VAR residuals being composed of contemporaneous shocks ε_t and (2) instrument z_t being composed of the target shock and orthogonal measurement error. Finally, from the derivation in section 2.1.4 of Stock and Watson (2018), the target shock $\varepsilon_{1t} = \boldsymbol{\varphi}'\mathbf{v}_t$ where $\boldsymbol{\varphi} = \frac{\mathbf{H}'_{\cdot 1}\mathrm{Cov}(\mathbf{v}_t)^{-1}}{\mathbf{H}'_{\cdot 1}\mathrm{Cov}(\mathbf{v}_t)^{-1}\mathbf{H}_{\cdot 1}}$ from invertibility. Estimators are constructed using the corresponding sample moments, where the estimated target shock is noted as $\widehat{\varepsilon}_{1t}$.

Since the sample length of instrument z_t is shorter than the VAR sample length, $\mathbf{H}_{j1}$ is estimated using the instrument's sample. Equipped with estimator $\widehat{\mathbf{H}}_{j1}$, $\varepsilon_{1,t}$ is estimated using the VAR residual of the longer VAR sample (maximizing power).

The lag lengths for both the monetary and oil applications is set to $p = 12$. The monetary VAR parameters ($\mathbf{c}, \mathbf{A}_1, \ldots, \mathbf{A}_p$) are estimated using data from July 1979 to to June 2019. The monetary impact impulse responses $\mathbf{H}_{\cdot 1}$ are estimated using Federal Funds Futures surprises from January 1990 to June 2019. To sidestep concerns about time variation in parameters, the estimated monetary shocks $\widehat{\varepsilon}_{1,t}$ from January 1990 to June 2019 are used for the later analyses. The oil VAR parameters are estimated using data from January 1975 to December 2017, with impact impulse responses estimated using oil futures surprises data from April 1983 to December 2017. The estimated oil shock is from January

1975 to December 2017.

2. IMPULSE RESPONSE FUNCTION ESTIMATION: To compute the impulse response function of some outcome variable x_{nt}, $n \in \{1, \ldots, N\}$ to the target shock, we use an "internal instrument" approach (Plagborg-Møller and Wolf, 2021). That is, the IRF is estimated using a 12-lag VAR using vector $[\widehat{\varepsilon}_{1t}, \mathbf{y}_t', x_{nt}]'$ using the longest possible sample with non-missing observations, where impulse responses on impact are identified using a recursive causal ordering; that is, the estimated shock from Step 1 is ordered first.[62] For each bootstrap draw (details in the following paragraph), impulse responses for horizons $h \geq 1$ are computed by propagating the shock forward with the VAR model. To help account for the well-known small-sample biases prevalent in VARs, our final point estimate is the bootstrap bias-corrected IRF. Separate VARs are estimated for each outcome variables of interest to avoid the curse of dimensionality in finite samples.

   The confidence intervals for the IRF figures are computed using the moving block bootstrap of Jentsch and Lunsford (2019). Block lengths are determined by their rule of thumb $l = 5.03T^{1/4}$ for sample length T with 10,000 bootstrap iterations. For each outcome variable, the bootstrap draws are stored. For comparability with the IRF typically reported in time-series publications, we report 68% and 90% Hall bootstrap confidence intervals in the figures.

   The block bootstrap is also used for computing the standard errors on the welfare effects. The asymptotic variances of the impulse responses (for each horizon) and asympototic covariances *across* horizons are estimated using the bootstrap draws.

Consistent with Gertler and Karadi (2015) and Känzig (2021), our VAR's are estimated in levels. The specific variables and transformations are described in section B.7. Doing so allows for non-stationary variables to enter and is conservative in the sense that cointegrating relationships need not be specified. In response to a particular shock, the resulting horizon h IRF can be interpreted as the expected outcome variable with the shock relative to trend—specifically, the counterfactual outcome were the shock to have not occurred.

C.2 Standard Errors for Welfare Effects

The first-order change in money-metric welfare described in Proposition 1 is subject to sampling uncertainty. In this section, we quantify uncertainty arising from the estimation of the IRFs, conditioning on the cross-sectional moments.

We begin by establishing several preliminary results. Store the impulse responses of variable n for horizons 0 through h to some shock in vector $\mathbf{\Psi}^n = [\Psi_0^n, \Psi_1^n, \ldots, \Psi_h^n]'$. Store the impulse responses of variable 1 through N in vector $\mathbf{\Psi} = [\mathbf{\Psi}^{1\prime}, \mathbf{\Psi}^{2\prime}, \ldots, \mathbf{\Psi}^{N\prime}]'$. Assumption 1 (below) establishes consistency and asymptotic normality of the IRF estimator for outcome variable n $\widehat{\mathbf{\Psi}}^n$. These conditions are satisfied under standard time series techniques. Note the $(h+1) \times$

[62] One could alternately estimate direct regressions of the outcome variable on the estimated shock $\widehat{\varepsilon}_t$, but this is subject to attenuation bias coming from finite sample measurement error.

$(h+1)$ matrix $\mathbf{V}_{nn}$ stores the asymptotic covariance matrix of the IRFs of outcome variable n across horizons $0, \ldots, h$, which can be estimated by the bootstrap. Applying the Delta method, Proposition 5 (below) establishes that a linear combination of the IRFs with known weights is itself asymptotically normal.

Assumption 1 (Consistency and asymptotic normality)**.** *For true impulse response $\mathbf{\Psi}^n$, estimator $\widehat{\mathbf{\Psi}}^n$ is consistent $\widehat{\mathbf{\Psi}}^n \xrightarrow{p} \mathbf{\Psi}^n$ and is asymptotically normal $\sqrt{T}(\widehat{\mathbf{\Psi}}^n - \mathbf{\Psi}^n) \xrightarrow{d} N(0, \mathbf{V}_{nn})$ for asymptotic variance matrix $\mathbf{V}_{nn}$ as $T \to \infty$.*

Proposition 5. *Impose Assumption 1, let $\sqrt{T}(\widehat{\mathbf{\Psi}} - \mathbf{\Psi}) \xrightarrow{d} N(0, \mathbf{V})$, and take weights b_{it} as given. For welfare estimator $\widehat{W} = \sum_n \sum_t b_{nt} \widehat{\Psi}_t^n = \mathbf{b}'\widehat{\mathbf{\Psi}}$, denote the true welfare as $W = \sum_n \sum_t b_{nt} \Psi_t^n$. Then, the welfare estimator $\sqrt{T}(\widehat{W} - W) \xrightarrow{d} N(0, V_W)$ when $V_W = \mathbf{b}'V\mathbf{b} > 0$.*

Proof. The result immediately follows from an application of the Delta method. $\square$

Below are two notes on Proposition 5:

1. Within the framework established in Proposition 1, weights b_{nt} are computed taking data on cross-sectional expenditures and discounting $R_{0\to t}^{-1}$ as known. The weights also take into account the monthly-to-quarterly aggregation of impulse responses.

2. Proposition 5 can be extended to the case where $\widehat{\mathbf{\Psi}}^n$ are estimated on differing sample lengths across outcome variable N. Let $T_n = \omega_n T$ give the length of the sample used to estimate the IRF of variable n, where T is the length of the largest sample. That is, as T grows, T_n grows proportionally according to $\omega_n \in (0, 1]$. Then, for diagonal matrix function $\mathrm{diag}(\cdot)$,

$$\begin{pmatrix} \sqrt{T_1}(\widehat{\mathbf{\Psi}}_1 - \mathbf{\Psi}_1) \\ \sqrt{T_2}(\widehat{\mathbf{\Psi}}_2 - \mathbf{\Psi}_2) \\ \vdots \\ \sqrt{T_n}(\widehat{\mathbf{\Psi}}_n - \mathbf{\Psi}_n) \end{pmatrix} = \mathrm{diag}(\sqrt{\omega_1}, \ldots, \sqrt{\omega_n})\sqrt{T}(\widehat{\mathbf{\Psi}} - \mathbf{\Psi}).$$

   Pre-multiplying the above display by $\mathrm{diag}(\sqrt{\omega_1}, \ldots, \sqrt{\omega_n})^{-1}$ and applying Slutsky's lemma characterizes the asymptotic distribution of $\sqrt{T}(\widehat{\mathbf{\Psi}} - \mathbf{\Psi})$ so Proposition 5 can be applied.

To estimate V_W, an estimate of the asymptotic covariance matrix of impulse responses V is required. From the ordering of $\mathbf{\Psi}$, the matrix V can be decomposed into a collection of $(h+1) \times (h+1)$ block matrices:

$$V = \begin{bmatrix} V_{11} & V_{21}' & \ldots & V_{n1}' \\ V_{21} & V_{22} & \ldots & V_{n2}' \\ \vdots & & \ddots & \\ V_{N1} & V_{N2} & \ldots & V_{NN} \end{bmatrix}.$$

Here, $V_{nn'}$ stores the asymptotic variance matrix corresponding to the IRFs of variables n and n'. Recall that the internal VAR strategy described in the main text uses the longest possible sample for each outcome variable, so the IRFs across outcome variables are estimated under differing sample lengths. Note V_{nn} can be computed from the resulting bootstrap draws. Computing $V_{nn'}$, however, requires a shared estimation sample. As a working assumption, we impose block-wise uncorrelatedness, so that $V_{nn'} = 0$ for $n \neq n'$. Doing so allows us to exploit all available information for each outcome variable.

To understand the practical implications of blockwise-uncorrelatedness, we consider an exercise where IRFs are estimated under a (shorter) shared estimation sample (1999:M1-2017:M12)—using the same random seed across outcome variables $n = 1, \ldots, N$, the same blocks in the block bootstrap are drawn. Doing so allows for the estimation of the full asymptotic covariance matrix, including off-diagonal blocks $V_{nn'}$ for $n \neq n'$. Figure C1 compares the confidence interval lengths for V_W computed under full information and under blockwise-uncorrelatedness for both the oil and monetary shocks. For both the oil shock and monetary shock, we find that the standard errors under blockwise-uncorrelatedness are generally conservative. For the some college and bachelor's+ groups, confidence intervals under blockwise-uncorrelatedness are larger than those computed under full information. The confidence intervals for the HS or less category computed under blockwise-uncorrelatedness are roughly the same length over the life cycle.

C.3 Estimating Squared Impulses

For this subsection, let $X_{1,t}$ and $X_{2,t}$ be two covariance stationary processes determined by $K < \infty$ mutually uncorrelated white noise shocks with absolutely summable weights $\sum_{l=0}^{\infty} |\theta_{i,k,l}| < \infty$:

(C1)
$$X_{i,t} = \sum_{l=0}^{\infty} \sum_{k=0}^{K} \theta_{i,k,l} \varepsilon_{k,t-l} \quad i = 1, 2, \quad \varepsilon_{k,t} \sim WN(0, \sigma_k^2), \quad \mathbb{E}[\varepsilon_{m,t}\varepsilon_{n,t-l}] = 0 \text{ for } m \neq n \text{ and } l \geq 0.$$

The moving average coefficients $\theta_{i,k,l}$ can also be interpreted as the IRF of shock k on variable i at horizon l.

The following Lemma shows that the impulse of the product of $X_{1,t}$ and $X_{2,t+m}$ equals the product of their associated IRFs.

Lemma 2. *For $X_{1,t}$ and $X_{2,t}$ defined in equation C1,*

$$\mathbb{E}[X_{1,t}X_{2,t+m} | \varepsilon_{k^*,t-h} = 1] - \mathbb{E}[X_{1,t}X_{2,t+m} | \varepsilon_{k^*,t-h} = 0] = \theta_{1,k^*,h}\theta_{2,k^*,m+h}.$$

FIGURE C1: Relative Confidence Interval Length over the Life Cycle and by Household Education

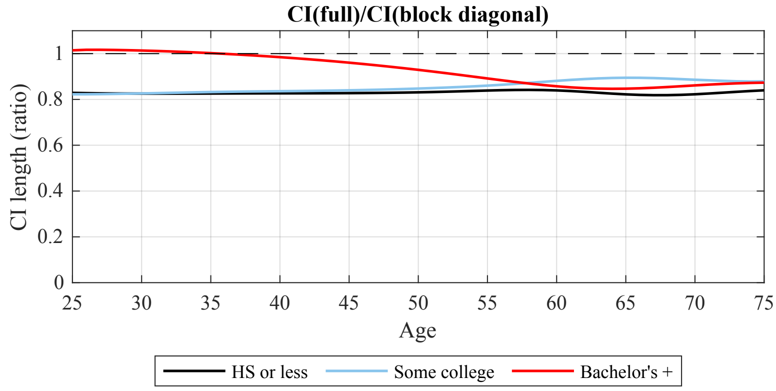

PANEL A: OIL SHOCK

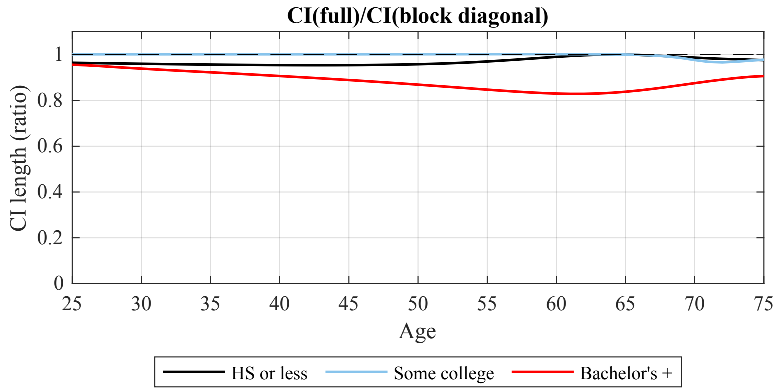

PANEL B: MONETARY SHOCK

Notes: The above panels show the ratio of confidence interval lengths computed using the full impulse response covariance matrix ("full") and blockwise-uncorrelatedness ("block diagonal"). The specifications are estimated using data from 1999:M1–2017:M12.

Proof. The conditional covariance of $X_{1,t}$ and $X_{2,t+m}$ given $\varepsilon_{k^*,t-h} = e$ is

$$\begin{aligned}
\mathbb{E}\Big[X_{1,t}X_{2,t+m}|\varepsilon_{k^*,t-h} = e\Big] &= \mathbb{E}\Big[\big(\sum_{l=0}^{\infty}\sum_{k=0}^{K}\theta_{1,k,l}\varepsilon_{k,t-l}\big)\big(\sum_{l=0}^{\infty}\sum_{k=0}^{K}\theta_{2,k,l}\varepsilon_{k,t+m-l}\big)|\varepsilon_{k^*,t-h} = e\Big] \\
&= \mathbb{E}\Big[\big(\sum_{l=0}^{\infty}\sum_{k=0}^{K}\theta_{1,k,l}\varepsilon_{k,t-l}\big)\big(\sum_{l=-m}^{\infty}\sum_{k=0}^{K}\theta_{2,k,l+m}\varepsilon_{k,t-l}\big)|\varepsilon_{k^*,t-h} = e\Big] \\
&= \sum_{l=0}^{\infty}\sum_{k=0}^{K}\theta_{1,k,l}\theta_{2,k,l+m}\mathbb{E}\big(\varepsilon_{k,t-l}^2|\varepsilon_{k^*,t-h} = e\big).
\end{aligned}$$

The second line of the above reindexes the second summation. The third line follows from serial and mutual uncorrelatedness of the shocks. The result then follows. □

Applying this Lemma, we see that the squared impulse $\mathbb{E}[X_{i,t}^2|\varepsilon_{k^*,t-h} = 1] - \mathbb{E}[X_{i,t}^2|\varepsilon_{k^*,t-h} = 0]$ equals the square of the associated IRF.

Corollary 1. *For $X_{i,t}$ defined in equation C1,* $\mathbb{E}[X_{i,t}^2|\varepsilon_{k^*,t-h} = 1] - \mathbb{E}[X_{i,t}^2|\varepsilon_{k^*,t-h} = 0] = \theta_{i,k^*,h}^2$.

D. ROBUSTNESS

D.1 Second-Order Effects

Here, we detail our approach to sizing the behavioral elasticities required to compute second-order welfare effects. To size the consumption elasticities with respect to goods' prices, we assume a homothetic constant elasticity of substitution (CES) consumption aggregator in which no individual good occupies a large share of the consumption basket. Under this assumption, own-price elasticities are equal to the elasticity of substitution and cross-price elasticities are zero. We consider an elasticity of substitution of 4, in line with estimates from Hottman et al. (2016). We then make the extreme assumption that households spend all of any increase in income, so that income elasticities of consumption are unity. This implies elasticities with respect to dividends, labor income, and transfer income are 1, while the elasticity with respect to asset prices is -1, so that an increase in asset prices is taken as a pure income effect. We also weight the effects from each of these by the share of that income source with respect to total income.[63] This number is an upper bound on the size of these behavioral elasticities since households likely save a portion of income increases; this assumption therefore gives a stronger chance for the second-order effect to be large.

Second, we use a high estimate of 2 for the Frisch elasticity of labor supply.[64] To size the response of labor supply to goods' prices, dividends, and transfer income, we follow Cesarini et al. (2017) and consider the largest labor supply elasticity reported: -0.2. In particular, we assume that prices lead to pure income effects, so the elasticity with respect to the prices of different goods is positive. We also weight the effect of each category by the share of that category in the consumption basket.[65] Finally, we follow Chodorow-Reich et al. (2021) and consider an elasticity of labor supply to asset prices of 0.035.[66]

Third, to bound the asset accumulation elasticities with respect to own asset prices we follow Gabaix et al. (2023) and assume an elasticity of -0.02.[67] Also, since households have to be indifferent between buying at price Q when dividends are D and buying at λQ when dividends are λD for any $\lambda \in \mathbb{R}^+$, the elasticity of asset accumulation to own dividends must be equal in magnitude, but opposite in sign, so we set it to 0.02. As we do for consumption categories, we set cross-price elasticities for asset classes to be zero. In addition, we assume asset accumulation responds one-to-one with income, so that the elasticities of asset accumulation to goods' prices (since we assume goods' price changes are pure income effects), labor income, and transfer income are all unity. This contradicts what we assumed for consumption, but our aim is to make the second-order effects as large as is reasonable. We also note that assuming a

[63] Note that from the budget constraint, total income is the sum of labor income, transfers, and dividends. This implies that the shares do not add up to 1 since the "income" coming from asset sales is not part of income.

[64] Chetty, Guren, Manoli, and Weber (2013) suggests a much lower elasticity is more consistent with microeconomic studies, but our goal here is to give the greatest change for second-order effects to matter.

[65] This share is computed using the cross-sectional averages for each category.

[66] Note this estimate captures the general equilibrium response of hours to stock market wealth.

[67] The authors find a (wealth-weighted) response of asset purchases 0.1% to a 10% change in the stock market. To give second-order effects as large a chance as possible to overturn our results, we double this implied elasticity.

constant behavioral elasticity may be problematic in extreme crisis or inflationary episodes.

D.2 Robustness: Projecting "No-shock" Consumption, Wage and Asset Holdings

Our framework requires projections for each of the components in Proposition 1 forward through time from the moment of the identified impulse. In our baseline results, we assume 2019 is a steady state: prices, wages, and dividends streams in expecation would remain fixed at their 2019 levels. However, different products have experienced different trend inflation rates, 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 life cycle. 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 $\pi^{ag}_{C,j}$. That is, we estimate the following regression for each good j and each age and education group using CEX data:

$$\ln\left(p_{jt}c^{a,g}_{jt}\right) = \pi^{a,g}_{C,j} \cdot t + \varepsilon^{a,g}_{j,t}. \tag{D1}$$

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

$$p_{jt}c^{ag}_{jt} = p_{j0}c^{(a+t)g}_{j0} \cdot (1 + \pi^{ag}_{C,j})^t.$$

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 earnings, we follow a similar approach. Given a life cycle profile of earnings at $t = 0$, we compute earnings at t as

$$W^{ag}_t = W^{(a+t)g}_0 (1 + \pi^{ag}_W)^t$$

where π^{ag}_W is trend wage inflation of age a households in education group g, estimated similarly to equation (D1).

For assets, we proceed as follows. The observed variables in the SCF are $Q_{k0}N^{ag}_{k0}$. Following the approach mentioned in Section 3 in the main text, we can estimate $Q_{k0}\Delta N^{ag}_{k0}$. To obtain the $D_{k0}N^{ag}_{k0}$ series for equity, note we can rewrite

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

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

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

The dividend yield for bonds is simply the bond yield, while the dividend yield for equities is publicly available information. Again, we estimate the trend in dividends for each asset class (π_k^D) similarly to equation (D1). 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.

D.3 Asset Values in the Utility Function

When asset *values* $Q_{kt}N_{kt}$, as opposed to quantities, appear in the utility function, an additional term appears in the welfare formula. Specifically, suppose the problem of the consumer is

$$V(\{N_{k0}\}_k) = \max_{\{\{c_{jt}^{ag}(s_t)\}_j, L_t^{ag}(s_t), \{N_{kt}^{ag}(s_t)\}_k\}_{t=0,s}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta_t^{ag}\delta_t^{ag} U(C_t^{ag}(s_t), \{Q_{kt}(s_t)N_{kt}^{ag}(s_t)\}_{k=1}^K, L_t^{ag}(s_t)),$$

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

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

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

$$\lim_{T\to\infty} \mathbb{E}_0[R_{0\to T}^{-1}N_{kT}^{ag}Q_{kT}] \geq 0, \qquad \forall k \in \{0,\ldots,K\}.$$

Then, as both aggregate and idiosyncratic risk become small, the welfare change from an impulse to element n of the fundamental shock vector is:

$$\begin{aligned} dV = \sum_t R_{0\to t}^{-1}\Bigg(& \underbrace{-\sum_j p_{j,t}c_{jt}^{ag}\Psi_{n,t}^{p,j}}_{\text{Consumption Price Changes}} + \underbrace{W_t^{ag}L_t^{ag}\Psi_{n,t+h}^{W^{ag}}}_{\text{Labor Income Changes}} + \\ & + \sum_k [\underbrace{N_{kt-1}^{ag}D_{kt}\Psi_{n,t}^{D,k}}_{\text{Asset Income Changes}} - \underbrace{Q_{kt}\Delta N_{kt}^{ag}\Psi_{n,t}^{Q,k}}_{\text{Asset Price Changes}}] + \underbrace{T_t^{ag}\Psi_{n,t}^{T^{ag}}}_{\text{Transfer Income Changes}} \Bigg) \\ & + \sum_t \underbrace{R_{0\to t}^{-1}\left(1 - R_{0t}^{-1}\mathbb{E}_0[\frac{Q_{kt+1}}{Q_{kt}}]\right) Q_{kt}N_{kt}\Psi_{n,t}^{Q,k}}_{\text{Asset Value Fluctuations}}. \end{aligned} \tag{D2}$$

The intuition is as follows: asset value fluctuations impact welfare according to the marginal utility of that asset. If households are to hold them, assets that have a lower expected return have a higher marginal utility in the zero-risk limit, which is all that matters to a first order. To see this, note

$$\frac{dV}{d\sigma} = \sum_t \beta_t \delta_t \sum_{s_t} \pi_t(s_t)\lambda_t(s_t)\Bigg(-\sum_j p_{j,t}(s_t)c_{jt}(s_t)v^p_{jt} + W_t(s_t)e^i_t(s^i_t)L_t(s_t)v^W_t$$
$$+ T_t(s_t)v^T_t(s_t) + \sum_k \Big[N_{kt-1}(s_{t-1}))D_{kt}(s_t)v^D_{kt}(s_t) - Q_{kt}(s_t)\Delta N_{kt}(s_t)v^Q_{kt}(s_t)\Big]\Bigg)$$
$$+ \sum_t \beta_t \delta_t \sum_{s_t} \pi_t(s_t) \sum_{k\neq 0}\Bigg(U_{N_k}(C_t(s_t), \{Q_{kt}N_{kt}(s_t)\}_k, L_t(s_t)))N_{kt}(s_t)Q_{kt}v^Q_{kt}(s_t)\Bigg).$$

Now the FOC for N_{kt} for $k \neq 0$ is

$$\beta_t\delta_t U_{N_k}(s_t)Q_{kt}(s_t) = \beta_t\delta_t\lambda_t(s_t)Q_{kt}(s_t) - \beta_{t+1}\delta_{t+1}\mathbb{E}_t[\lambda_{t+1}(s_t)Q_{kt+1}(s_t)|s_t],$$

or, as aggregate and idiosyncratic risk becomes small,

$$U_{N_k} = \lambda_t\left(1 - R^{-1}_{0,t}\mathbb{E}_0\left[\frac{Q_{kt+1}}{Q_{kt}}\right]\right)$$

Following the strategy of the general proof above then gives (D2).

E. Two-Asset HANK Model Appendix

This section proides further details on the two-asset HANK model used to to validate the feasible set approach to computing welfare as described in section 9. It first describes the model environment, before providing additional details on how we compute welfare changes within the model. Finally, it presents additional results from the model, such as the model-implied impulse responses of various objects and welfare changes throughout the state space.

E.1 Model Description

The model is the same as that of Auclert et al. (2021) in both specification and (baseline) calibration. Our description follows theirs closely.

Households. Households are subject to uninsurable idiosyncratic labor earnings risk and may save either in a liquid or an illiquid account. We let the quantity of liquid savings held by household i in period t be given by b_{it}, whereas the quantity of illiquid savings is given by a_{it}. Illiquid assets are subject to a convex portfolio adjustment cost $\Phi_t(a_{it}, a_{it-1})$. Liquid assets pay a coupon rate r^b_t while illiquid assets pay a coupon r^a_t. Households supply labor N_t, which is the same for all households and determined by a labor union to be described below. Household earnings are the product of their realization of idiosyncratic productivity e^i_t, an aggregate wage w_t, labor supply N_t, and the net-of-tax rate $(1-\tau_t)$. Each period, the household observes their realization of labor earnings risk e^i_t as well as their holdings of the two assets and choose their consumption c_{it} and next period's holdings of the two assets. The Bellman equation summarizing the household problem is[68]

$$\text{(E1)} \qquad V_t(e^i_t, b_{it-1}, a_{it-1}) = \max_{c_{it}, b_{it}, a_{it}} \frac{c_{it}^{1-\sigma}}{1-\sigma} - \varphi \frac{N^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}} + \beta \mathbb{E}_t V_{t+1}(e_{it+1}, b_{it}, a_{it})$$

$$\text{s.t.} \quad c_{it} + a_{it} + b_{it} = (1-\tau_t) w_t N_t e^i_t + (1+r^b_t) b_{it-1} + (1+r^a_t) a_{it-1} - \Phi_t(a_{it}, a_{it-1})$$

$$a_{it} \geq 0, \qquad b_{it} \geq \underline{b}.$$

We specify the adjustment cost of illiquid assets as

$$\Phi_t(a_{it}, a_{it-1}) = \frac{\chi_1}{\chi_2} \left| \frac{a_{it} - (1+r^a_t) a_{it-1}}{(1+r^a_t) a_{it-1} + \chi_0} \right|^{\chi_2} \left((1+r^a_t) a_{it-1} + \chi_0 \right),$$

where $\chi_0, \chi_1 \geq 0, \chi_2 > 1$. Finally, we assume idiosyncratic earnings evolve according to

$$e_{it+1} = \rho e^i_t + \sigma_e \epsilon_{it}, \qquad \epsilon_{it} \sim \mathcal{N}(0,1).$$

Financial Intermediary. A representative, competitive financial intermediary takes liquid and illiquid deposits from households and invests in firm equity at price p_t and in government bonds B^g_t. It performs liquidity transformation at proportional cost $\omega \int b_{it} di$. A no-arbitrage condition ensures that the economywide ex-ante return on assets $\mathbb{E}[1+r_{t+1}]$ equals the expected returns on nominal government bonds and equity. These returns are passed onto households, subject to intermediation costs, so that

$$\mathbb{E}_t[1+r^a_{t+1}] = \frac{1+i_t}{\mathbb{E}_t[1+\pi_{t+1}]} = \frac{\mathbb{E}_t[d_{t+1}+p_{t+1}]}{p_t} = \mathbb{E}_t[1+r^b_{t+1}] + \omega$$

where i_t is the nominal interest rate on government bonds, π_t is inflation in the price level and d_t is dividends paid on equity. Whereas the above equation holds ex-ante due to no-arbitrage, ex-post returns are subject to surprise inflation and capital gains. We assume capital gains (movements in p_t) accrue to the illiquid account, so that

$$1 + r^a_t = \Theta_p \left(\frac{d_t + p_t}{p_{t-1}} \right) + (1-\Theta_p) \left(\frac{1+i_{t-1}}{1+\pi_t} \right)$$

for Θ_p the share of equities in the illiquid portfolio.

Firms. Final goods Y are produced by a competitive final goods producer that aggregates a continuum of intermediate goods, indexed by j, with a constant elasticity of substitution $\mu/(\mu-1) > 1$. Intermediate goods are produced by monopolistically competitive producers that operate a constant-returns Cobb-Douglas production function in capital k_{jt} and labor n_{jt}: $y_j = A_t k^{\alpha}_{jt-1} n^{1-\alpha}_{jt}$. Firms choose their capital stock subject to quadratic adjustment costs

[68] Note that we have written the household problem in its real form; we could alternately have multiplied by an aggregate price index to make the problem nominal.

$\zeta\left(k_{jt}/k_{jt-1}\right)k_{jt-1}$ with $\zeta(x) \equiv x-(1-\delta)+\frac{(x-1)^2}{2\delta\epsilon_I}$, where $\delta > 0$ is the depreciation rate and $\epsilon_I > 0$.

Firms set the price of their products with quadratic adjustment costs

$$\psi_t(p_{jt}, p_{jt-1}) = \frac{\mu}{\mu-1}\frac{1}{2\kappa_p}[\ln(p_{jt}/p_{jt-1})]^2 Y_t,$$

for κ_p a parameter governing the degree of price rigidity in the economy. Solving and linearizing the firm's optimal pricing problem under the symmetric equilibrium gives the following price Phillips Curve for aggregate inflation π_t:

$$\text{(E2)} \qquad \ln(1+\pi_t) = \kappa_p\left(mc_t - \frac{1}{\mu}\right) + \frac{1}{1+r_{t+1}}\frac{Y_{t+1}}{Y_t}\ln(1+\pi_{t+1}),$$

where $mc_t \equiv (1-\alpha)\frac{w_t}{A_t}\left(\frac{K_{t-1}}{N_t}\right)^{\alpha}$ is the marginal cost of production in period t. Aggregate investment is given by $I_t = K_t-(1-\delta)K_{t-1}+\zeta(K_t/K_{t-1})K_{t-1}$. Dividends equal output net of investment, labor costs and price adjustment costs: $d_t = Y_t - w_tN_t - I_t - \psi_t$. Tobin's Q and the capital stock evolve according to

$$\text{(E3)} \qquad Q_t = 1+\frac{1}{\delta\epsilon_I}\frac{K_t-K_{t-1}}{K_{t-1}}$$

$$\text{(E4)} \qquad (1+r_{t+1})Q_t = \alpha\frac{Y_{t+1}}{K_t}mc_t - \left[\frac{K_{t+1}}{K_t}-(1-\delta)+\frac{1}{2\delta\epsilon_I}\left(\frac{K_{t+1}-K_t}{K_t}\right)^2\right]+\frac{K_{t+1}}{K_t}Q_{t+1}.$$

Unions. A labor-aggregating firm combines a continuum of differentiated labor services with constant elasticity of substitution $\mu_w/(\mu_w-1) > 1$. We assume each household type supplies every type of labor. Wages are set by a type-specific labor union that sets wages to maximize the average utility of households, taking as given their consumption-savings decisions. There is a quadratic adjustment cost on the wages of an arbitrary type of labor k: $\psi_t^w(w_{kt}, w_{kt-1}) = \frac{\mu_w}{\mu_w-1}\frac{1}{2\kappa_w}[\ln(w_{kt}/w_{kt-1})]^2$, where κ_w is a parameter summarizing the degree of wage rigidity in the economy. In the symmetric equilibrium, aggregate wage inflation $1+\pi_t^w = (1+\pi_t)w_t/w_{t-1}$ evolves according to the wage Phillips Curve, derived by linearizing the union's wage-setting problem:

$$\text{(E5)} \qquad \ln(1+\pi_t^w) = \kappa_w\left(\varphi N_t^{1+\nu} - \frac{(1-\tau_t)w_tN_t}{\mu_w}\int e_t^i c_{it}^{-\sigma}di\right)+\beta\ln(1+\pi_{t+1}^w).$$

Policy. The fiscal authority spends G_t, issues one-period nominal bonds B_t^g and balances it budget every period by raising taxes, so that $\tau_t w_t N_t = r_t B_t^g + G_t$. Monetary policy behaves according to a Taylor rule so that the nominal rate i_t is

$$i_t = r_t^* + \phi_\pi\pi_t + \phi_y(Y_t - Y_{ss}).$$

The Fisher equation is $1 + r_t = (1 + i_t)/(1 + \pi_t)$. A monetary shock is a shock to r_t^*.

Market Clearing. The final good is used for consumption, investment, liquidity transformation and adjustment costs, yielding the goods market clearing condition

$$Y_t = \int (c_{it} + \omega b_{it-1} + \Phi(a_{it}, a_{it-1}))di + G_t + I_t + \psi_t.$$

Asset market clearing implies that total saving by the household equals the value of firm equity and government bonds:

$$p_t + B_t^g = \int a_{it} + b_{it} di.$$

Calibration. A period is a quarter. We calibrate our model following Auclert et al. (2021) so that the steady state matches some key features of the data. This calibration is summarized in Table E1. We calibrate the disutility of labor and steady-state TFP so that aggregate output and labor supply equal 1 in steady state. The markup is chosen so that total wealth, inclusive of stock market wealth, is equal to fourteen times aggregate GDP. We set households' discount factor to target an economy-wide interest rate of 1.25%, and the scale of the illiquid asset's adjustment cost so that household holdings of liquid assets is 104% of GDP. The idiosyncratic income process is chosen to generate a cross-sectional standard deviation of log earnings equal to 0.5, with autocorrelation 0.966.

E.2 Computation Details: Steady State and Impulse Responses

We solve the model using the Sequence-Space Jacobian package provided by Auclert et al. (2021). This package solves the household problem following the endogenous gridpoint method of Carroll (2006). We discretize households' state space so that they have three grid points for idiosyncratic earnings, fifty points on a liquid asset grid, and seventy grid points on the illiquid asset grid. We then compute impulse responses to shocks following the sequence space Jacobian methodology described in Auclert et al. (2021). This methodology computes the derivative of perfect-foresight equilibrium mappings between aggregate sequences around a steady state. Therefore, the methodology we employ calculates the first-order impulse responses to a given path of shocks to our exogenous variables.

Carroll's method for solving household problems provides a computationally efficient algorithm for computing policy functions and marginal values of assets. However, it does not directly return households' value functions, which is the key input to measuring welfare changes. One approach to computing value functions would be to perform a value function iteration, but this is computationally intense and subject to numerical errors. We therefore derive an alternative method to computing value functions.

Carroll's method returns $V_b(e, b, a)$ and $V_a(e, b, a)$ – the marginal value of an additional liquid and illiquid asset, respectively – and policy functions $c^*(e, b, a), b^*(e, b, a), a^*(e, b, a)$ at every point on the discretized state space. Our task is therefore to turn these marginal values into

TABLE E1: Baseline Two-Asset HANK Model Calibration

Parameter	Description	Value	Target
Households			
β	Discount Factor	0.983	r = 0.0125
σ	Coefficient of Relative Risk Aversion	1	
χ_0	Portfolio adj. cost pivot	0.25	
χ_1	Portfolio adj. cost scale	9.803	$B = 1.04Y$
χ_2	Portfolio adj. cost curvature	2	
$\underline{b}$	Borrowing constraint	0	
ρ_e	Autocorrelation of idiosyncratic earnings	0.966	
σ_e	Standard deviation of idiosyncratic earnings	0.92	Std(log earnings) = 0.5
Labor Unions			
φ	Disutility of Labor	0.634	N = 1
ν	Inverse Frisch Elasticity of Labor Supply	1	
μ_w	Steady-state Wage Markup	1.1	
κ_w	Slope of Wage Phillips Curve	0.1	
Firms			
Z	Steady-state TFP	0.468	$Y = 1$
α	Capital Share	0.33	$K = 10Y$
μ	Steady-state markup	1.015	Wealth $\equiv p + B^g = 14Y$
δ	Depreciation	0.02	
κ_p	Slope of Price Phillips Curve	0.1	
Financial Intermediary			
ω	Intermediation cost/Liquidity Premium	0.005	
Policy			
τ	Steady-state Labor Income Tax	0.356	Budget Balance
G	Government spending	0.2	
B^g	Bond Supply	2.8	
ϕ_π	Taylor Rule Coefficient on π	1.5	
ϕ_y	Taylor Rule Coefficient on Output	0	

values. To this end, we employ the Fundamental Theorem of Calculus and note:

$$V_t(e,b,a) = \int_{\underline{b}}^{b} V_{b,t}(e,\tilde{b},a)d\tilde{b} + K_{b,t}(e,a) = \int_{\underline{a}}^{a} V_{a,t}(e,b,\tilde{a})d\tilde{a} + K_{a,t}(e,b). \tag{E6}$$

To recover the constants of integration, consider the value function for a given earnings level e when $b = \underline{b}$ and $a = \underline{a}$. By equation (E6), we have:

$$V_t(e,\underline{b},\underline{a}) = K_{b,t}(e,\underline{a}) = K_{a,t}(e,\underline{b}).$$

Now consider the value function when $a = \underline{a}$ but b is unrestricted. Equation (E6) implies

$$V_t(e,b,\underline{a}) = \int_{\underline{b}}^{b} V_{b,t}(e,\tilde{b},\underline{a})d\tilde{b} + K_{b,t}(e,\underline{a}) = K_{a,t}(e,b).$$

Using the fact that $K_{b,t}(e,\underline{a}) = K_{a,t}(e,\underline{b})$, this implies the following updating rule for $K_{a,t}(e,b)$:

$$K_{a,t}(e,b) = K_{a,t}(e,\underline{b}) + \int_{\underline{b}}^{b} V_{b,t}(e,\tilde{b},\underline{a})d\tilde{b}.$$

Analogously, we have the following updating rule for $K_{b,t}$:

$$K_{b,t}(e,a) = K_{b,t}(e,\underline{a}) + \int_{\underline{a}}^{a} V_{a,t}(e,\underline{b},\tilde{a})d\tilde{a}. \tag{E7}$$

Therefore, since Carroll's method returns estimates of $V_{a,t}(\cdot)$ and $V_{b,t}(\cdot)$, one can calculate any $K_{a,t}(e,b)$ and $K_{b,t}(e,a)$ as long as we have an estimate of $V_t(e,\underline{b},\underline{a})$.

To make progress, we use the definition of the value function:

$$V_t(e,b,a) = u(c_t^*(e,b,a)) - \nu(N_t^*) + \beta\mathbf{E}\left[V_{t+1}(e',b_t^*(e,b,a),a_t^*(e,b,a))\right],$$

where again, we have estimates of the optimal policy functions $c^*(\cdot), b^*(\cdot), a^*(\cdot)$ from Carroll's method. Evaluating this at $\underline{b},\underline{a}$ and using equations (E6) and (E7) gives us

$$\begin{aligned}
V_t(e,\underline{b},\underline{a}) &= u(c_t^*(e,\underline{b},\underline{a})) - \nu(N_t^*) + \beta\mathbb{E}\quad [V_{t+1}(e',b_t^*(e,\underline{b},\underline{a}),a_t^*(e,\underline{b},\underline{a}))] \\
&= u(c_t^*(e,\underline{b},\underline{a})) - \nu(N_t^*) + \beta\mathbb{E}\quad \left[\int_{\underline{a}}^{a_t^*(e,\underline{b},\underline{a})} V_{a,t+1}(e',b_t^*(e,\underline{b},\underline{a}),\tilde{a})d\tilde{a} + K_{a,t+1}(e',b_t^*(e,\underline{b},\underline{a}))\right] \\
&= u(c_t^*(e,\underline{b},\underline{a})) - \nu(N_t^*) + \beta\mathbb{E}\quad \left[\int_{\underline{a}}^{a_t^*(e,\underline{b},\underline{a})} V_{a,t+1}(e',b_t^*(e,\underline{b},\underline{a}),\tilde{a})d\tilde{a}\right. \\
&\qquad \left. + \int_{\underline{b}}^{b_t^*(e,\underline{b},\underline{a})} V_{b,t+1}(e',\tilde{b},\underline{a})d\tilde{b} + K_{a,t+1}(e',\underline{b})\right].
\end{aligned}$$

Let's now consider a steady state where $V_t(\cdot) = V(\cdot)$ for all t. Then, since $V(e,\underline{b},\underline{a}) = K_a(e,\underline{b})$, the above is a linear equation in $K_a(e,\underline{b})$, which can be solved by inverting some matrices. To see this, rewrite the above equation in vector form and solve, where notationally $\vec{X}(b,a)$ denotes a vector whose j^{th} entry corresponds to the j^{th} value on the grid for e:

(E8)

$$\vec{K}_a(\underline{b}) = (I-\beta\Pi)^{-1}\left[u(\vec{c^*}(\underline{b},\underline{a})) - \nu(N_t^*) + \beta\Pi\left(\int_{\underline{a}}^{\vec{a^*}(\underline{b},\underline{a})} \vec{V}_a(\vec{b^*}(\underline{b},\underline{a}),\tilde{a})d\tilde{a} + \int_{\underline{b}}^{\vec{b^*}(\underline{b},\underline{a})} \vec{V}_b(\tilde{b},\underline{a})d\tilde{b}\right)\right].$$

Using the policy functions $c^*(\cdot), b^*(\cdot), a^*(\cdot)$ and marginal value functions $V_a(\cdot), V_b(\cdot)$, one can then compute the constants of integration in equation (E6) for the lowest values of asset holdings using equation (E8), then use the updating rules (E7) to get the constants of integration for other values of (b,a), and finally use equation (E6) to get the estimate of the steady-state value function in a computationally efficient manner. For additional computational efficiency, we linearly interpolate the marginal value functions between the grid points of asset holdings, which allows us to analytically compute integrals of the marginal values.

E.3 Computation Details: Value Changes

This section describes our approach to computing value changes from shocks in the model.

E.3.1 Backward Induction for the True Welfare Change

Here we describe our approach to computing the full value change from a particular shock within the model. To fix ideas, we describe the process for a monetary shock, understanding that an analogous approach is used for a TFP shock.

1. Compute the steady state following the approach laid out above. Let $V_{SS}(\cdot)$ denote the steady-state value function.

2. Employ the sequence-space Jacobian methodology to compute the impulse response of all macro variables to the monetary shock. Specifically, compute the impulse response of wages w_t, labor supply N_t, taxes τ_t, and returns on liquid assets r_t^b and illiquid assets r_t^a.

3. Assume that the economy returns to steady state T periods after the beginning of the shock. For our exercise, we choose $T = 100$, so that the economy returns to steady state 25 years after the initial monetary shock. This implies $V_T(\cdot) = V_{SS}(\cdot)$.

4. Solve the household problem (E1) backwards for $V_{T-1}(\cdot)$ given that $V_T(\cdot) = V_{SS}(\cdot)$ and given the implied impulse response for $w_t, r_t^a, r_t^b, \tau_t$ and N_t. Note we invoke certainty equivalence to assume perfect foresight in aggregate variables.

5. Iterate on step 4 above to solve for $V_{t-1}(\cdot)$ given $V_t(\cdot)$ and aggregate impulse responses. Repeat to get an estimate of $V_0(\cdot)$.

6. Compute the value change as

$$dV^{FULL}(e, b, a) = V_0(e, b, a) - V_{SS}(e, b, a).$$

This algorithm yields the full value change of the monetary shock, inclusive of higher-order effects, risk, and occasionally binding constraints. The units of the value change is utils. We therefore calculate money-metric value changes in the model as

$$dV^{MM,FULL}(e, b, a) = dV^{FULL}(e, b, a)/u'(c^*(e, b, a)).$$

Finally, we convert money-metrics into shares of consumption by computing

$$dV^{MM,FULL}(e, b, a)/c^*(e, b, a)$$

E.3.2 Feasible Set Approach to Welfare

The following algorithm computes welfare in the model following the feasible set approach of Proposition 1 for an individual who begins at the point $x_0 \equiv (e_0, b_0, a_0)$ on the state space. We simulate a sample of S individuals who all begin at x_0 using the model's steady state. That is, we generate S paths of idiosyncratic earnings of length T, and compute paths of consumption c_{it}, liquid asset holdings b_{it} and illiquid asset holdings a_{it} using the steady-state policy functions $c^*(\cdot), b^*(\cdot), a^*(\cdot)$. From these paths of policy functions, we further generate paths of labor earnings $w_{ss}N_{ss}(1-\tau_{ss})e_t^i$, illiquid asset income $(1+r_{ss}^a)a_{it}$, and liquid asset income $(1+r_{ss}^b)b_{it}$. Finally, we calculate Euler equation wedges τ_t^{EE} as the gap between average consumption growth among people who begin at point x_0 and $\beta(1+r_{ss}^b)$:

$$\tau_t^{EE} = \frac{\mathbb{E}[c_{t+1}|x_0]}{\mathbb{E}[c_t|x_0]} \cdot \frac{1}{\beta(1+r_{ss}^b)} - 1.$$

Given these paths simulated from the steady state, as well as the estimated IRFs for w_t, N_t, τ_t, r_t^b and r_t^a, we can then compute the value change for an individual i as[69]

$$dV_i^{FS} = \sum_{t=0}^{T} \left(\frac{1}{1+r_{ss}^b}\right)^t \prod_{s=0}^{t} (1+\tau_s^{EE})^{-1} \Big(\underbrace{w_{ss}N_{ss}(1-\tau_{ss})e_t^i d\ln(w_t N_t(1-\tau_t))}_{\text{Labor Income}} + \underbrace{r_{ss}^b b_{it} d\ln r_t^b + r_{ss}^a a_{it} d\ln r_t^a}_{\text{Portfolio}} \Big).$$

Note that the constraint effect in the workhorse model is zero since there is no direct effect of the shock on the value of the constraint. We do so for $S = 1000$ households per initial point in the state space, then average over the S households to get an estimate of the average welfare effect for someone whose initial state is (e_0, b_0, a_0). We compute the feasible set approach for 1) households at the borrowing constraint for both liquid and illiquid assets, 2) every decile of the unconstrained liquid and illiquid asset distribution and 3) every value of idiosyncratic earnings realizations. Given three gridpoints for idiosyncratic earnings, this yields a value change implied by the feasible set approach for 3 $\times$ 11 $\times$ 11 = 363 initial points on the state space. When calculating the comparisons in Table 2, we compare the full value change against the feasible set approach at these 363 grid points, weighting by their mass in the steady-state distribution. For the calibrated shocks studied in Table 3, we compute the feasible set approach for the full state space.

To compare against our baseline results in the main text, we do not adjust for idiosyncratic risk in our model's feasible set approach exercises. The reason is that idiosyncratic risk-adjustments require particular functional forms on utility, which our baseline formula does not. Abstracting from the risk-adjustment allows us to see the bias that ignoring this feature introduces. Table 2 suggests this bias is small for reasonable calibrations of idiosyncratic risk.

Nevertheless, the model permits a validation of our estimates of the risk-adjustment factors Θ. Given our simulated path of c_{it} and an assumed utility function, we can compute a simulated path for $u'(c_{it})$. We then use these marginal utility paths to compute the adjustment factors Θ's in the model. We compute these as we do in the data: by constructing the cross-sectional covariance between marginal utilities of consumption with idiosyncratic earnings and asset returns for everyone who begins at a certain point in the state space as of period 0. We present these results below.

[69]Note that, since the workhorse model features only one consumption good, one could add and subtract the consumption price effect, turning the labor income and portfolio channels into their nominal counterparts.

E.4 Additional Model Results

This section reports additional results from the two-asset HANK model. Figure E1 reports the impulse responses of various prices implied by the two-asset HANK model in response to different shocks. The solid red line is the response to a 25 basis point cut policy rates, which decays with AR(1) persistence 0.4. The solid black line is the response to a 1 percent decline in aggregate TFP which decays with persistence 0.9. These solid lines are the inputs to Table 2. The dashed red line is the response to a 25 basis point monetary shock as identified from the data. The policy rate is assumed to follow the IRF reported in Panel C of Figure 1. The dashed black line is the response to a 10% increase in the price of oil, modeled here as a TFP shock. The path for the shock matches the impulse response of TFP to the Känzig (2021) oil supply shock studied in the text. These dashed lines are the inputs to the exercise of Table 3.

Figure E2 plots estimated values for the idiosyncratic risk-adjustment factors Θ for each point on the initial state space. We assume log utility for this exercise, as we do for our exercise in the data. We find Θ^{w} runs between -0.4 and -0.2 and is largest in magnitude for borrowing constrained households and households with low initial earnings. Meanwhile, Θ^{r_a} runs between 0 and -0.2 whereas Θ^{r_b} runs between 0 and -0.35. The smaller adjustment factor for asset a reflects its illiquidity. These numbers validate those found within the data.

Figure E3 shows money-metric value changes, scaled by steady-state one-quarter consumption, across the state space. The left panels (A, C and E) report value changes in response to expansionary monetary policy shocks, whereas the right panels (B, D and F) report changes in response to a 1% TFP decline. All shocks are assumed to decay according to an AR(1). The top panels (A and B) report value changes for those with the lowest level of idiosyncratic earnings as of period 0, while the middle (C and D) and bottom (E and F) panels report the value changes for those with middle and highest levels of discretized idiosyncratic earnings, respectively. The figure is a 3-D plot of value changes for each point on the initial state space as of time 0 following the specification of the "baseline" column of Table 2. The figure shows that value changes are largest for those closest to the borrowing constraints, but largest for high-earning households, because the labor income effect is strongest for these high-earning households.

Figure E4 plots the change in money-metric welfare—as a share of consumption—as a result of an inflationary oil supply shock (panels A and B) and monetary shock (panels C and D) against a household's log four-year consumption. Panels A and C construct these plots within the model. Panels B and D report the welfare changes obtained from applying Proposition 1, where each dot represents an age $\times$ education. Panels A and C report model-implied estimates of value changes, where each dot is an initial point on the discretized state space. The size of the dots is proportional to the share of the population contained within that dot at the steady-state distribution, while the horizontal axis for the model plots is log steady-state consumption. To generate a comparable horizontal scale, we construct consumption relative to the mean in both figures.

The model generates significantly larger welfare effects of these shocks. Furthermore, house-

FIGURE E1: Model-Implied Impulse Response Functions

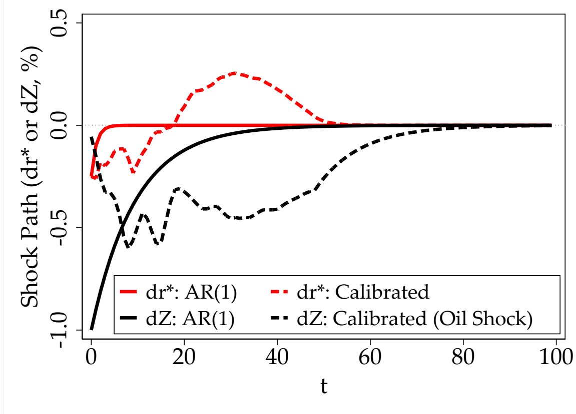

PANEL A: SHOCKS dr^* AND dZ

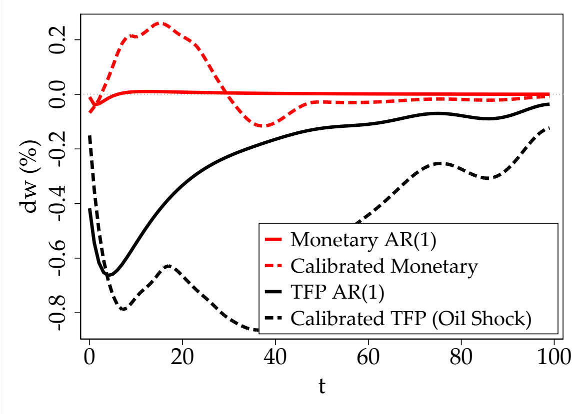

PANEL B: REAL WAGES w_t

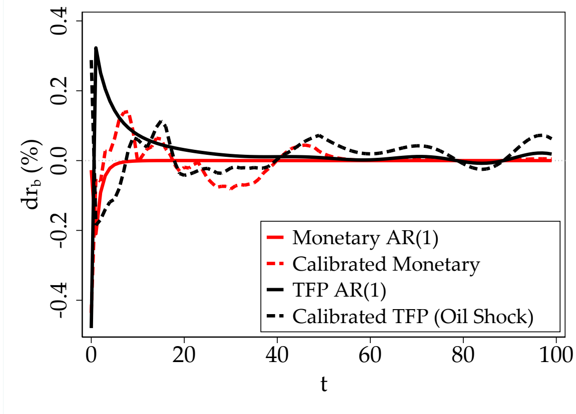

PANEL C: LIQUID RETURN r_t^b

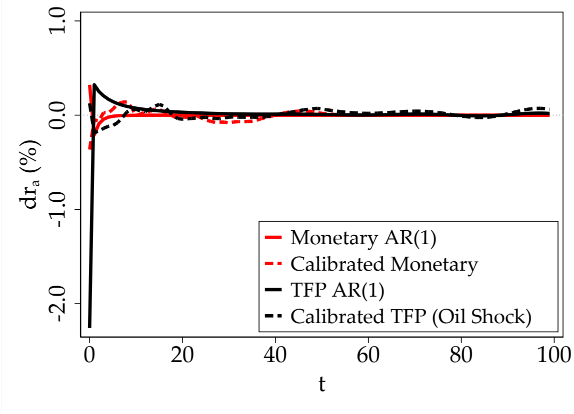

PANEL D: ILLIQUID RETURN r_t^a

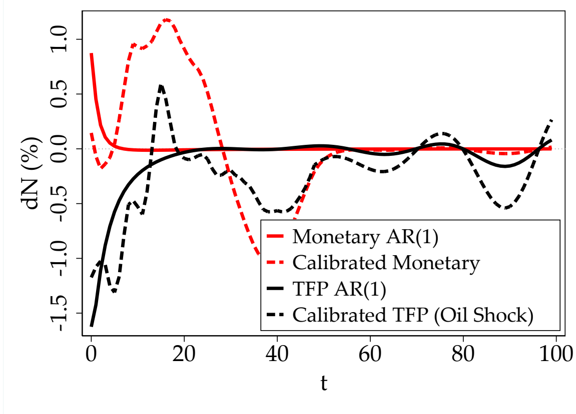

PANEL E: LABOR SUPPLY N_t

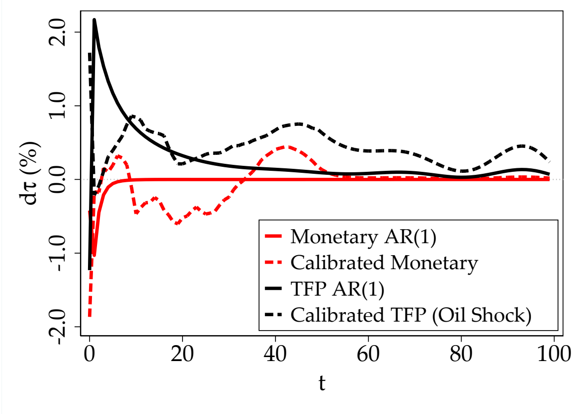

PANEL F: TAX RATE τ_t

Notes: Figure reports equilibrium impulse response functions in a benchmark two-asset HANK model. Red lines indicate responses to monetary cuts, black lines are responses to negative TFP shocks. Solid lines consider shocks which decay according to an AR(1) with persistence 0.4 (monetary) or 0.9 (TFP). Dashed lines consider shocks calibrated to impulse responses to Gertler and Karadi (2015) monetary shocks or Känzig (2021) oil shocks as estimated from the data. Plots represent percentage point deviations from steady-state values.

FIGURE E2: Distribution of Model-Implied Risk-Adjustment Factors Θ

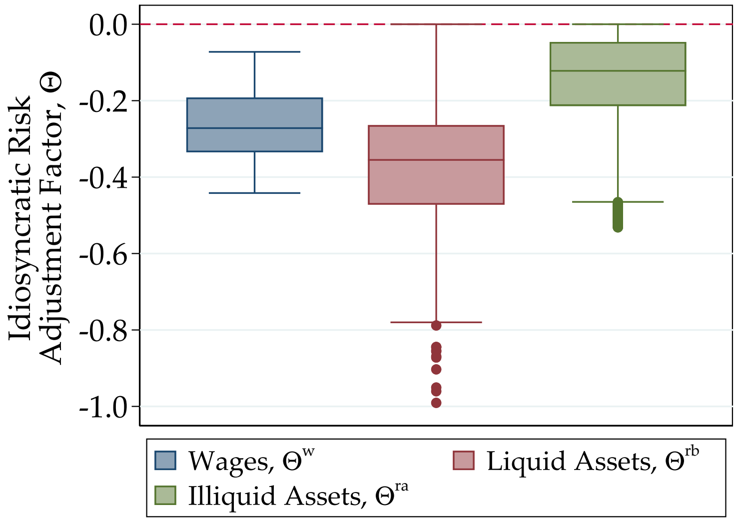

Notes: Figure plots the distribution – across different initial points on the state space – of estimated risk-adjustment factors Θ within the two-asset HANK model. The box represents the interquartile range of the estimated Θ with the line in the middle of the box representing the median. The whiskers represent the min and max of the distribution of Θ's, or the 25th or 75th percentile plus or minus 1.5 times the interquartile range, whichever is tighter. Dots are outliers from these inferred "min" and "max". The left blue box plot represents the risk-adjustment factor for labor earnings Θ^w, the middle red box represents the factor for liquid assets Θ^{r_b}, and the right green box plots the factor for illiquid assets Θ^{r_a}.

holds with higher earnings in the model have higher consumption and larger welfare responses to shocks, particularly from monetary shocks. But given an earnings level, households with more consumption (based on their initial asset position) are less affected by both oil and monetary shocks. This stands in stark contrast to the data, which exhibit a highly non-monotone relationship between consumption and value changes. This difference reflects the complex asset accumulation patterns found in the data.

FIGURE E3: Estimated Value Changes in HANK Model

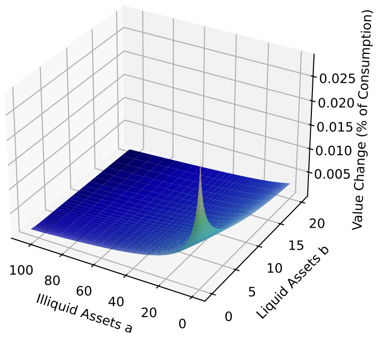

PANEL A: MONETARY SHOCK, LOW EARNINGS

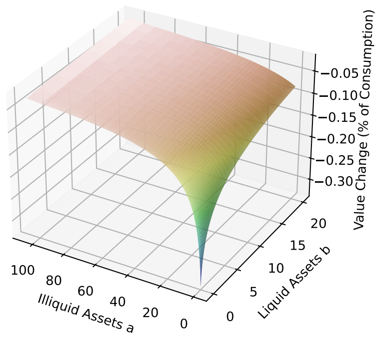

PANEL B: OIL SHOCK, LOW EARNINGS

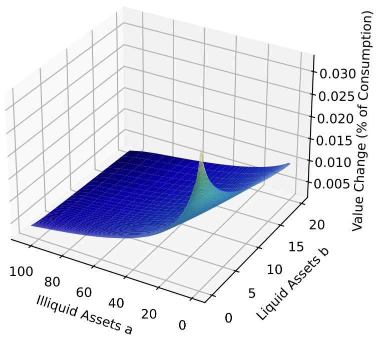

PANEL C: MONETARY SHOCK, MID EARNINGS

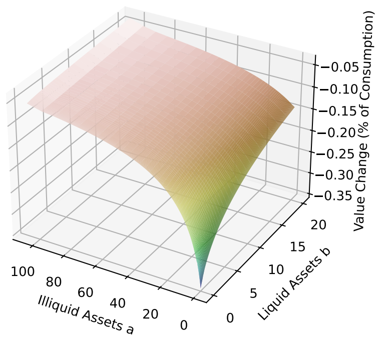

PANEL D: OIL SHOCK, MID EARNINGS

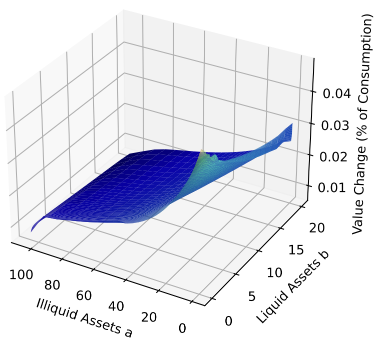

PANEL E: MONETARY SHOCK, HIGH EARNINGS

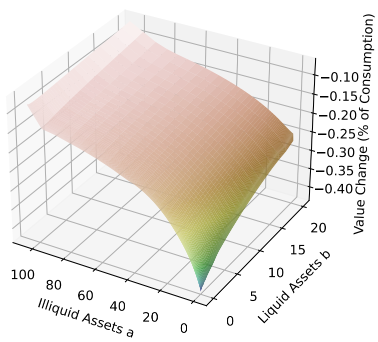

PANEL F: OIL SHOCK, HIGH EARNINGS

Notes: Figure plots model-implied value changes across initial point in the state space, calculated via backward induction, in response to a 25 basis point reduction in monetary policy rates (Panels A, C and E) or a 1% decline in aggregate TFP (Panels B, D and F). Shocks are assumed to decay according to an AR(1) process with persistence 0.4 (monetary shocks) or 0.9 (TFP shocks). Darker colors indicate more negative (or less positive) welfare changes. Panels A and B are plots for the lowest realization of idiosycnratic earnings, Panels C and D for the middle realization, whereas Panels E and F plot for the highest realization. All plots are based on our baseline calibration (e.g., log utility, calibrated idiosyncratic risk, tight borrowing constraints, etc.).

FIGURE E4: Regressivity of Inflationary Monetary and Oil Shocks in the Model and Data

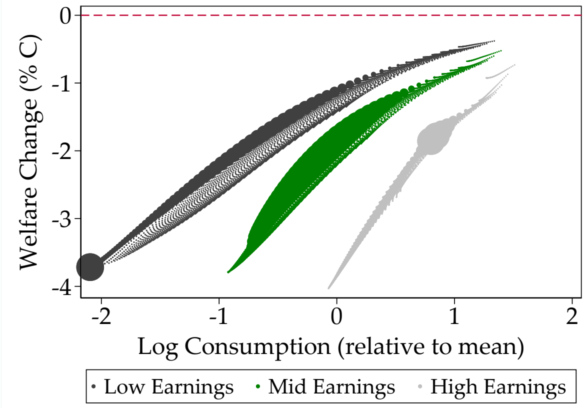

PANEL A: OIL SHOCK – MODEL

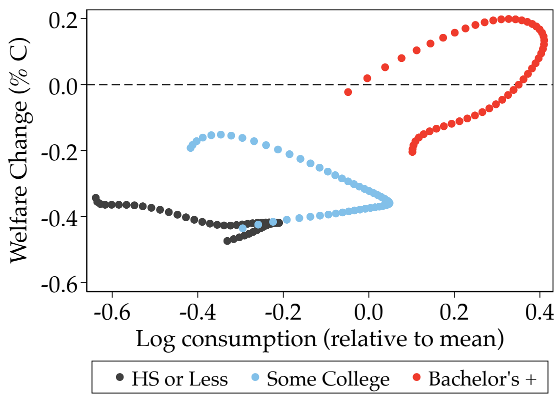

PANEL B: OIL SHOCK – DATA

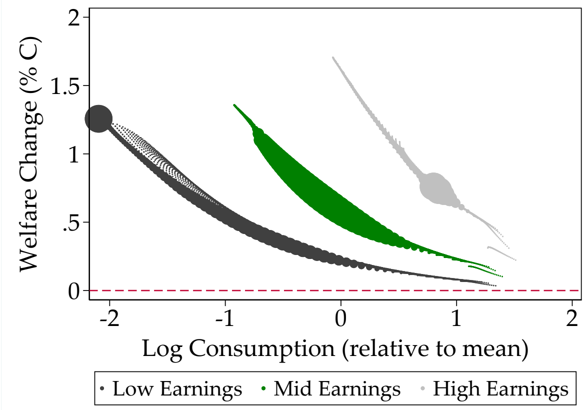

PANEL C: MONETARY SHOCK – MODEL

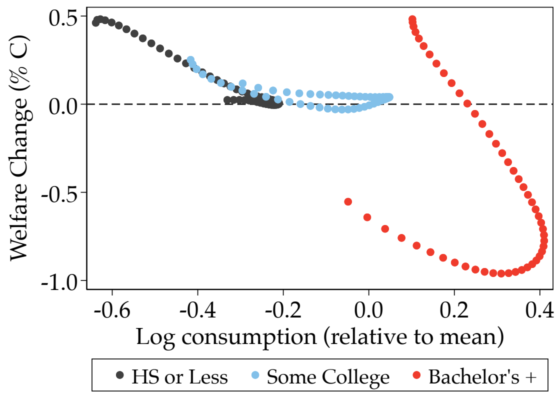

PANEL D: MONETARY SHOCK – DATA

Notes: Figure plots demeaned log consumption against welfare changes from oil shocks (Panels A and B) and monetary shocks (C and D) in the model (Panels A and C) and data (B and D). The data changes reflect those calculated in section 7.1: each dot is a different age $\times$ education group. The model plots reflect welfare changes implied by the two-asset HANK model after a sequence of monetary shocks equivalent to those we estimate in response to Gertler and Karadi (2015) shocks, or a sequence of TFP shocks that we estimate as the response to the Känzig (2021) oil shocks. Each dot in the model is a point in the state space, whose size is proportional to the point's mass in the steady-state distribution. Welfare changes are scaled relative to four-year consumption.

FIGURE F1: Impulse Responses for an Oil Price Shock

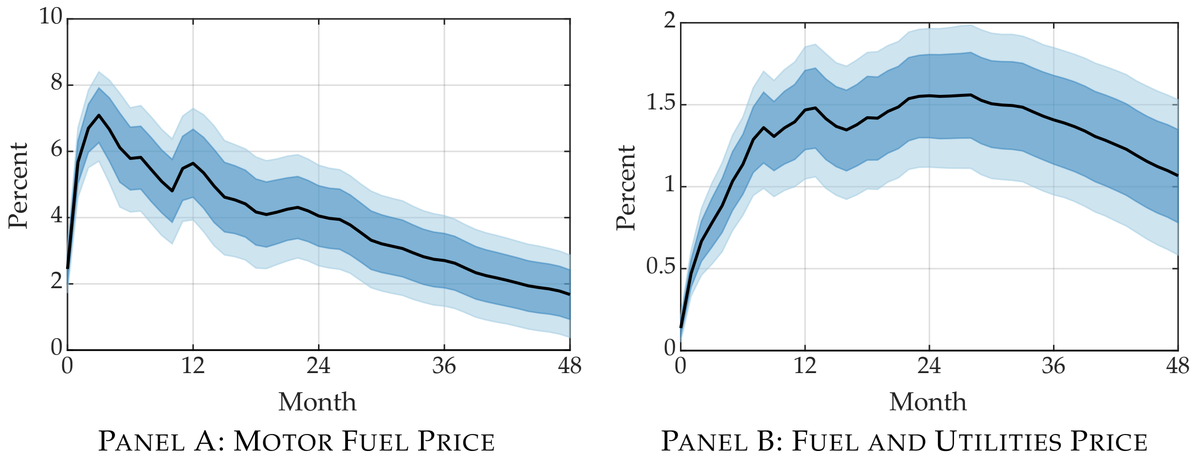

Notes: Figure plots cumulative IRFs to inflationary oil supply news shocks constructed by Känzig (2021). Shocks normalized to represent a 10% increase in the real West Texas Intermediates (WTI) crude oil price in high frequency windows around OPEC supply announcements. IRFs estimated using the "internal instrument" SVAR procedure explained in section 4. Panels A and B report the IRFs of the CPI categories for Motor Fuel and Fuel and Utilities, respectively. 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 and light blue regions represent 68% and 90% confidence intervals, respectively.

F. ADDITIONAL TABLES AND FIGURES

This section presents a number of additional results. Figure F1 reports the full IRF of motor fuel and fuel and utilities CPI prices in response to the Känzig (2021) oil shock. Figures F2 and F3 report the full impulse responses for stock prices, dividends, house prices, and corporate bond yields in response to inflationary oil shocks and monetary shocks, respectively.

Figure F4 reports the shares of consumption of motor fuel, fuels and utilities, and public transportation.

Figure F5 reports the life-cycle accumulation profiles of total net wealth (Panel A) and total equity, bonds, and housing (Panel B).

Figure F6 reports the transfer income channel of welfare losses from inflationary oil and monetary shocks over the life cycle for our three education groups.

Figures F7 and F8 report the total four-year money-metric welfare losses from inflationary oil and monetary shocks, respectively, for our three education groups over the life cycle. Negative numbers indicate welfare gains. These are analogous to Figures 7 and 8 in the main text, but are not scaled by four-year consumption.

Table F1 reports money-metric welfare changes in response to a 10% oil price shock identified following Känzig (2021) and a 25 basis point reduction in nominal interest rates following Gertler and Karadi (2015) for our three education groups and for three age groups. It further decomposes the welfare changes into the consumption, labor income, portfolio and transfers

FIGURE F2: Impulse Response of Asset Prices to an Inflationary Oil Shock

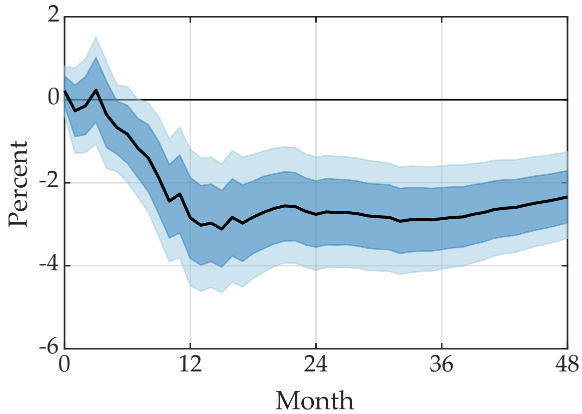

PANEL A: CUMULATIVE STOCK MARKET RETURN, EXCL. DIVIDENDS

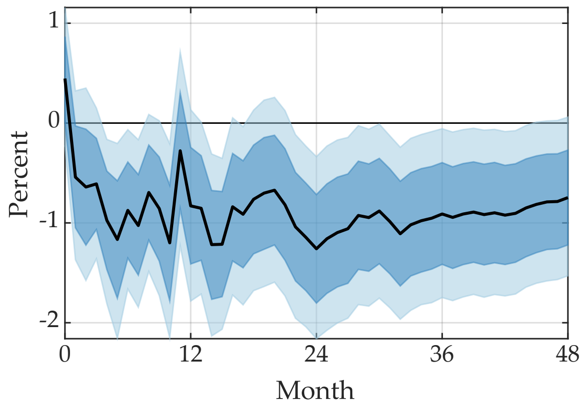

PANEL B: DIVIDEND PAYOUTS

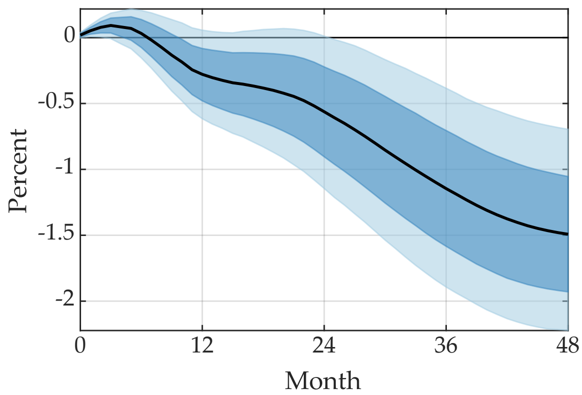

PANEL C: CASE-SHILLER HOME PRICE INDEX (HPI)

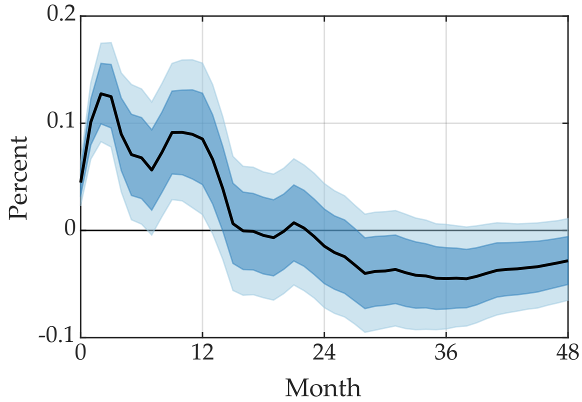

PANEL D: MOODY'S AAA CORPORATE BOND YIELDS

Notes: Figure plots 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 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. Dark and light blue regions represent 68% and 90% confidence intervals, respectively.

channels.

Figure F9 reports the p-value from a test that the welfare effects of oil (Panel A) and monetary (Panel B) shocks are the same for high school or less households and the more educated household groups. The p-value for the test for some college groups are given by the light blue line, while the red line plots the associated p-value for the bachelor's plus group. The vertical axis is limited to 0.3; therefore, p-values which rise above 0.3 are not shown.

Figure F10 plots estimated risk-adjustment factors Θ^x by age and education group. These are calculated as described in section 7.2.

Figure F11 reports inputs to the borrowing constraint analysis of section 7.3. Panel A reports

TABLE F1: Money-Metric Welfare Effects of Inflationary Oil Supply and Monetary Policy Shocks, by education and age group (2019 $)

	Oil Supply Shock					Monetary Policy Shock					
Household Group	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)	Total Consumption (11)
HS or Less											
22-35 y.o.	-$536	-$242	-$121	+$14	-$885	-$1 112	+$868	+$249	+$22	+$26	$198 706
35-50 y.o.	-$523	-$240	-$138	+$26	-$875	-$1 208	+$883	+$310	+$40	+$25	$208 893
51+ y.o.	-$362	-$69	-$397	+$163	-$665	-$1 031	+$277	+$912	+$252	+$410	$167 604
Some College											
22-35 y.o.	-$495	-$439	+$8	+$18	-$909	-$1 125	+$1 241	-$1	+$27	+$142	$226 395
35-50 y.o.	-$533	-$458	-$12	+$32	-$970	-$1 354	+$1 332	+$73	+$50	+$101	$267 012
51+ y.o.	-$341	-$135	-$247	+$191	-$533	-$1 314	+$432	+$659	+$296	+$72	$220 415
College+											
22-35 y.o.	-$445	-$440	+$1 256	+$5	+$376	-$1 467	+$1 422	-$2 495	+$8	-$2 532	$302 888
35-50 y.o.	-$392	-$461	+$1 463	+$13	+$622	-$2 099	+$1 513	-$2 691	+$20	-$3 258	$378 855
51+ y.o.	-$248	-$146	-$11	+$182	-$222	-$2 094	+$503	+$932	+$284	-$375	$336 445

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 F3: Impulse Response of Asset Prices to an Expansionary Monetary Policy Shock

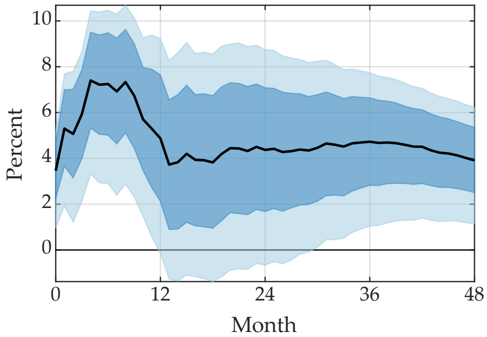

PANEL A: CUMULATIVE STOCK MARKET RETURN, EXCL. DIVIDENDS

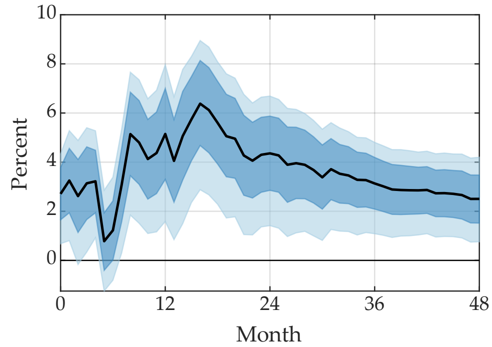

PANEL B: DIVIDEND PAYOUTS

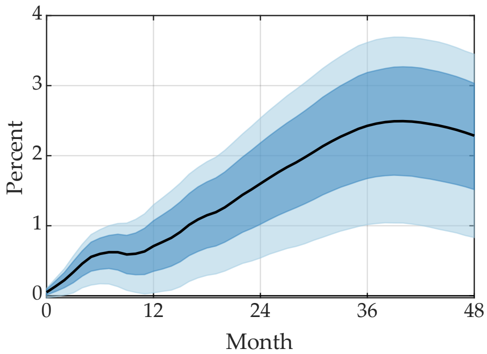

PANEL C: CASE-SHILLER HOME PRICE INDEX (HPI)

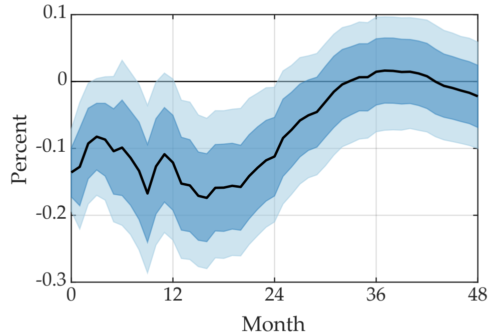

PANEL D: MOODY'S AAA CORPORATE BOND YIELDS

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 and light blue regions represent 68% and 90% confidence intervals, respectively.

the estimated Euler wedge constructed following equation (9). Panel B reports the share of households we classify as constrained in the 2019 SCF. Households are considered constrained if their net worth is non-positive. Wedges are larger for those with at least a bachelor's degree, who experience steeper lifetime consumption growth (see Figure 5) and do not perfectly smooth consumption.[70]

Figure F12 reports money-metric welfare losses up to age 80 from oil and monetary shocks, as a share of total consumption up to age 80. Whereas our baseline analysis calculates welfare losses over a four-year horizon, effectively assuming that IRFs fade to zero after four years, this exercise assumes that IRFs stay forever elevated at their level four years after the shock. For instance, stock prices remain permanently 4% higher after the monetary shock as seen

[70]These hump shaped consumption profiles could be generated by non-homothetic preferences, which we assume away here for measurement, and by child-rearing patterns.

FIGURE F4: Life-Cycle Consumption Shares

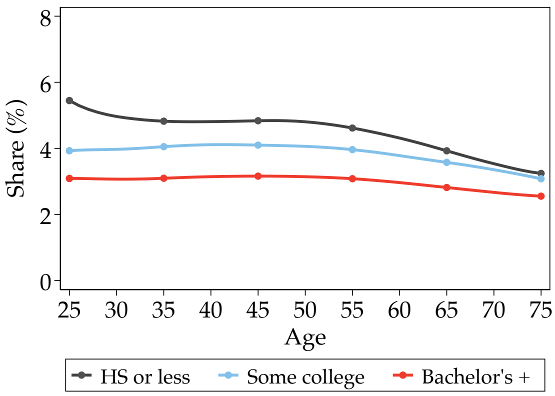

PANEL A: MOTOR FUEL

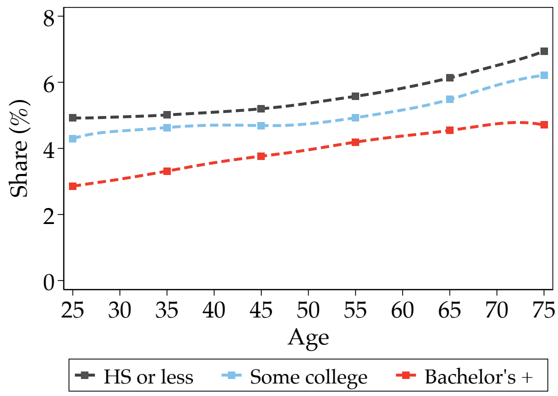

PANEL B: FUELS AND UTILITIES

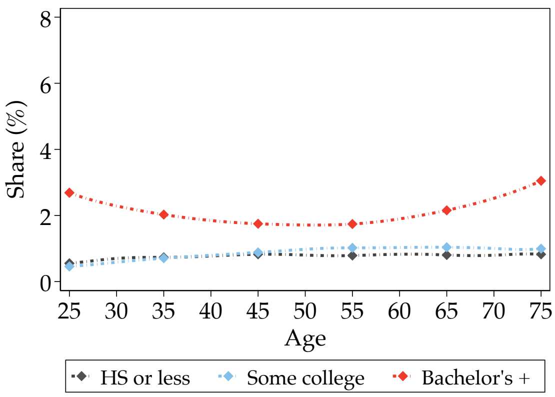

PANEL C: PUBLIC TRANSPORTATION

Notes: Figure reports shares of consumption of motor fuel (Panel A), fuels and utilities (Panel B), and public transportation (Panel C). All panels average expenditure shares by age, obtained from the CEX.

from Panel A of Figure F3. Likewise, we project consumption forward up to age 80 to put these welfare responses in similar units to those in the main text. We do not compute standard errors for these estimates, as there is no theoretically consistent way to account for estimation uncertainty for these long impulse responses which we simply assert. We find little difference between the welfare effects after four years and the welfare effects after thirty years, at least in consumption share space.

Figure F13 plots four-year money-metric welfare by income quintile. As usual, we scale by four-year consumption. To construct this plot, we first need to determine how to obtain income quintiles from all cross sectional data sources. After these are defined, welfare calculations follow the same procedure as with educational groups. We proceed as follows. All five surveys contain household income information and thus obtaining income quintiles is, in principle, straightforward. However, the purpose of educational groups (which we wish to maintain for income groups) was capturing *permanent* differences in life cycle earnings. Thus, given the inherent life cycle behavior of income, creating these groups naïvely is problematic.

FIGURE F5: Life-Cycle Accumulation of All Assets

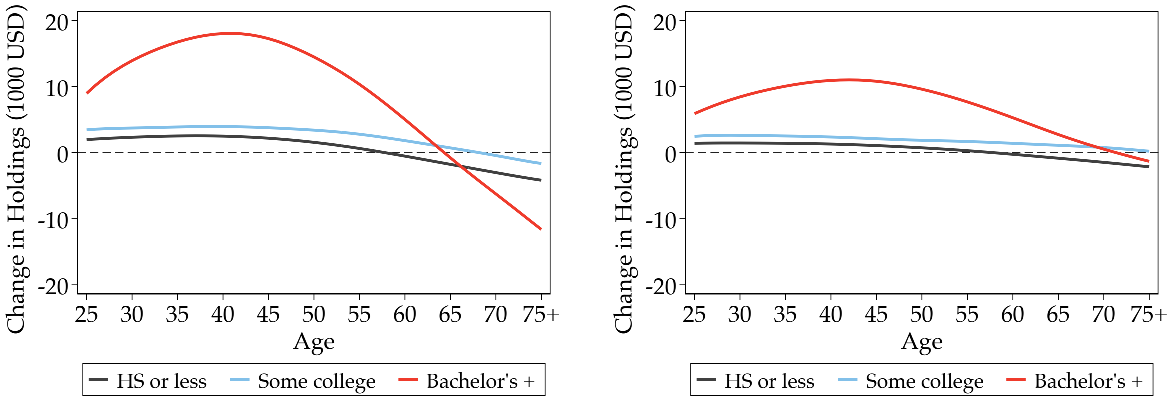

PANEL A: ALL ASSETS

PANEL B: EQUITY, BONDS, AND HOUSING

Notes: 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" includes equity, (corporate and non-corporate) bonds, housing, vehicles, liquid assets, business wealth, and other financial and non-financial assets.

Instead, for each survey we create quintiles of household income *within* each age.[71] Then, we assume the obtained groups correspond to permanent household groups, which we can use to compute welfare effects. Throughout, we calculate welfare following Proposition 1 and so do not account for borrowing constraints or risk in these figures.

Panels A and B present losses for oil shocks, while Panels C and D report losses for monetary shocks. Panels A and C report total welfare effects summing across all of the channels, while Panels B and D report welfare losses focusing on each channel in turn. The figure shows that the patterns for the top quintile of earners appears very similar to those for college-educated households in our baseline analysis. Similarly, lower earnings appear quite similar to less educated households.

We think this exercise is less compelling than our baseline analysis using education groups for a two reasons. First, many households do not have any labor income for a variety of reasons, such as unemployment, schooling, or retirement, which makes them hard to classify by income. Second, forecasting future consumption and asset holdings by current income is difficult, since current income may be a poor predictor of permanent income. Since education is a fixed characteristic of a household, our synthetic cohort approach to projecting choices is more likely to be valid when households are grouped by education. Nevertheless, it is heartening that the patterns by income are consistent with those we document by education.

Finally, Figure F14 shows the robustness of our results to using different waves of the SCF to

[71] All surveys allow us to compute these groups directly using household income, with the CPS being the only exception. The CPS has household income binned, which precludes using it directly to create quintiles, as some ages might not have 5 defined income groups. To overcome this hurdle, we add to the household income variable the weekly household income, divided by twice the maximum weekly income in the survey. This allows us to order households within the bins of the household income variable.

FIGURE F6: Government Transfer Channel of Oil Price and Monetary Shocks

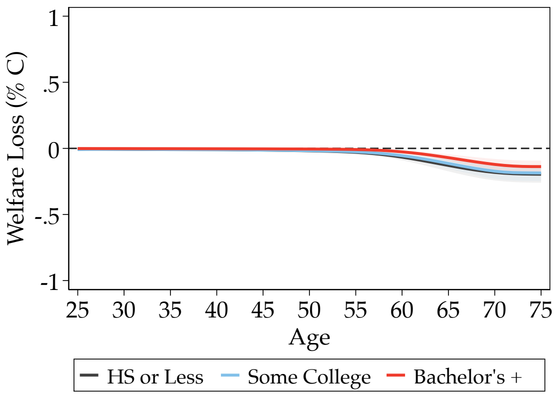

PANEL A: OIL SHOCKS

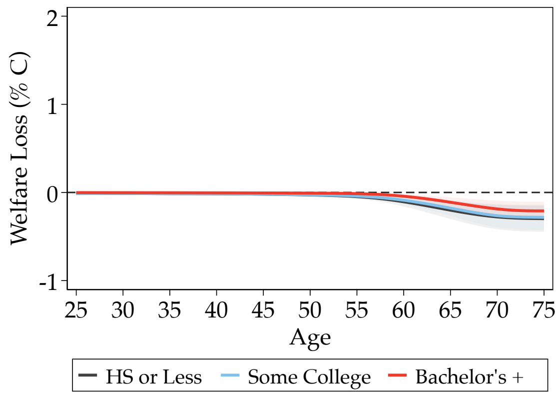

PANEL B: MONETARY SHOCKS

Notes: Figure reports the money-metric welfare loss arising from changes in government transfer income that result from inflationary oil price shocks (Panel A) or monetary shocks (Panel B). Vertical axis is normalized to be a share of consumption. Negative numbers represent welfare gains. Shaded regions represent 90% confidence bands.

construct the asset portfolios by age. Our baseline uses the 2019 SCF. We present estimates using 2016, 2013 and 2010. Figure F15 does the same for the monetary shocks.

G. THE DETERMINANTS OF RISK-ADJUSTMENT SHIFTERS Θ

Here we consider a simple parametric example to shed light on the determinants of the risk-adjustment shifters in Proposition 2.

Under complete markets, Θ tend to zero because the marginal utility of consumption is equated across states. Similarly, if the agent has sufficient savings in liquid assets, they become close to zero because the agent can (to a degree) self-insure through accumulating savings.

Let us consider the opposite polar case, where the agent is completely exposed to fluctuations in labor income risk and consumes hand-to-mouth. Moreover, suppose they supply one unit of their labor endowment in each state, so that across states $C_t(s_t) = W_t e_t^i(s_t)$, where we assume aggregate risk vanishes. What is Θ_t^W in this case?

If their utility gained from consumption takes the constant relative risk aversion form and is separable from preferences over leisure and asset holdings, then $U_C(s_t) = C(s_t)^{-\nu}$, where $\nu > 0$ governs the household's degree of risk aversion. As such,

$$\Theta_t^W = Cov\left(\frac{1/e_t^i(s_t)^\nu}{\mathbb{E}_0(1/e_t^i(s_t)^\nu)}, \frac{e_t^i(s_t)}{\mathbb{E}_0(e_t^i(s_t))}\right).$$

To examine the determinants of this, suppose $X = e_t^i(s_t) \sim LogNormal(\mu, \sigma)$. Then $Z = X^{-\nu} \sim$

FIGURE F7: Money-Metric Welfare Loss of Inflationary Oil Price Shocks (in 2019 $)

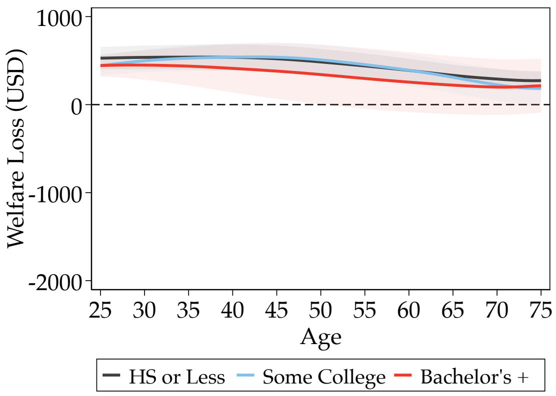

PANEL A: CONSUMPTION CHANNEL

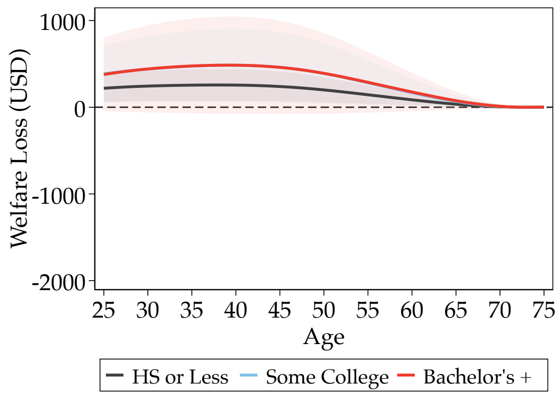

PANEL B: LABOR-INCOME CHANNEL

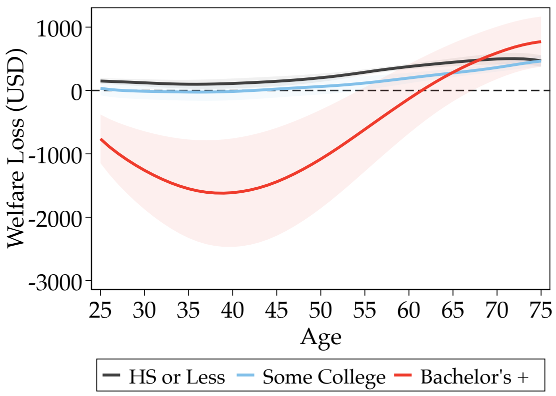

PANEL C: PORTFOLIO CHANNEL

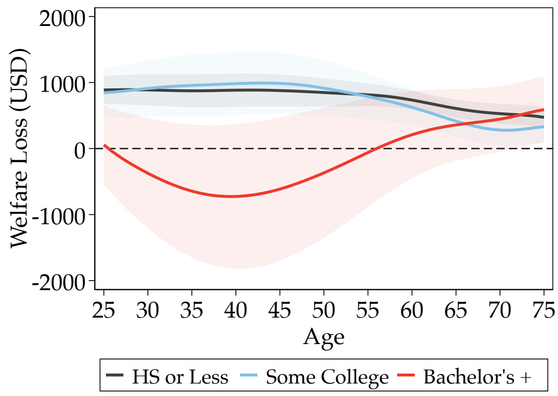

PANEL D: TOTAL WELFARE CHANGE

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, and the portfolio channel, respectively. A negative number represents a welfare gain. Shaded regions represent 90% confidence bands.

FIGURE F8: Money-Metric Welfare Loss of Inflationary Monetary Shocks (in 2019 $)

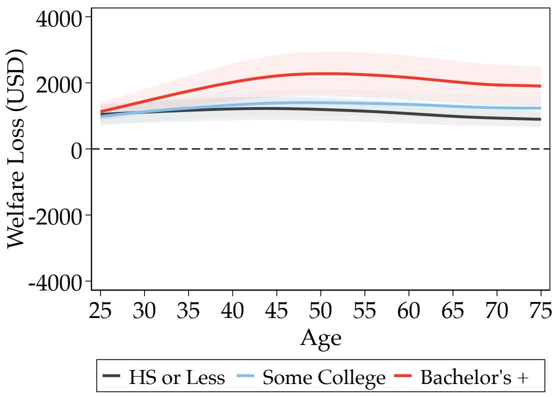

PANEL A: CONSUMPTION CHANNEL

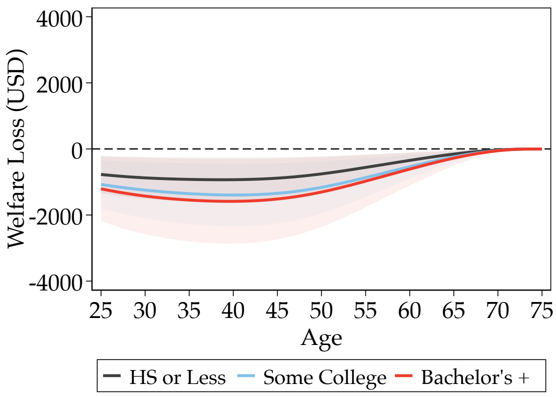

PANEL B: LABOR-INCOME CHANNEL

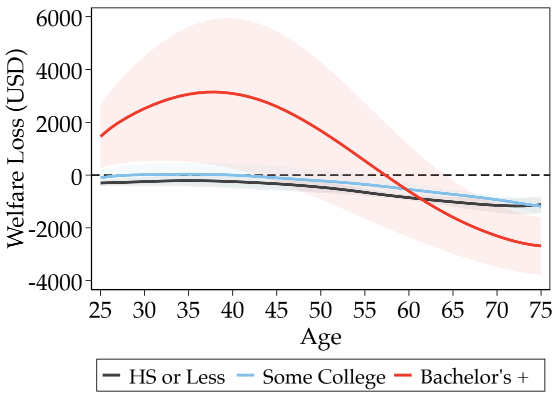

PANEL C: PORTFOLIO CHANNEL

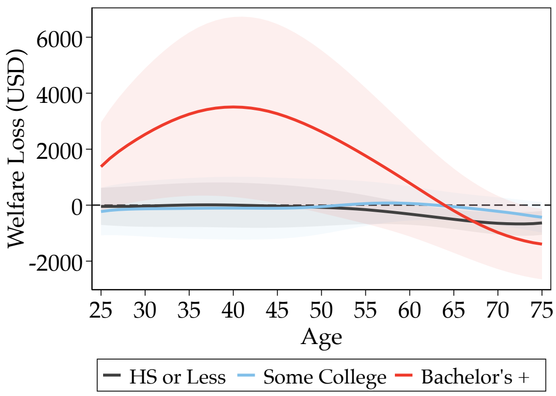

PANEL D: TOTAL WELFARE CHANGE

Notes: Figure shows the estimated money-metric welfare effects of a 25 basis point cut to the one-year Treasury yield. Panels A-C split the effect into the consumption channel, the labor-income channel. A negative number represents a welfare gain. Shaded regions represent 90% confidence bands.

FIGURE F9: Statistical Test for Different Welfare Losses Across Groups

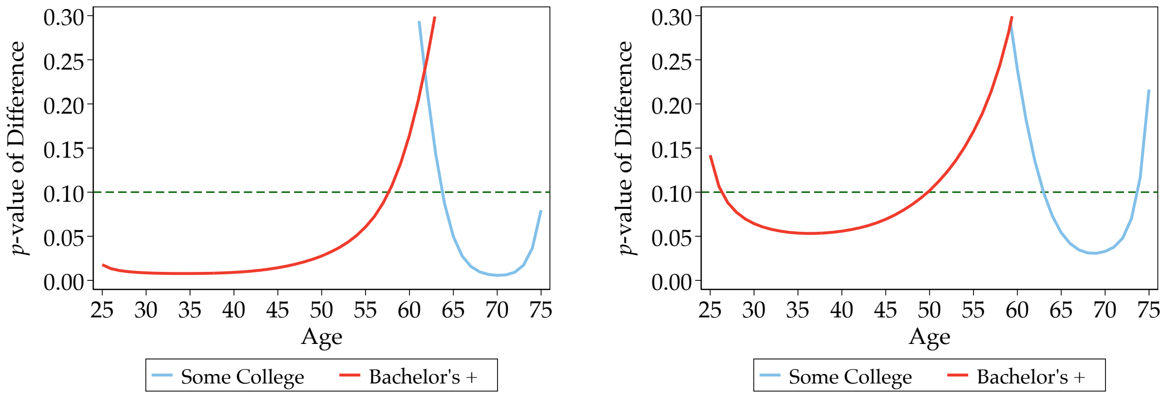

PANEL A: OIL SHOCK PANEL B: MONETARY SHOCK

Notes: The panels plot the p-values of the hypothesis test where the null hypothesis is that the pointwise difference in welfare between a given educational category (some college or Bachelor's +) and HS or less is zero for a particular age. These are computed using the Delta method, accounting for uncertainty in the estimation of the underlying time series impulse response functions. p-values exceeding the y-axis limit of 0.3 are not shown.

$LogNormal(-\nu\mu, \nu\sigma)$. Under this assumption,

$$\begin{aligned}\Theta_t^W = \frac{Cov(X,Z)}{\mathbb{E}_0(X)\mathbb{E}_0(Z)} &= \frac{\mathbb{E}_0(XZ)}{\mathbb{E}_0(X)\mathbb{E}_0(Z)} - 1 \\ &= \frac{\exp\left((1-\nu)\mu + \frac{(1-\nu)^2\sigma^2}{2}\right)}{\exp\left(\mu + \frac{\sigma^2}{2}\right)\exp\left(-\nu\mu + \frac{\nu^2\sigma^2}{2}\right)} - 1 \\ &= \exp\left(-\nu\sigma^2\right) - 1,\end{aligned}$$

where we have used the fact that $XZ = X^{1-\nu}$. Note this quantity is bounded between zero and -1. First note that as σ^2 goes to zero, the risk-adjustment term goes to zero as well. The same occurs as ν goes to zero. If the household is risk-neutral or if risk is small, there is no risk-adjustment term even if households are hand-to-mouth.

Now suppose the household's labor income becomes very idiosyncratically risky ($\sigma \to \infty$) or the agent grows infinitely risk-averse ($\nu \to \infty$), and Θ_t^W approaches -1. In this case, the labor income terms drop out of Proposition 2. The intuition is as follows. As the household grows extremely risk-averse and subject to large idiosyncratic income risk, small changes in *aggregate* wages have a very small effect on household well-being. In effect, idiosyncratic earnings risk swamps any shift in wages caused by a small aggregate shock, so small changes in the value of wage payments in every state is discounted more and more heavily.

More generally, there are three determinants of the idiosyncratic risk-adjustment factors. First, the larger the idiosyncratic risk, the larger in (absolute value) the risk-adjustment factors. Second, the more risk-averse are households, the larger these factors will be. Finally, the less able households are to insure against these idiosyncratic risks, e.g. due to incomplete markets, the larger these factors will be.

FIGURE F10: Estimated Θ^x by Age Group

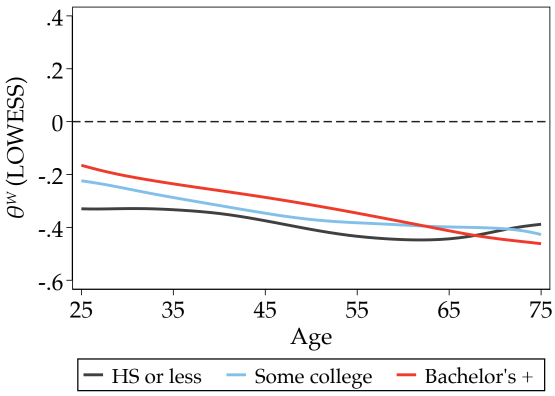

PANEL A: LABOR INCOME Θ_t^W

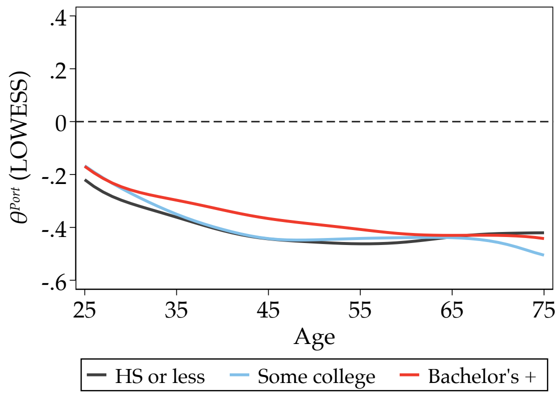

PANEL B: PORTFOLIO INCOME Θ_t^{Port}

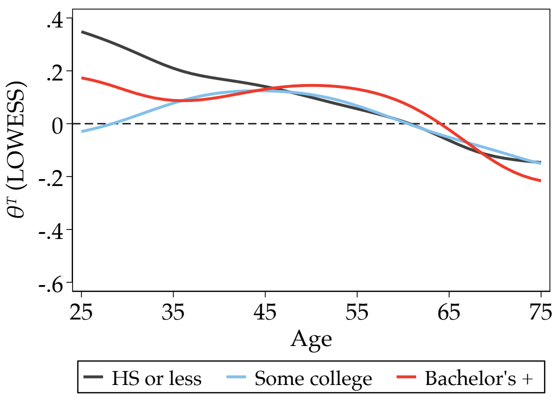

PANEL C: TRANSFER INCOME Θ_t^T

Notes: Figure reports estimated risk-adjustment factors Θ^x throughout the life cycle for our three education groups. These are estimated by taking the covariance of the demeaned reciprocal of consumption expenditures with demeaned labor income (Panel A), portfolio income (Panel B), and transfer income (Panel C), all in the CEX. Covariances calculated within 5-year age groups by education to maximize power, then LOWESS smoothed over the full life cycle.

FIGURE F11: Estimated Euler Equation Wedges and Constrained Household Shares

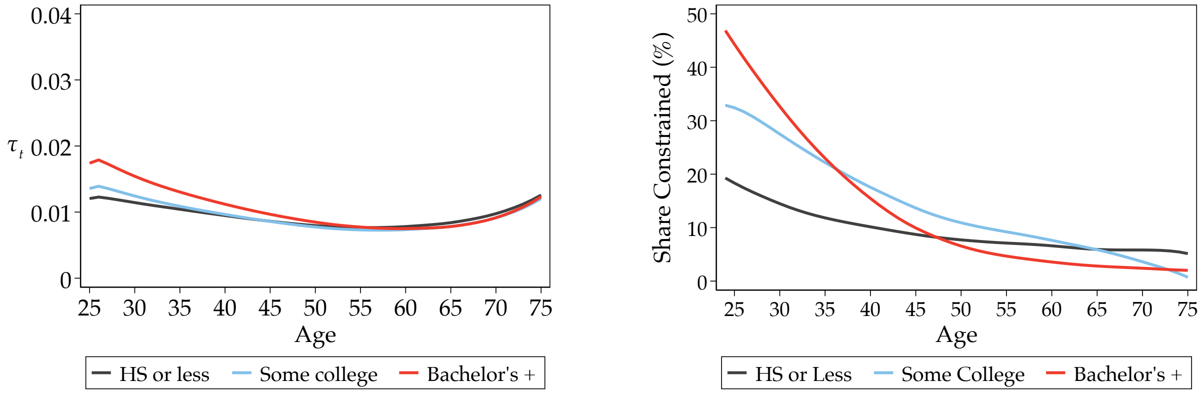

PANEL A: ESTIMATED WEDGES τ PANEL B: SHARE OF CONSTRAINED HOUSEHOLDS

Notes: Panel A plots the estimated τ from (9) in the CEX for our three groups of consumers by age, assuming log utility. Death rates are taken from the Period Life Table for 2019 from the Social Security Administration. Panel B shows the life cycle share of constrained households over the life cycle, after applying a LOWESS smoother individually for each educational group. A constrained household is defined as one whose net worth is non-positive. Data is from the Survey of Consumer Finances 2019.

G.1 Robustness to Greater Risk Aversion

As the above parametric example makes clear, greater risk aversion makes the risk-adjustment shifters larger in absolute value. For our main empirical results, we assumed log separability between consumption and the other components of utility. Here we explore how adding more curvature to the marginal utility of consumption changes the empirically estimated Θ's.

Specifically, let us now assume

$$U(C_t, L_t, \{N_{kt}\}) = \frac{C_t^{1-\nu} - 1}{1-\nu} + h(L_t, \{N_{kt}\}),$$

where the consumption aggregator over the j goods is again homothetic within group and age cells. We can then write

$$\begin{aligned} \Theta_t^W &\equiv Cov\left(\frac{U_C(s_t)/P_t(s_t)}{\mathbb{E}_0[U_C/P_t]}, \frac{W_t(s_t)e(s_t)L_t(s_t)}{\mathbb{E}_0[W_t e_t L_t]}\right) \\ &= Cov\left(\frac{C_t(s_t)^{-\nu}}{\mathbb{E}_0[C_t^{-\nu}]}, \frac{W_t(s_t)e(s_t)L_t(s_t)}{\mathbb{E}_0[W_t e_t L_t]}\right) \\ &= Cov\left(\frac{(P_t C_t(s_t))^{-\nu}}{\mathbb{E}_0[(P_t C_t)^{-\nu}]}, \frac{W_t(s_t)e(s_t)L_t(s_t)}{\mathbb{E}_0[W_t e_t L_t]}\right), \end{aligned}$$

where the second and third equalities follow via homotheticity.

Figure G1 shows what happens to the estimated Θ^W's when we move from log separability to a coefficient of relative risk aversion of 2. This makes the risk-adjustment shifters substantially more negative and increases them by around 50% in absolute value. However, the relative

FIGURE F12: Estimated Lifetime Welfare Loss from Inflationary Monetary and Oil Price Shocks

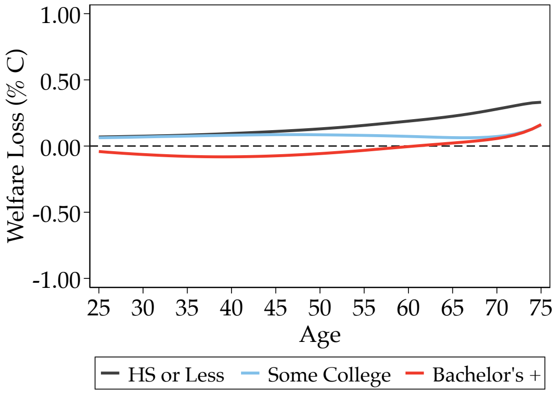

PANEL A: OIL SHOCK – TOTAL

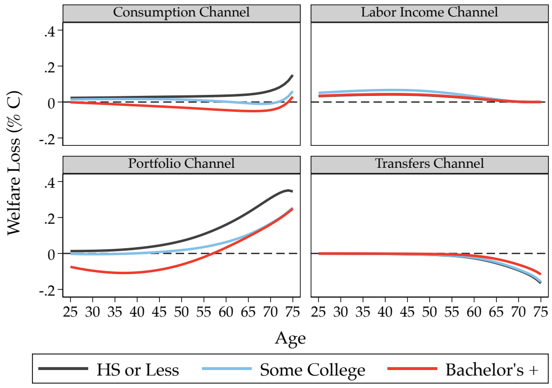

PANEL B: OIL SHOCK – CHANNELS

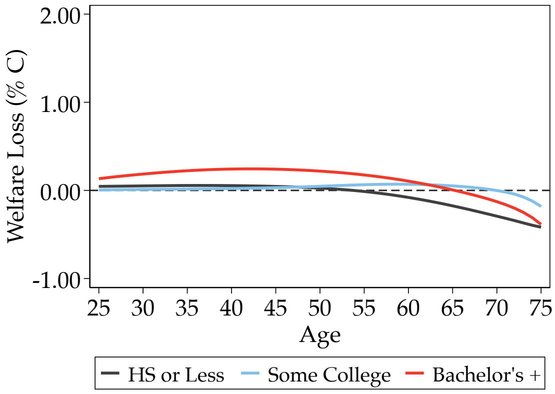

PANEL C: MONETARY SHOCK – TOTAL

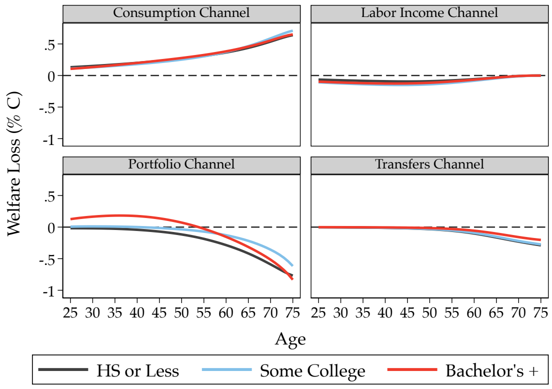

PANEL D: MONETARY SHOCK – CHANNELS

Notes: Figure plots the money-metric welfare losses up to age 80 from oil (Panels A and B) and monetary shocks (Panels C and D). Panels A and C report total welfare losses, while Panels B and D report losses arising from our four channels. Impulse responses are assumed to be constant following four-years. Money-metric welfare losses are scaled by total consumption up to age 80.

FIGURE F13: Estimated Welfare Losses By Income Group

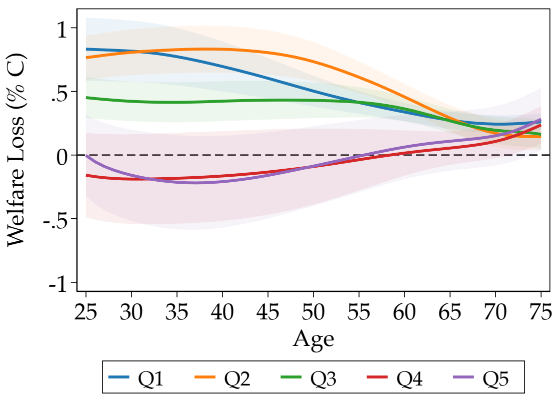

PANEL A: OIL SHOCK – TOTAL

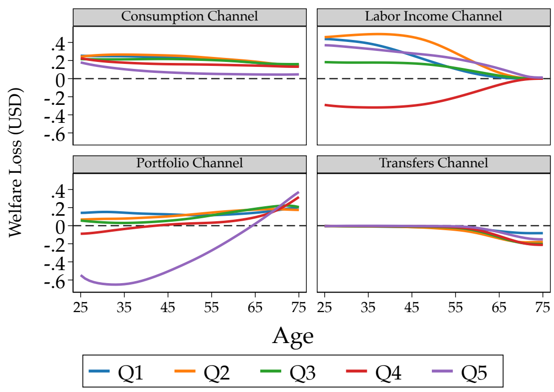

PANEL B: OIL SHOCK – CHANNELS

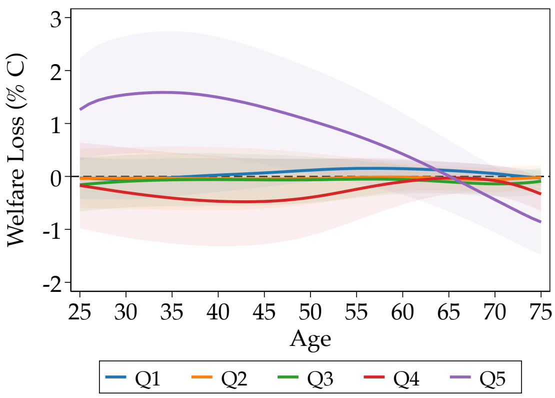

PANEL C: MONETARY SHOCK – TOTAL

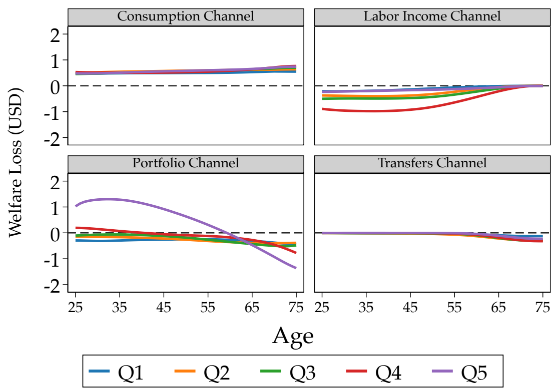

PANEL D: MONETARY SHOCK – CHANNELS

Notes: Figure plots four-year welfare losses, as a share of four-year consumption, by income quintile in 2019. Each line corresponds to a different income level. Panels A and C report total welfare losses of oil and monetary shocks, respectively, while Panels B and D do the same broken down by channel. Negative numbers indicate welfare gains. Shaded regions represent 90% confidence bands.

FIGURE F14: Robustness to Using Other SCF Waves: Oil Shocks

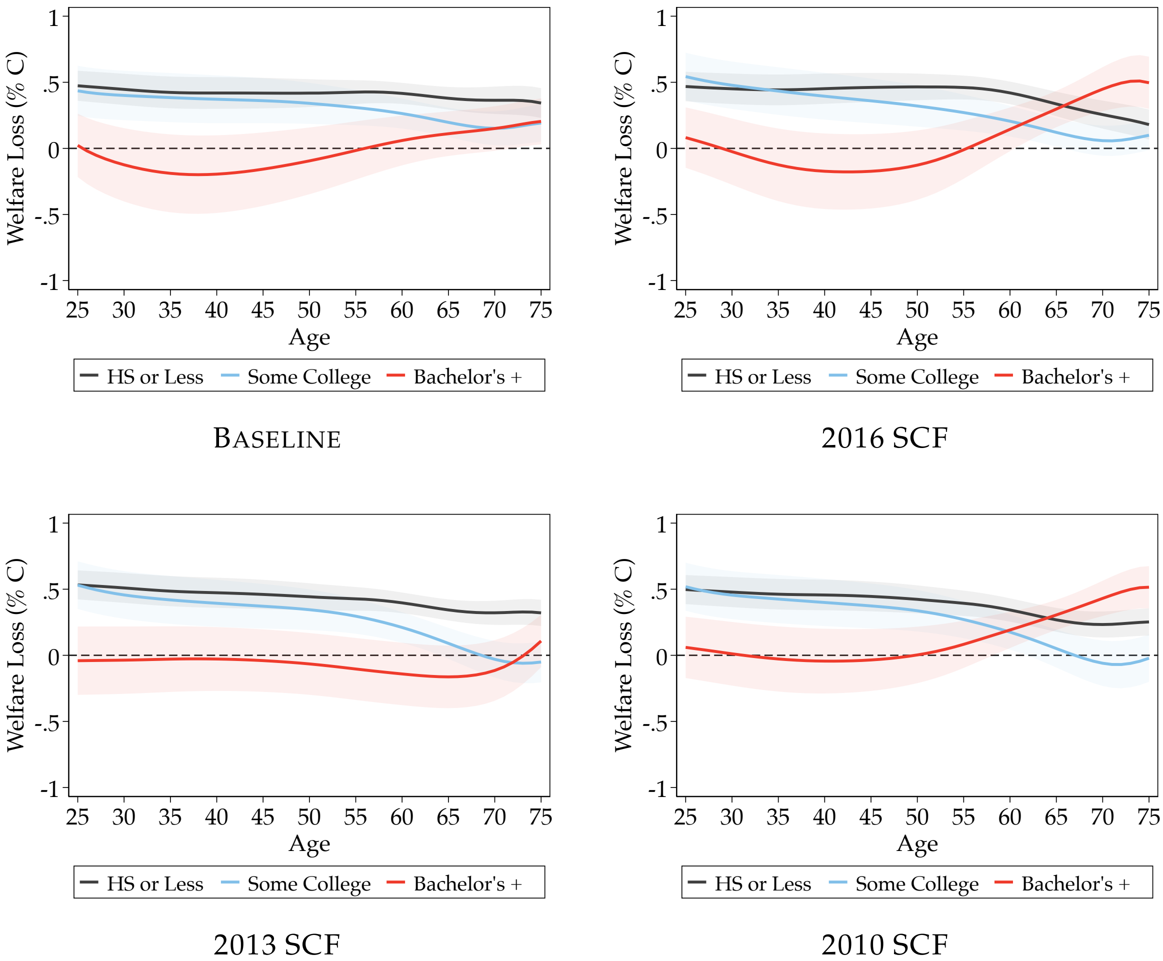

Notes: Figure shows the estimated money-metric welfare effects of a 10% increase in the West Texas Intermediates Crude Oil price driven by announced reductions in OPEC oil supply using different waves of the SCF to construct the baseline asset accumulation profiles. Negative numbers indicate welfare gains. Shaded regions represent 90% confidence bands.

patterns across age and education bins are mostly unchanged. As a result, plugging these new shifters into the formula of Proposition 2 mainly serves to scale down the welfare estimates proportionally, but leave the relative statements about regressivity unchanged.

FIGURE F15: Robustness to Using Other SCF Waves: Monetary Shocks

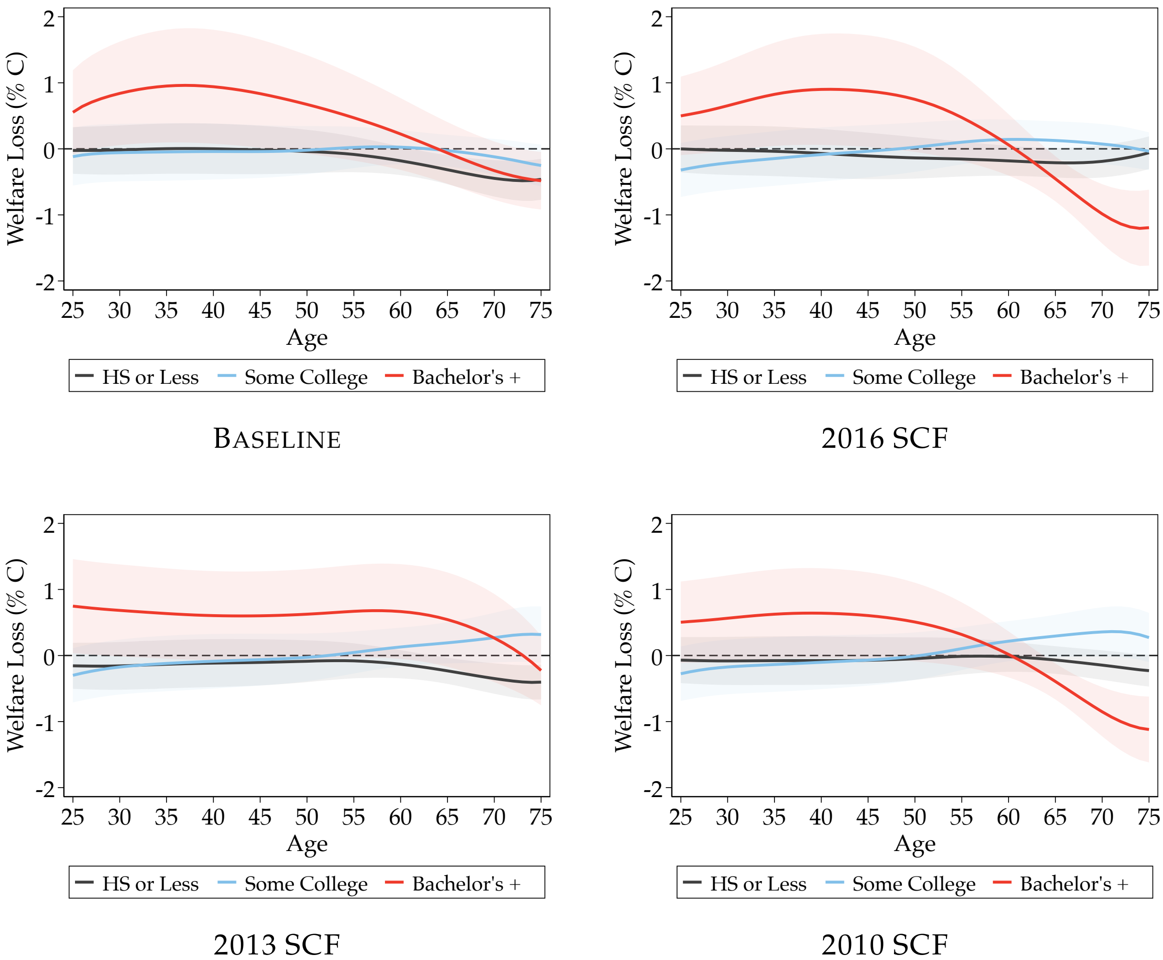

Notes: Figure shows the estimated money-metric welfare effects of a 25 basis point cut to the one-year Treasury yield using different waves of the SCF to construct the baseline asset accumulation profiles. Negative numbers indicate welfare gains. Shaded regions represent 90% confidence bands.

FIGURE G1: Risk-Adjustment Shifters for Labor Income Under Larger Risk Aversion

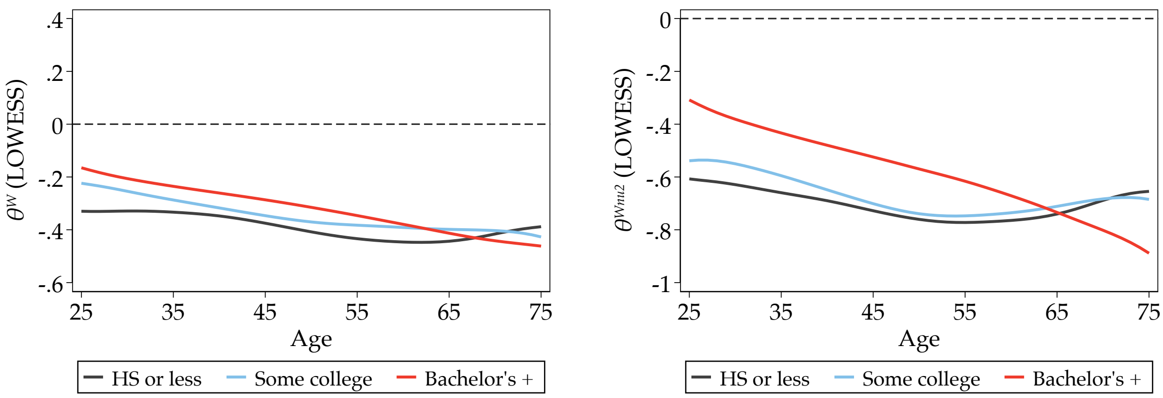

PANEL A: CRRA COEFF. $\nu = 1$ PANEL B: CRRA COEFF. $\nu = 2$

Notes: Figure reports estimated risk-adjustment factors Θ^x throughout the life cycle for our three education groups. These factors are estimated by taking the covariance of the demeaned reciprocal of consumption expenditures with demeaned labor income in the CEX. Covariances calculated within five-year age groups by education to maximize power, then LOWESS smoothed over the full life cycle. Panel A sets the risk aversion coefficient ν to 1, and Panel B sets it at 2.