IMPERFECT RISK SHARING AND THE BUSINESS CYCLE*

David Berger
Luigi Bocola
Alessandro Dovis

This article studies the macroeconomic implications of imperfect risk sharing implied by a class of New Keynesian models with heterogeneous agents. The models in this class can be equivalently represented as a representative-agent economy with wedges. These wedges are functions of households' consumption shares and relative wages, and they identify the key cross-sectional moments that govern the impact of households' heterogeneity on aggregate variables. We measure the wedges using U.S. household-level data and combine them with a representative-agent economy to perform counterfactuals. We find that deviations from perfect risk sharing implied by this class of models account for only 7% of output volatility on average but can have sizable output effects when nominal interest rates reach their lower bound. *JEL Codes:* E32, E44.

I. Introduction

To what extent are households' heterogeneity and deviations from perfect risk sharing important for aggregate fluctuations? Building on the influential quasi-aggregation result in Krusell and Smith (1998), for a long time the consensus view in macroeconomics was that these features were not critical drivers of the business cycle. With the Great Recession, however, several economists suggested that shocks and frictions at the household level, such as tighter credit constraints or heightened income risk, could foster households' precautionary savings and explain the observed persistent slump in aggregate demand. Despite recent advances in understanding how these economic forces play out in general-equilibrium models, quantifying their role is still an open question. This is because the impact of these microeconomic frictions on macroeconomic aggregates depends

*We have benefited from the detailed comments of Robert Barro, Pol Antràs, four anonymous referees, and Fernando Alvarez, Adrien Auclert, Anmol Bhandari, Loukas Karabarbounis, Greg Kaplan, Patrick Kehoe, Dirk Krueger, Guido Lorenzoni, Alisdair McKay, Kurt Mitman, Ben Moll, Luigi Pistaferri, Richard Rogerson, Joe Vavra, and participants at several seminars and conferences. We thank Aniket Baksy and Sara Casella for excellent research assistance. Luigi Bocola thanks the Alfred P. Sloan foundation for generous funding.

The Quarterly Journal of Economics (2023), 1765–1815. https://doi.org/10.1093/qje/qjad013.
Advance Access publication on March 17, 2023.

on modeling choices that are hard to discipline empirically, like the assumed set of risk-sharing mechanisms available to households, the nature of their idiosyncratic risk, and the timing and distribution of fiscal transfers.[1]

In this article, we propose a method to quantify the importance of imperfect risk sharing for the business cycle and help researchers discipline these modeling choices. We extend a result by Nakajima (2005) and show that the aggregate implications of a large class of economies with incomplete markets can be equivalently studied in a representative-agent economy with wedges. We measure these wedges using micro data and combine them with the representative-agent model to evaluate the contribution of imperfect risk sharing for the U.S. business cycle. We find that these frictions account for only 7% of output volatility on average, but they can have much larger effects when nominal interest rates are at the zero lower bound. Indeed, we find that these frictions were key determinants of the depth and persistence of the Great Recession.

We apply our methodology to a class of New Keynesian models with heterogeneous agents. To keep the analysis transparent, the "macro block"—the details of production, nominal rigidities, the conduct of monetary policy, and so on—of these economies is the same as in the standard three-equations model. The households' decision problem, or "micro block," is modeled in a more flexible way: households face idiosyncratic risk, and their ability to smooth these shocks depends on the risk-sharing mechanisms available to them, for example, the set of financial assets they can trade, their financial constraints, and the presence of redistributive fiscal transfers. We formulate these features so as to nest most of the specifications considered in the literature including incomplete-market models in the Bewley-Huggett-Aiyagari tradition and those with endogenous debt limits as in Kehoe and Levine (1993) and Alvarez and Jermann (2000).

1. A prominent example is Kaplan and Violante (2014), which shows that the consumption response to fiscal transfers is very different if households can trade one liquid asset or one liquid and one illiquid asset. Other modeling choices, which are inconsequential in a representative-agent economy, matter in heterogeneous-agent economies, such as the timing and distribution of the fiscal transfers (Kaplan, Moll, and Violante 2018), how profits get distributed across households (Broer et al. 2020), and the cyclicality of idiosyncratic risk and access to liquidity (Werning 2015).

Our analysis builds on an equivalence result between these economies with heterogeneous agents and the canonical representative-agent New Keynesian model. For any heterogeneous-agent economy in our class, the macroeconomic aggregates—output, inflation, and nominal interest rates—satisfy the equilibrium conditions of a representative-agent economy with two wedges: one affecting the discount factor of the stand-in household and one affecting her labor supply. We call this the RA representation.

The discount factor wedge captures the failure of aggregation when consumption risk is not perfectly shared across households. To explain why the RA representation features this wedge, consider a household that is on her Euler equation in a heterogeneous-agent economy. With isoelastic preferences and complete markets, her consumption is a constant fraction of aggregate consumption at every point in time. Thus, the Euler equation also holds for aggregate consumption, and there is no need for a discount factor wedge in the RA representation. If risk sharing is not perfect, instead, consumption inequality can vary over time and households' consumption is no longer proportional to that of the aggregate. Therefore, the Euler equation evaluated using aggregate consumption does not typically hold, and the RA representation features a discount factor wedge.

The RA representation also features a wedge in the labor supply condition because of compositional changes in hours worked. To understand this result, consider an economy in our class and suppose that there is an increase in the cross-sectional dispersion of labor productivity. Because of substitution effects, high-productivity households will work more and low-productivity households will work less, a change in the composition of the labor force that leads to an increase in aggregate hours when measured in efficiency units. These compositional changes in worked hours that occur in the heterogeneous-agent economy are captured in the RA representation by a wedge in the labor supply condition.

We derive expressions for these wedges as functions of households' consumption shares and relative wages and emphasize two key properties. First, the mapping between the wedges and the micro allocation is the same for every model in our class. This implies that we do not need to take a stand on the details of the micro block for measuring the wedges, as long as we have observations on households' consumption and wages. Second, the wedges are a sufficient statistic for how households' heterogeneity affects aggregate variables, in the sense that shocks and frictions at the micro level matter for the aggregate if and only if they generate time variation in these two statistics. This makes the wedges ideal empirical targets for the analysis of incomplete-market economies.

We use panel data from the Consumer Expenditure Survey to measure the two wedges for the U.S. economy over 1992–2017. The labor supply wedge does not display much variation at business cycle frequencies, and it mostly reflects the secular increase in labor income inequality over this time. The discount factor wedge is, instead, countercyclical, and it displays a persistent increase after the Great Recession.

We then turn to study the aggregate implications of these movements using the RA representation. It is well known that an increase in the discount factor can induce sizable output drops in representative-agent New Keynesian models, especially when the zero lower bound constraint binds (Christiano, Eichenbaum, and Rebelo 2011).[2] Indeed, these models typically need large increases in the discount factor to explain episodes of low interest rates, inflation, and output, for example, the U.S. Great Recession. Frictions impeding risk sharing may be the root cause of these dynamics, and the increase in the discount factor wedge that we document provides qualitatively some support to this view. An important question is whether these movements are large enough to be quantitatively important.

To address this question, we use the realization of the wedges and data on output, inflation, and nominal interest rates to jointly estimate the structural parameters of the RA representation and the stochastic process of the wedges. Given the estimated parameters, we use the RA representation to construct the counterfactual path for aggregate variables in an economy with complete financial markets—that is, an economy with time-invariant consumption shares. We show that the presence of complete financial markets reduces the standard deviation of output by only 7%, suggesting that deviations from perfect risk sharing contribute little on average to business cycle fluctuations.

2. This is especially true for New Keynesian models that, unlike the one we study here, feature capital accumulation. Away from the zero lower bound, an increase in the discount factor typically generates a comovement problem between consumption and investment, as first suggested by Barro and King (1984) for neoclassical models. At the zero lower bound, this does not happen because the higher discount factor can lead to higher real interest rates.

This is because under the estimated monetary policy rule, increases in the discount factor wedge are offset by a reduction in nominal interest rates, with little effects on aggregate demand.

To further explore the role of the policy rule, we perform the counterfactual during the Great Recession, an episode where the monetary authority was constrained by the zero lower bound. We find that imperfect risk sharing explains one-third of the observed output losses and helps account for the slow recovery. This result underscores the importance of accounting for constraints on monetary policy when evaluating the aggregate implications of imperfect risk sharing.[3]

We conclude by studying what feature of the micro data is responsible for the rise in the discount factor wedge during the Great Recession and what this tells us about the economic mechanisms behind our results. In our application, the discount factor wedge is the sample average of the inverse change in the consumption shares for a group of households that we identify as financially unconstrained. We show that the increase in this statistic is driven mostly by an increase in the dispersion of consumption growth for these households. This pattern cannot be rationalized by models with simple forms of heterogeneity, such as the "two-agent" New Keynesian model studied in Galí, López-Salido, and Vallés (2007) and Bilbiie (2008). However, it is consistent with models in the Bewley-Huggett-Aiyagari tradition, specifically those that have emphasized heightened consumption risk and precautionary savings as a key driver of the fall in aggregate consumption during the Great Recession.

The economic literature has proposed two main mechanisms that can explain the increase in consumption risk of unconstrained households: higher volatility of their labor income, as in Bayer et al. (2019), versus a deterioration of the risk-sharing mechanisms available to them, such as a tightening of credit constraints, as in Guerrieri and Lorenzoni (2017). To explore which of these two mechanisms better accounts for the observed patterns, we study the relation between consumption and income during the Great Recession. We find that the data favor the second explanation, as the documented increase in the dispersion of consumption growth for unconstrained households is mostly due to higher sensitivity of consumption to income rather than an increase in the dispersion of the latter.

3. Schaab (2020) discusses this point in a heterogeneous-agent New Keynesian model with a zero lower bound constraint.

I.A. Related Literature

Our research contributes to a growing literature that introduces heterogeneous agents and incomplete financial markets in New Keynesian models. Thanks to recent computational advances, researchers have used these environments to study how frictions impeding risk sharing across households affect the transmission mechanism of monetary and fiscal policy and more generally the business cycle.[4] Closely related to our work are the papers of Bayer, Born, and Luetticke (2020), Auclert, Rognlie, and Straub (2020), and Bilbiie, Primiceri, and Tambalotti (2022), who estimate medium-sized New Keynesian models with heterogeneous agents. These are fully structural models and can be used to perform a variety of counterfactuals. However, they require the researchers to specify the details of the micro block—such as the set of financial assets available to households, their borrowing constraints, and the risk they face.

The main contribution of this study is to recognize that we can be agnostic about these details of the micro block when performing some of these counterfactuals, as long as we observe households' consumption choices. Specifically, our methodology allows one to assess the business cycle effects of imperfect risk sharing. We think that this is important for two reasons. First, our approach is less subject to misspecifications of the micro block; by nesting most of the frameworks studied in this literature, it provides a benchmark calculation. Second, we identify two sufficient statistics for the aggregate implications of households' heterogeneity in a broad class of models, and we show how to measure them using panel data. In this sense, our contribution is complementary to structural modeling because researchers could fruitfully use these wedges as a calibration target to discipline their models. We believe this is important because, unlike structural modeling,

4. See Gornemann, Kuester, and Nakajima (2016); McKay, Nakamura, and Steinsson (2016), McKay and Reis (2016), Auclert et al. (2017), Kaplan, Moll, and Violante (2018), Bhandari et al. (2018), and Hagedorn, Manovskii, and Mitman (2019) for recent contributions on monetary and fiscal policy. Guerrieri and Lorenzoni (2017) and Jones, Midrigan, and Philippon (2018) study how credit constraints can affect aggregate demand, while Challe et al. (2017), Den Haan, Rendahl, and Riegler (2018), McKay (2017), Ravn and Sterk (2017), Heathcote and Perri (2018), and Bayer et al. (2019) focus on idiosyncratic income risk. See Krueger, Mitman, and Perri (2016) and Kaplan and Violante (2018) for literature reviews.

our approach cannot be used for welfare assessment or policy analysis.[5]

The counterfactuals that we perform are related to the business cycle accounting methodology of Chari, Kehoe, and McGrattan (2007). There are two main differences between these procedures. First, in our approach the wedges are measured using household-level observations, rather than being chosen to replicate the observed path of aggregate data. By doing so, we are able to isolate the wedges due to imperfect risk sharing.[6] Second, our main quantitative experiment constructs the path for macroeconomic variables in a counterfactual economy with complete financial markets. This is not equivalent to the approach of Chari, Kehoe, and McGrattan (2007), which assesses the effects of specific wedges on the business cycle.

The idea that heterogeneous-agent economies admit an RA representation with wedges was developed by Maliar and Maliar (2001, 2003) for complete-markets economies following the insight of Constantinides (1982), and by Nakajima (2005) and more recently by Werning (2015) and Debortoli and Galí (2017, 2022) for economies with incomplete markets. In general, these time-varying wedges are endogenous equilibrium objects in the heterogeneous-agent economy under consideration, although recent papers derive the mapping between primitives and the wedges in some specific economies (Werning 2015; Acharya and Dogra 2018). Krueger and Lustig (2010) provide conditions under which the discount factor wedge is constant over time and is therefore irrelevant for aggregate fluctuations. This article is the first to exploit this representation to quantify the macroeconomic implications of imperfect risk sharing.

Finally, there is a connection between our article and the literature that evaluates the asset-pricing implications of models with heterogeneous households. See Brav, Constantinides, and Geczy (2002), Vissing-Jørgensen (2002), Krueger, Lustig, and Perri (2008), and Kocherlakota and Pistaferri (2009). The goal of

5. The latter is due to the endogeneity of the wedges to the policy regime. In this sense, our methodology is valid only for ex post evaluations similar to Arkolakis, Costinot, and Rodríguez-Clare (2012) in the international trade literature.

6. For example, if we were to only use aggregate data to measure the wedge in the Euler equation of the RA representation we would not be able to distinguish the discount factor wedge due to imperfect risk sharing from any other friction that would take the form of an Euler equation wedge. By using micro data, we can separately identify the former.

these papers is to estimate the stochastic discount factor with micro data given a particular form of market incompleteness. This is similar to the construction of the discount factor wedge in our approach. Clearly the scope of our analysis differs from these papers.

The article is organized as follows. Section II introduces the class of heterogeneous-agent economies at the center of our application. Section III derives the RA representation and discusses why this representation is a useful tool for the evaluation of heterogeneous-agent models. Section IV discusses how we estimate the preference wedges using panel data and approximate their stochastic process. In Section V we measure the preference wedges and combine these series with the RA representation to measure the role of imperfect risk sharing for the U.S. business cycle. We finish this section by discussing which models are most consistent with the patterns we identify in the micro data. Section VI concludes.

II. New Keynesian Models with Heterogeneous Agents

We introduce a class of New Keynesian models with heterogeneous households. The models in this class share the same specification for households' preferences, technology, market structure, and the conduct of monetary policy—elements that are borrowed from the prototypical "three equations" New Keynesian framework. However, they can differ in the details of the households' decision problem, such as the cyclicality of idiosyncratic risk faced by households, the set of assets they can trade, their financial constraints, as well as the timing and distribution of fiscal transfers.

1. *Environment.* Time is discrete and indexed by $t = 0, 1, \ldots$. The economy is populated by a continuum of households, final-good producers, intermediate-good firms, and the monetary authority. There are two types of states: aggregate and idiosyncratic. We denote the aggregate state by z_t and the idiosyncratic state by v_t, both of which are potentially vector valued. Let $z^t = (z_0, z_1, \ldots, z_t)$ be a history of realized aggregate states up to period t and $v^t = (v_0, v_1, \ldots, v_t)$ be a history of idiosyncratic states up to period t. We let $s_t = (z_t, v_t)$ and $s^t = (z^t, v^t)$. Let $\Pr(s^t|s^{t-1})$ be the probability of a history s^t given s^{t-1}. We assume that $\Pr(s^t|s^{t-1}) = \Pr(v^t|v^{t-1}, z^t)\Pr(z^t|z^{t-1})$ and allow for the possibility that the aggregate states affect the distribution of the idiosyncratic states. To reduce the notation, the initial idiosyncratic

state v_0 also indexes the initial distribution of assets for this agent and the initial aggregate state z_0 indexes the initial distribution of these variables. Thus, without loss of generality, we express all the individual variables as a function of history s^t and, in a symmetric equilibrium, aggregate quantities and prices are functions of z^t.

Households are infinitely lived and have preferences over consumption, $c(s^t)$, and hours worked, $l(s^t)$, given by

$$\sum_{t=0}^{\infty}\sum_{s^t}\beta^t \Pr(s^t|s_0)\tilde{\theta}(z^t)U(c(s^t),l(s^t)), \tag{1}$$

where β is the discount factor and $\tilde{\theta}(z^t)$ is a shock to the marginal utility of consumption and disutility of labor defined recursively as $\tilde{\theta}(z^{t+1})=\theta(z_{t+1})\tilde{\theta}(z^t)$. We further assume that the period utility is given by

$$U\left(c,l\right)=\frac{c^{1-\sigma}-1}{1-\sigma}-\chi\frac{l^{1+\psi}}{1+\psi}, \tag{2}$$

with $\sigma>0$ and $\psi>0$.

The final good is produced combining differentiated intermediate goods according to the technology

$$Y(z^t)=\left(\int_0^1 y_j(z^t)^{\frac{1}{\mu}}dj\right)^{\mu}, \tag{3}$$

where μ is related to the (constant) elasticity of substitution across varieties, ε, by the following, $\mu=\frac{\varepsilon}{\varepsilon-1}$. The intermediate inputs indexed by j are produced using labor

$$y_j(z^t)=A(z_t)n_j(z^t), \tag{4}$$

where $A(z_t)$ is an aggregate technology shock, common across firms, and $n_j(z^t)$ is labor in efficiency units used by the producer of intermediate good j. Feasibility requires that

$$\int n_j(z^t)dj=\sum_{v^t}\Pr(v^t|z^t)e(v_t)l(v^t,z^t), \tag{5}$$

where the right side is the supply of labor in efficiency units. Each individual v^t is associated to a particular level of efficiency $e(v_t)$: hiring more high-efficiency types, holding total hours worked fixed, results in higher output produced by the firm. This individual-specific productivity shock $e(v_t)$ generates idiosyncratic labor income risk for households.

We now describe the market structure for this economy with a particular emphasis on the households side.

2. *Households.* Households begin the period with financial assets and they work for intermediate-good producers. They choose consumption, new financial positions, and labor to maximize their expected lifetime utility.

We model financial markets in a flexible way. Households can trade a risk-free nominal bond. We denote by $b(s^t)$ the position taken today by a household and by $1 + i(z^t)$ the nominal return on the bond. Households can also trade a set $\mathcal{K}$ of additional assets, with the nominal payout of a generic asset $k \in \mathcal{K}$ given by $R_k(s^t, s_{t+1})$. We let $q_k(s^t)$ be the price of the asset. This formulation allows for different types of financial assets: individual Arrow securities, shares of the intermediate-good firms, complex financial derivatives, and so on. We let $a_k(s^{t-1})$ be the holdings of assets k that a household with history v^{t-1} has accumulated after an aggregate history z^{t-1}. Trades in these additional financial assets potentially require transaction costs $\mathcal{T}(\{a_k(s^{t-1})\}_{k\in\mathcal{K}}, \{a_k(s^t)\}_{k\in\mathcal{K}}, s^t)$ that can depend on the inherited portfolio $\{a_k(s^{t-1})\}_{k\in\mathcal{K}}$, the new portfolio $\{a_k(s^t)\}_{k\in\mathcal{K}}$, and s^t. The transaction costs do not apply to $b(s^t)$, so our framework does not nest limited-participation economies where agents must pay a fixed cost to change their position in nominal bonds.

In addition, we allow for a number of constraints that potentially restrict the financial positions that households can choose,

$$H(b(s^t), \{a_k(s^t)\}_{k\in\mathcal{K}}, s^t) \geqslant 0, \tag{6}$$

for some vector-valued function H. We refer to the set of constraints in inequality (6) as trading restrictions. We assume that purchasing risk-free nominal bonds weakly relaxes the trading restrictions, $\frac{\partial H(b,\{a_k\}_{k\in\mathcal{K}},s^t)}{\partial b} \geqslant 0$.

The set of assets $\mathcal{K}$, the transaction costs $\mathcal{T}$, and the trading restrictions in inequality (6) are a flexible way of representing different sets of risk-sharing mechanisms available to households. Our formulation nests the economy with complete financial markets, when the set of assets spans all possible aggregate and idiosyncratic histories and there are no transaction costs or trading restrictions. In addition, it encompasses as special cases a large class of models with incomplete financial markets: the Bewley-Huggett-Aiyagari economy, the two-assets economy in Kaplan and Violante (2014) and Kaplan, Moll, and Violante (2018), the endogenous debt limits in Alvarez and Jermann (2000), or the various restrictions on asset trading in Chien, Cole, and Lustig (2011, 2012). Note also that the H function can depend on s^t, so we allow for aggregate and idiosyncratic shocks to affect households' financial constraints. Moreover, our formulation allows for heterogeneity in households' access to assets other than the risk-free nominal bond, a property that is critical to account for the wealth distribution in the data.

Given initial asset holdings, households choose $\{c(s^t), l(s^t), b(s^t), \{a_k(s^t)\}_{k\in\mathcal{K}}\}$ to maximize their utility subject to the trading restrictions in inequality (6) and the nominal budget constraint,

$$\begin{aligned} P(z^t)c(s^t) &+ \frac{b(s^t)}{1+i(z^t)} + \sum_{k\in\mathcal{K}} q_k(s^t)a_k(s^t) \\ &+ P(z^t)\mathcal{T}(\{a_k(s^{t-1})\}_{k\in\mathcal{K}}, \{a_k(s^t)\}_{k\in\mathcal{K}}, s^t) \\ &\leqslant W(z^t)e(v_t)l(v^t, z^t) - T(s^t) + b(s^{t-1}) \\ &+ \sum_{k\in\mathcal{K}} R_k(s^{t-1}, s_t)a_k(s^{t-1}), \end{aligned}$$

where $W(z^t)$ is the nominal wage per efficiency unit and $T(s^t)$ are lump-sum taxes.

Because of the assumption that $\frac{\partial H}{\partial b} \geqslant 0$, a necessary condition for optimality is

$$(7) \qquad \frac{1}{1+i(z^t)} \geqslant \beta \sum_{s_{t+1}} \left\{ \frac{\Pr(s^{t+1}|s^t)\theta(z_{t+1})}{1+\pi(z^{t+1})} \left[\frac{c(s^t, s_{t+1})}{c(s^t)} \right]^{-\sigma} \right\},$$

where $\pi(z^{t+1}) = \frac{P(z^{t+1})}{P(z^t)} - 1$ is the net inflation rate. This condition must hold with equality if the trading restrictions on the nominal bond do not bind. For the rest of the article, we assume that there always exists an agent for whom equation (7) holds as an equality. Because $\frac{\partial H}{\partial b} \geqslant 0$, equation (7) holds with equality for households with the highest valuation for the risk-free bond.

The condition for the optimality of labor supply is

$$\chi l(s^t)^{\psi} = w(z^t)e(v_t)c(s^t)^{-\sigma}, \tag{8}$$

where $w(z^t) = \frac{W(z^t)}{P(z^t)}$ is the real wage per efficiency unit.

3. *Final-Good Producers.* The final good is produced by competitive firms that operate the production function in equation (3). From their decision problem, we can derive the demand function for the jth variety

$$y_j(z^t) = \left(\frac{P_j(z^t)}{P(z^t)}\right)^{\frac{\mu}{1-\mu}} Y(z^t) \tag{9}$$

where $P_j(z^t)$ is the price of variety j and $P(z^t) = [\int P_j(z^t)^{\frac{1}{1-\mu}} dj]^{1-\mu}$ is the price index.

4. *Intermediate-Good Producers.* Each intermediate good is supplied by a monopolistically competitive firm. The monopolist of variety j operates the technology equation (4). The firm faces quadratic costs to adjust their prices relative to the inflation target of the monetary authority $\bar{\pi}$,

$$\frac{\kappa}{2}\left[\frac{P_j(z^t)}{P_j(z^{t-1})(1+\bar{\pi})} - 1\right]^2. \tag{10}$$

The problem of firm j is to choose its price $P_j(z^t)$ given its previous price $P_j(z^{t-1})$ to maximize the present discounted value of real profits. As is well known, state prices in incomplete market economies are not uniquely determined. The issue is even more complex in our framework because we are purposefully not fully specifying the set of assets available and the trading restrictions in inequality (6). We resolve this indeterminacy by assuming that the firm discounts future profits using the real state price

$$Q(z^{t+1}) = \beta \max_{v^t}\left\{\Pr(z^{t+1}|z^t)\theta(z_{t+1})\sum_{v_{t+1}}\Pr(v^{t+1}|z^{t+1}, v^t) \times \left[\frac{c(z^{t+1}, v^{t+1})}{c(z^t, v^t)}\right]^{-\sigma}\right\}. \tag{11}$$

That is, firms discount future profits using the marginal rate of substitution of the agent that values dividends in a given aggregate state next period the most. This would be the equilibrium state price if all agents could trade aggregate Arrow securities.

The firm's problem can be written recursively as

$$V(P_j, z^t) = \max_{p_j, y_j, n_j} \frac{p_j y_j}{P(z^t)} - w(z^t) n_j(z^t) - \frac{\kappa}{2}\left[\frac{p_j}{P_j(1+\bar{\pi})} - 1\right]^2 + \sum_{z^{t+1}} Q(z^{t+1}|z^t) V(p_j, z^{t+1}), \tag{12}$$

subject to the production function (4) and the demand function (9).

The solution to the firm's problem together with symmetry across firms requires that the following version of the New Keynesian Phillips curve holds in equilibrium

$$\tilde{\pi}(z^t) = \frac{1}{\kappa(\mu-1)} Y(z^t)\left[\mu \frac{w(z^t)}{A(z_t)} - 1\right] + \sum_{z^{t+1}} Q(z^{t+1}|z^t)\tilde{\pi}(z^{t+1}), \tag{13}$$

where we define $\tilde{\pi}(z^t) = [\frac{\pi(z^t)-\bar{\pi}}{1+\bar{\pi}}] \times [\frac{\pi(z^t)+1}{1+\bar{\pi}}]$ and $\frac{w(z^t)}{A(z_t)}$ is the real marginal cost for producing a unit of the final good.

5. *Monetary and Fiscal Policy.* The monetary authority follows a standard Taylor rule

$$1 + i(z^t) = \max\left\{[1 + i(z^{t-1})]^{\rho_i}\left[(1+\bar{i})\left(\frac{1+\pi(z^t)}{1+\bar{\pi}}\right)^{\gamma_\pi}\right]^{1-\rho_i} \times \exp\{\varepsilon_m(z_t)\}, 1\right\}, \tag{14}$$

where $1 + \bar{i} = \frac{1+\bar{\pi}}{\beta}$ is the nominal interest in a deterministic steady state of the model and $\varepsilon_m(z_t)$ is a monetary shock. Because of the zero lower bound, whenever the interest rate predicted by the Taylor rule is negative, the monetary authority sets the nominal interest rate to zero.

The evolution of the aggregate supply of the nominal bond, $B(z^t)$, and taxes, $T(s^t)$, must satisfy the government budget constraint,

$$(15) \qquad B(z^{t-1}) = \frac{B(z^t)}{1+i(z^t)} + \sum_{v^t} \Pr(v^t|z^t) T(z^t, v^t).$$

6. *Equilibrium.* In equilibrium, the labor market, goods markets, and financial markets clear. Specifically, market clearing in the nominal bond market requires that

$$(16) \qquad \sum_{v^t} \Pr(v^t|z^t) b(z^t, v^t) = B(z^t).$$

For the other assets, given the asset supply $\bar{a}_k(z^t)$, market clearing requires that

$$(17) \qquad \sum_{v^t} \Pr(v^t|z^t) a_k(z^t, v^t) = \bar{a}_k(z^t).$$

Moreover, because firms' equity is the only asset in positive net supply other than the nominal risk-free bond, the supply of the other assets available and their returns R_k must satisfy the following restriction to ensure that agents' budget constraints are consistent with the aggregate resource constraint:

$$(18) \qquad \sum_{k\in\mathcal{K}} \sum_{v^t} \Pr(v^t|z^t) R_k(z^t, v^t) \bar{a}_k(z^{t-1}) - \sum_{k\in\mathcal{K}} q_k(z^t) \bar{a}_k(z^t) = D(z^t),$$

where $D(z^t) = [P(z^t) - \frac{W(z^t)}{A(z_t)}] Y(z^t) - \frac{\kappa}{2}[\frac{1+\pi(z^t)}{1+\bar{\pi}} - 1]^2$ are aggregate nominal firm profits.

We can then define an equilibrium for this economy.

DEFINITION 1. Given an asset structure $(\mathcal{K}, \mathcal{T}, R_k, \bar{a}_k, H)$, the distribution of initial assets, and lagged prices, an equilibrium is a set of households' allocations $\{c(s^t),\ l(s^t),\ b(s^t),\ a_k(s^t)\}$, a fiscal policy $\{B(z^t),\ T(s^t)\}$, prices $\{P(z^t),\ W(z^t),\ 1 + i(z^t),\ Q(z^t),\ q_k(z^t)\}$, and aggregates $\{C(z^t),\ Y(z^t)\}$ such that (i) the households' allocation solves the households' decision problem, (ii) the price for the final good solves equation (12) with $P(z^t) = P_j(z^t)$, (iii) the state price is given by equation (11), (iv) the nominal interest rate satisfies the Taylor rule equation (14), (v) the government budget constraint (15) is satisfied, and (vi) markets clear in that equations (16)–(18)

hold, and

$$Y(z^t) = C(z^t) + \frac{\kappa}{2}\left[\frac{\pi(z^t) - \bar{\pi}}{1+\bar{\pi}}\right]^2 + \mathcal{T}(z^t),$$

where aggregates are given by

$$Y(z^t) = A(z_t)\sum_{v^t} \Pr(v^t|z^t)e(v_t)l(v^t, z^t),$$

$$C(z^t) = \sum_{v^t} \Pr(v^t|z^t)c(z^t, v^t),$$

and $\mathcal{T}(z^t)$ are the aggregate transaction costs,

$$\mathcal{T}(z^t) = \sum_{v^t} \Pr(v^t|z^t)\mathcal{T}(\{a_k(s^{t-1})\}_{k\in\mathcal{K}}, \{a_k(s^t)\}_{k\in\mathcal{K}}, s^t).$$

III. The RA Representation

We show that the aggregate variables in the class of New Keynesian models just described can be equivalently derived from the equilibrium conditions of a representative-agent economy with wedges. We refer to this as the RA representation. Section III.A derives the RA representation, and Section III.B discusses how we can use it for counterfactual analysis. Sections III.C and III.D study specific economies nested in our framework, which admits an analytical mapping between the model primitives and the wedges. This analysis will be useful for building intuition on how different types of shocks and frictions studied in the literature affect the wedges in the RA representation.

III.A. Equilibrium Representation

Toward establishing the RA representation for our class of New Keynesian models, let us define the following statistics:

$$\beta(v^t, z^{t+1}) \equiv \sum_{v_{t+1}} \Pr(v_{t+1}|v^t, z^{t+1})\left(\frac{c(z^{t+1}, v^t, v_{t+1})}{C(z^{t+1})}\frac{C(z^t)}{c(z^t, v^t)}\right)^{-\sigma}$$

(19)

$$\omega(z^t) \equiv \left[\sum_{v^t} \Pr(v^t|z^t) \left(\frac{c(z^t, v^t)}{C(z^t)} \right)^{\frac{-\sigma}{\psi}} e\,(v_t)^{\frac{1+\psi}{\psi}} \right]^{-\psi}. \tag{20}$$

We then have the following proposition where we assume that the aggregate transaction costs, $\mathcal{T}(z^t)$, are negligible.

PROPOSITION 1. Suppose that $\{C(z^t), Y(z^t), \pi(z^t), i(z^t), Q(z^{t+1})\}$ are part of an equilibrium of a heterogeneous-agent economy described in Section II. Then, they must satisfy the aggregate Euler equation,

$$\frac{1}{1+i(z^t)} = \beta \max_{v^t} \sum_{z_{t+1}} \Pr(z^{t+1}|z^t) \times \left[\frac{\theta(z_{t+1})\beta(v^t, z^{t+1})}{1+\pi(z^{t+1})} \left(\frac{C\left(z^{t+1}\right)}{C(z^t)} \right)^{-\sigma} \right], \tag{21}$$

the Phillips curve,

$$\tilde{\pi}(z^t) = \frac{Y(z^t)}{\kappa\,(\mu-1)} \left[\mu \frac{\chi Y(z^t)^{\psi} C(z^t)^{\sigma} \omega(z^t)}{A(z_t)^{1+\psi}} - 1 \right] + \sum_{z^{t+1}} Q(z^{t+1}|z^t)\tilde{\pi}(z^{t+1}), \tag{22}$$

the Taylor rule equation (14), the resource constraint,

$$Y(z^t) = C(z^t) + \frac{\kappa}{2} \left[\frac{\pi(z^t) - \bar{\pi}}{1+\bar{\pi}} \right]^2, \tag{23}$$

and

$$Q(z^{t+1}|z^t) = \beta \max_{v^t} \left\{ \beta(v^t, z^{t+1}) \Pr(z^{t+1}|z^t) \theta(z_{t+1}) \left(\frac{C(z^{t+1})}{C\,(z^t)} \right)^{\sigma} \right\}, \tag{24}$$

given $\{\beta(v^t, z^{t+1}), \omega(z^t)\}$ defined in equations (19) and (20).

The proof for this result is straightforward, and it extends the result of Nakajima (2005) to an economy with nominal rigidities.

The aggregate Euler equation (21) is obtained by substituting the households' marginal rate of substitution in equation (7) using the definition of $\beta(v^t,\ z^{t+1})$ and noting that under our assumptions the Euler equation holds for agents with the highest valuation of the bond—those attaining the maximum in equation (21).

For the aggregate Phillips curve equation (22) we proceed as follows. We raise both sides of households' labor supply condition (8) by $\frac{1}{\psi}$, multiply them by $e(v_t)C(z^t)^{\frac{\sigma}{\psi}}$ and average across households to obtain

$$\chi^{\frac{1}{\psi}}\left[\sum_{v^t}\Pr(v^t|z^t)e(v_t)l(s^t)\right]C(z^t)^{\frac{\sigma}{\psi}}$$
$$= w(z^t)^{\frac{1}{\psi}}\left[\sum_{v^t}\Pr(v^t|z^t)\left(\frac{c(z^t,\,v^t)}{C(z^t)}\right)^{\frac{-\sigma}{\psi}}e\,(v_t)^{\frac{1+\psi}{\psi}}\right].$$

We can use the production function (4) and the definition of $\omega(z^t)$ in equation (20) to express the real wage as

$$w(z^t) = \chi\left[\frac{Y(z^t)}{A(z_t)}\right]^{\psi}C(z^t)^{\sigma}\omega(z^t).$$

We substitute the above expression in equation (13) to obtain the aggregate Phillips curve equation (22).

The equilibrium conditions (14), (21), (22), and (23) are equivalent to those of a New Keynesian representative agent with two wedges: one in the Euler equation and one in the Phillips curve. We refer to the first as the discount factor wedge, and to the second as the labor supply wedge. These wedges are functions of the individual allocations in the original heterogeneous-agent economy and do not typically have a structural interpretation.

Before moving forward, we discuss some potential limitations of our framework. First, although we have been flexible on the households' decision problem, we made restrictive assumptions about other aspects of the model. For instance, there is no capital accumulation in this economy, wages are perfectly flexible, all movements in labor are at the intensive margin, and we have taken a stand on some of the aggregate shocks driving the economy—the technology, monetary, and preference shocks. In Online Appendix A we show that it is relatively straightforward to derive an equivalent of Proposition 1 for economies with a more complex macro block. Specifically, we derive the RA representation in a model that features capital accumulation, one with frictions in financial intermediation as in Gertler and Karadi (2011) and Gertler and Kiyotaki (2010), and for a small open economy model where asset prices matter for the amount of borrowing as in Mendoza (2010). All these economies admit an RA representation similar to the one in Proposition 1, with a discount factor and labor supply wedge as defined in equations (19) and (20). Intuitively, this is because market incompleteness in these economies matters for firms' production choices only through its effect on their stochastic discount factor.[7] Thus, these model ingredients do not affect the mapping between the micro-level allocations and the wedges, but they entail a different propagation of these wedges to aggregate outcomes.

Second, we have restricted households' preferences to be separable and isoelastic. While the definition of the wedges relies on this assumption, it is worth pointing out that we can obtain similar RA representations for different sets of preferences. For example, in Online Appendix C.1 we consider an economy where households' preferences allow for corner solutions in worked hours. We show that this economy admits an RA representation as in our benchmark economy, but with a slightly different labor supply wedge. In addition, Online Appendix A.4 studies an economy where households have preferences over durable and nondurable consumption goods. Also in this case we derive the RA representation and find it to be quite similar to that of our benchmark. Interestingly, the two are identical when the underlying heterogeneous-agent economy features no dispersion across households in the ratio between durable and nondurable consumption.

Third, we have assumed that firms discount future profits using the marginal rate of substitution of the agent that values dividends the most.[8] This choice is arbitrary, but it does not

7. This property does not hold in economies where the joint distribution of capital and entrepreneurial ability directly matters for production choices, as for instance in Buera and Moll (2015). Differently from our framework, the RA representation of these economies will also feature an "efficiency wedge" in the terminology of Chari, Kehoe, and McGrattan (2007).

8. An alternative would be to follow the approach in Makowski (1983a, 1983b) and let the state price used by the firm be the valuation of the agent that maximizes the firm's stock value for any possible choices for the firm, assuming that agents can trade stocks without frictions. See Bisin, Clementi, and Gottardi (2014) for an illustration of the appealing implications of this approach.

affect the definition of the wedges in the RA representation. That is, different assumptions on firms' capital structure choices and dividend distribution policies affect households' consumption, but do not alter the mapping between households' consumption choices and the wedges.

III.B. Using the RA Representation for Model Evaluation

Proposition 1 has two main implications. The first implication is that the wedges summarize all the information from the micro block of the model that is needed to characterize the behavior of aggregate variables. That is, they are sufficient statistics for the specifics of the model regarding the set of assets traded, the transaction costs and trading restrictions faced by households, the fiscal policy $\{B(z^t), T(s^t)\}$, and the nature of their idiosyncratic risk. The second implication is that the mapping between individual allocations and the wedges is invariant to the details of the micro block, in the sense that for all the economies nested in the environment of Section II, the relation between $\{\beta(v^t, z^{t+1}), \omega(z^t)\}$ and the households' allocation is the same and it is given by equations (19) and (20). These two properties, in turn, make the RA representation a useful device for the empirical analysis of New Keynesian economies with heterogeneous agents.

First, suppose that we have a procedure to measure $\{\beta(v^t, z^{t+1}), \omega(z^t)\}$ using observations on households' consumption shares and labor productivities. Because these are sufficient statistics for how the micro block affects macroeconomic variables, they also provide an informative empirical target for the calibration/evaluation of a specific model. Researchers that wish to use a specific model nested in this class should make sure that their model fits the statistical properties of the wedges observed in the data.

Second, we can use the RA representation to address a question that goes back to the seminal work of Krusell and Smith (1998)—to what extent imperfect risk sharing across households matters for aggregate fluctuations. To address this question, one needs to compare the behavior of observed macroeconomic variables to those that would arise in a world where households could perfectly insure their idiosyncratic risk. These counterfactuals can be computed using the RA representation, as the next result shows.

PROPOSITION 2. Consider the class of heterogeneous-agent economies described in Section II and suppose that financial markets are complete—the set of assets $\mathcal{K}$ contains Arrow securities contingent on the realizations of the aggregate and idiosyncratic state, and there are no trading restrictions other than a nonbinding no-Ponzi condition. Then, if $\{C^{\text{cm}}(z^t), Y^{\text{cm}}(z^t), \pi^{\text{cm}}(z^t), i^{\text{cm}}(z^t), Q^{\text{cm}}(z^{t+1})\}$ are part of an equilibrium they satisfy equations (14), (21), (22), (23), and (24) with

$$\beta^{\text{cm}}(v^t, z^{t+1}) = 1 \tag{25}$$

$$\omega^{\text{cm}}(z^t) = \left[\sum_{v^t} \Pr(v^t|z^t)\varphi(v_0)^{\frac{-\sigma}{\psi}} e(v_t)^{\frac{1+\psi}{\psi}}\right]^{-\psi}. \tag{26}$$

This result is an application of the aggregation result of Constantinides (1982) that Maliar and Maliar (2001, 2003) extend to economies with endogenous labor supply. Given isoelastic preferences, the complete-market economy features constant consumption shares for households, implying that $\beta^{\text{cm}}(v^t, z^{t+1}) = 1$. So the Euler equation of the heterogeneous-agent economy with complete markets coincides with the Euler equation of the representative-agent economy. Aggregate labor supply in the complete-market economy, however, can differ from that of the representative-agent economy because of substitution effects due to idiosyncratic productivity shocks. These differences are captured by the wedge $\omega^{\text{cm}}(z^t)$.[9]

To explain how we can leverage Propositions 1 and 2 to quantify the macroeconomic effects of imperfect risk sharing, let us assume for now that we know the probability distribution of z^t and the stochastic process for $\{\theta(z^t), A(z^t), \varepsilon_m(z^t), \beta(v^t, z^{t+1}), \omega(z^t), \omega^{\text{cm}}(z^t)\}$. Given a realization of z^t, we can use the RA representation of Proposition 1 to obtain the underlying equilibrium path for aggregate variables—output, inflation, and nominal interest rates. We label these the actual paths. We can then compare these paths to the complete-markets paths computed by setting

9. To understand this expression, suppose that $\psi = 1$ and households have the same initial wealth, so that $\varphi(v_0) = 1$. In this case, $\omega(z^t) = \text{Var}[e(v_t)|z^t]^{-1}$. When the variance of $e(v_t)$ increases, the labor supplied by high-productivity households increases and the labor supplied by the low-productivity households declines because of a substitution effect. So labor supply goes up when measured in efficiency units. This effect is captured by a decline in $\omega(z^t)$ in the RA representation—an increase in the labor supply of the stand-in household.

$\beta(v^t, z^{t+1}) = 1$ and $\omega(z^t) = \omega^{\mathrm{cm}}(z^t)$ in the RA representation. These paths will differ because they feature a different realization and a different stochastic process for the wedges: by comparing them we are able to isolate the effect that imperfect risk sharing has for macroeconomic aggregates over the particular history z^t.

To carry out this counterfactual, we need a procedure to measure the realization of $\{\beta(v^t, z^{t+1}), \omega(z^t), \omega^{\mathrm{cm}}(z^t)\}$ using household-level data and to approximate their stochastic process. We discuss these two issues in Section IV and apply our framework to U.S. data in Section V. Before moving there, let us discuss how different economic mechanisms map into the wedges of the RA representation using some analytically tractable examples. We focus mostly on the discount factor wedge because it turns out to be empirically the dominant factor in our application to the U.S. economy.

Readers less interested in these examples can move directly to Section IV without losing the thread.

III.C. Households' Precautionary Savings and the Discount Factor Wedge

We illustrate Proposition 1 using two examples. In the first example we study an incomplete-markets economy where households' idiosyncratic income risk is time-varying. In this economy, an increase in the volatility of households' income leads to higher incentives to save for precautionary reasons, which can result in a fall in aggregate consumption and output—a mechanism studied by Bayer et al. (2019), Heathcote and Perri (2018), and Ravn and Sterk (2021) to explain the depth and persistence of the Great Recession. In the second example, instead, we consider an economy where households' precautionary saving motives are triggered by a time-varying borrowing constraint, as in Guerrieri and Lorenzoni (2017). In both examples, we show that these time-varying precautionary saving motives are captured in the RA representation by the discount factor wedge $\beta(v^t, z^{t+1})$.

1. *Labor Income Risk and the Discount Factor Wedge.* Let $\sigma = 1$ and the idiosyncratic productivity shocks evolve according to

$$\Delta \log[e(v_{t+1})] = -\frac{\sigma_e^2(z_t)}{2} + \varepsilon_{t+1} \qquad \varepsilon_{t+1}|z^{t+1} \sim \mathcal{N}\big(0, \sigma_e^2(z_t)\big).$$

That is, idiosyncratic productivity is a random walk with Gaussian shocks. The standard deviation of individual productivity growth varies over time with the aggregate state: when $\sigma_e^2(z_t)$ is high, households face higher income risk.

We assume that households can only trade the risk-free bonds in zero net supply and face the borrowing limit $b(s^t) \geqslant 0$.[10] These two assumptions imply that households cannot save in equilibrium.[11] Therefore, all households are hand-to-mouth and consume every period their after-tax income, $e(v_t)[w(z^t)l(s^t) - T(s^t)]$. Furthermore, assuming that lump-sum taxes are proportional to $e(v_t)$, $T(s^t) = e(v_t)T(z^t)$, we can verify from the labor supply condition (8) and $\sigma = 1$ that $l(s^t)$ is the same across individuals. So we have $c(s^t) = e(v_t)C(z^t)$ from the aggregate resource constraint.

Given the equilibrium consumption function, the relative marginal rates of substitution of the households are just functions of the idiosyncratic income process,

$$\left(\frac{c(z^{t+1}, v^t, v_{t+1})}{C(z^{t+1})}\frac{C(z^t)}{c(z^t, v^t)}\right)^{-1} = \frac{e(v_t)}{e(v_{t+1})}.$$

Substituting these expressions in equations (19) and (20) we can compute the implied $\beta(v^t, z^{t+1})$ and $\omega(z^t)$ of this economy:

$$\beta(v^t, z^{t+1}) = \exp\left\{\sigma_e^2(z_t)\right\}$$
$$\omega(z^t) = 1.$$

The discount factor wedge is the same across households, and it depends on the realization of the aggregate shock at date t: the higher the idiosyncratic income risk at date t, the more "patient" the stand-in household in the RA representation. The labor supply wedge is constant over time because households supply the same amount of labor in equilibrium.

This example is useful to understand how the interaction between idiosyncratic risk and incomplete financial markets can affect aggregate variables in New Keynesian models with incomplete markets. Suppose that households face higher idiosyncratic risk, that is $\sigma_e^2(z_t)$ increases. Because of incomplete

10. Because households cannot trade stocks of the firms, we also assume that profits of the intermediate-good producers are taxed by the government and rebated to households in proportion to the realization of idiosyncratic productivity.

11. The literature refers to this example with tight borrowing limits and bonds in zero net supply as the zero liquidity limit. See Werning (2015) and Ravn and Sterk (2017) for examples.

markets, households respond to heightened risk by increasing their demand for the risk-free bond. In equilibrium, households are hand-to-mouth, so aggregate savings cannot change. Thus, the increase in households' demand for savings must be met in equilibrium by a fall in real interest rates and/or by a decline in their income at date t. Whether these precautionary motives are mostly reflected in prices or quantities depends on the response of the monetary authority.

2. *Households' Borrowing Constraint and the Discount Factor Wedge.* We consider the same environment of the previous example but change the specification of the households' borrowing constraint. The debt limits are now given by $b(s^t) \geqslant -\phi_t(z)Y_t(z)$. We assume that z can take two values, $z \in \{z_L, z_H\}$. If $z = z_L$, then debt limits are tight forever as in the previous example, $\phi_t(z) = 0$ for all $t \geqslant 0$. If $z = z_H$, however, then debt limits are tight at $t = 0$, but agents can borrow a fraction $\phi > 0$ of aggregate output from period $t = 1$ onward. Thus, the realization of z at date $t = 0$ determines to what extent households can hedge future income shocks by borrowing. We simplify the idiosyncratic income process and assume that e_t can take two values with equal probability, e_H and e_L with $e_H > e_L$.

In period 0, because $\phi_0(z) = 0$ irrespective of the realization of z, all households are hand-to-mouth and consumption is proportional to the idiosyncratic shock, $c_0(e_0, z) = e_0C_0(z)$. The same is true from period 1 onward if $z = z_L$. Under the more relaxed debt limit, instead, households that have negative income shocks can borrow from households with positive income shocks, so the allocation will be different. In period 1, for example, we have $c(e_0, e_H, z_H) = [e_H - \Delta]C_1(z_H)$ and $c(e_0, e_L, z_H) = [e_L + \Delta]C_1(z_H)$ for $\Delta > 0$.[12] Thus, households' consumption at date $t = 1$ is less volatile when $z = z_H$ than when $z = z_L$.

As in the previous analysis, households' precautionary motives depend on the realization of the aggregate shock. If $z = z_L$, households face more volatile consumption in the future and have higher precautionary saving motives, which puts downward pressures on interest rates and output. If $z = z_H$, households have lower precautionary saving motives, so output and interest rates are higher.

The RA representation captures these effects via the discount factor wedge, which is higher when $z = z_L$ than when $z = z_H$. To

12. In equilibrium, $\Delta = \frac{b(e_0, e_H, z_H)}{C_1(z_H)} > 0$ if $\phi > 0$.

see why, consider the households that at date 0 have high income. These are those with the highest incentives to save and they are thus the ones achieving the maximum in equation (21). Using their consumption, we can express the discount factor wedge as

$$\beta\left(e_H, z_L\right) = \frac{1}{2}\frac{e_H}{e_H} + \frac{1}{2}\frac{e_H}{e_L} > \frac{1}{2}\frac{e_H}{e_H - \Delta} + \frac{1}{2}\frac{e_H}{e_L + \Delta} = \beta\left(e_H, z_H\right).$$

Higher precautionary saving motives induced by a tightening of future debt limits are isomorphic to a higher discount factor wedge in the RA representation.[13]

III.D. The Role of Hand-to-Mouth Consumers

A series of recent papers has emphasized the critical role of agents with high marginal propensities to consume (MPCs) for the amplification of aggregate shocks in New Keynesian models with incomplete markets; see Auclert, Rognlie, and Straub (2018) and Kaplan and Violante (2022). At first sight, it appears that these considerations are not factored in the RA representation: after all, the discount factor wedge is computed using only the consumption share of unconstrained households—those with low MPCs. As we show in this subsection, however, this is not true, as the discount factor wedge can depend on the behavior of constrained households because of general equilibrium.

To explore this point, we consider an economy with MPC heterogeneity. In this economy, a transfer from low- to high-MPC households increases output, an effect that is stronger the larger is the share of high-MPC households in the population. We then derive the RA representation and show two results. First, the discount factor wedge declines with the transfers. Second, the sensitivity of the discount factor wedge to the transfer increases in the share of high-MPC households.

The economy lasts two periods, $t = 0, 1$, and there are only two levels for idiosyncratic productivity in period 0, $e_0 \in \{e_L, e_H\}$. We let μ_L be the probability of drawing e_L. In period 1, all agents have $e_1 = 1$. As in the previous examples, we assume that households can trade only the nominal bond. The household's

13. The labor supply wedge is equal to one at time 0—as in the previous example—because households are hand-to-mouth and $\sigma = 1$.

budget constraints are,

$$c_{i,0} + b_i \leqslant e_i w_0 l_{i,0} + T_0$$
$$c_{i,1} \leqslant w_1 l_{i,1} + b_{i,1}(1+r) - T_{i,1}.$$

We assume that households face tight debt limits, $b_i \geqslant 0$, and prices are assumed to be fully sticky in period 0 and perfectly flexible in period 1, $\kappa_0 = \infty$ and $\kappa_1 = 0$. To simplify further, the monetary authority sets nominal interest rates so that $\beta(1+i)\frac{P_0}{P_1} = 1$, implying a constant real rate equal to $\frac{1}{\beta}$.

Fiscal policy works as follows. In period 0 the government taxes firms' profits and issues debt B, redistributing the proceeds in a lump-sum fashion to all households. In period 1 the government taxes firms' profits and households to repay the debt issued in period 0. The taxes are lump sum and type-specific. The government budget constraints are

$$T_0 = B + \Pi_0$$
$$\sum_i \mu_i T_{i,1} = B(1+r) + \Pi_1.$$

Market clearing requires that $B = \sum_i \mu_i b_i$. To simplify the algebra, we further assume that the type-specific taxes in period 1 are $T_{i,1} = (1 + r)b_i + \Pi_1$, so agents have the same consumption in the last period. We assume $\sigma = \psi = 1$ and choose χ so that $c_{i,1} = 1$ for all i. From this it follows that $C_1 = 1$.

A full characterization of this example is presented in Online Appendix B. Because $e_L < e_H$, the debt limit binds for low-productivity households if T_0 is small enough. These households are effectively hand-to-mouth, as they consume all of the additional transfer they may receive. In contrast, the consumption of high-productivity households is constant over time because $(1+r) = \frac{1}{\beta}$, and equal to 1 because $c_{H,1} = 1$. Thus, because consumption of high-productivity households does not move with transfers while the consumption of low-productivity households moves one-for-one with transfers, the higher is the share of agents with a high marginal propensity to consume, μ_L, the larger is the response of aggregate consumption and output to T_0. This is illustrated in the left panel of Figure I, where we set $\beta = 0.9$ for illustration.

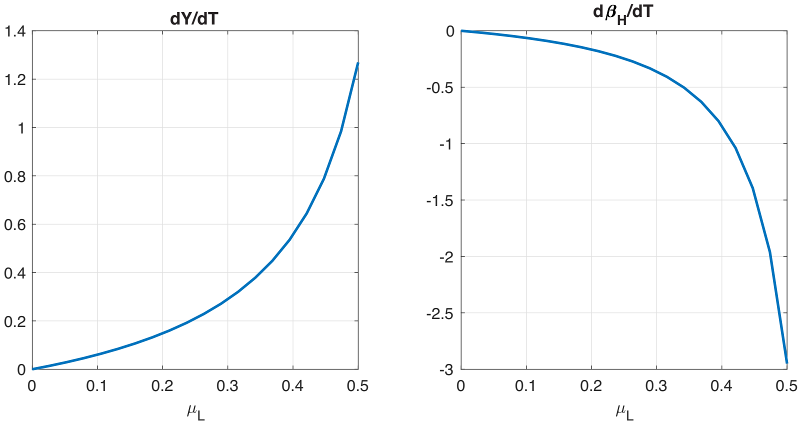

FIGURE I

The Sensitivity of Aggregate Consumption and the Discount Factor Wedge to Transfers

Let us consider how the RA representation captures these effects. The aggregate Euler equation is

$$\frac{1}{C_0} = \beta \hat{\beta}_H (1+r) \frac{1}{C_1}, \tag{27}$$

where $\hat{\beta}_H$ is the discount factor wedge for the unconstrained households—those with high productivity in period 0,

$$\hat{\beta}_H = \frac{C_1}{c_{H,1}} \frac{c_{H,0}}{C_0} = \frac{c_{H,0}}{C_0} = \frac{c_{H,0}}{(1-\mu_L)c_{H,0} + \mu_L c_{L,0}} = \frac{1}{(1-\mu_L) + \mu_L c_{L,0}}.$$

From the above expression we can see that the discount factor wedge depends on the consumption of constrained households and on their share in the population, even though we use only the consumption share of unconstrained households to compute it. We can then use the RA representation to study the aggregate effects of T_0. The higher the transfer, the higher the consumption of constrained households, the lower is the discount factor wedge $\hat{\beta}_H$. In addition, the higher the share of high-MPC households, μ_L, the larger the sensitivity of the discount factor wedge to T_0, a result that is illustrated in the right panel of Figure I. The fall in $\hat{\beta}_H$ then leads to an increase in aggregate consumption via equation (27).[14]

This example illustrates more generally how the RA representation captures the amplification mechanism that takes place in the canonical two-agent New Keynesian model.[15] As explained in Bilbiie (2020), this model produces amplification when the income of hand-to-mouth households responds more than one-for-one to a change in aggregate income. When aggregate income falls because of some shock, the consumption of hand-to-mouth households falls by more, and this sets in motion the amplification mechanism. The RA representation captures the amplification via the discount factor wedge: the fact that the consumption of hand-to-mouth households falls by more than aggregate income implies that the consumption share of the unconstrained households increases, which result in an increase in the discount factor wedge.

IV. Measuring the Wedges

As we discussed in Section III.B, we need a procedure to measure the wedges and approximate their stochastic process to use the RA representation for counterfactual analysis. We explain how we accomplish these steps. In Section IV.A we show how we can use panel data to estimate the realization of the wedges, and Section IV.B introduces a class of stochastic processes that we use to approximate their law of motion.

IV.A. Measuring the Wedges from Panel Data

Let us denote by i a household with a history of idiosyncratic shocks v^t. We assume that we observe a panel of N households' consumption choices and wage per worked hours, $\{c_{i,t}, w_{i,t}\}$ and the time series for aggregate consumption and wages, $\{C_t, W_t\}$. We also assume that we observe households' financial assets

14. In this example, $\beta(1 + r) = 1$ and $C_1 = 1$. So the aggregate Euler equation (27) can also be written as $\frac{1}{C_0} = \hat{\beta}_H$, describing an inverse relation between C_0 and $\hat{\beta}$.

15. This model features limited participation in the bond market, so it is not formally nested in our framework. However, it still has an RA representation in line with that of Proposition 1. See Online Appendix A.5 for this derivation.

and liabilities.[16] In what follows, we show how we can use these observations to recover the realization of the wedges.

Let us start with the discount factor wedge. From equation (19) we have that the discount factor wedge is the conditional expectation, across realizations of the idiosyncratic state, of the change in the consumption share between two periods raised to a power of $-\sigma$. If we observed multiple histories of consumption choices for the same households, we could compute this statistic for each household i by averaging across these consumption histories. This approach is clearly not feasible because we observe only one consumption path for each household, so we need an alternative way of estimating this conditional expectation.

We proceed as follows. For each household in the panel we compute the consumption shares $\varphi_{i,t} = \frac{c_{i,t}}{C_t}$. We then group households with similar characteristics at date t, leaving us with G groups. For each group $g \in G$, we compute the statistic

$$\beta_{g,t+1} \equiv \frac{1}{N_g} \sum_{i=1}^{N_g} \left(\frac{\varphi_{i,t+1}}{\varphi_{i,t}} \right)^{-\sigma}, \tag{28}$$

where N_g is the number of households in group g at time t.

The logic behind this approach builds on two premises. The first is that by grouping households along observable characteristics we are proxying for the individual history v^t. The second is that the size of the groups is large enough so that $\beta_{g,t+1}$ approximates the conditional expectation in equation (19) by the law of large numbers.

In our application in Section V, we consider partitions based on households' labor income and net worth: at date t we group households according to whether their labor income is above or below median income and, within these groups, whether their net worth is above or below the group median. Thus, for each t, we end up with four different groups of households of approximately equal size: low income/low net worth, low income/high net worth, high income/low net worth and high income/high net worth. For each group g, we use equation (28) to construct $\beta_{g,t+1}$. The rationale behind this partition is that income and net worth are

16. In our application these data series will be residualized to control for demographic factors and other household characteristics that are not included in our model.

sufficient statistics for an individual history v^t in benchmark incomplete-market economies.

Let us now turn to the measurement of ω_t. From equation (20) we can see that we need two household-level observations to construct this series: the consumption shares $\varphi_{i,t}$ and the idiosyncratic productivity $e_{i,t}$. In the class of models described in Section II, idiosyncratic productivity equals the ratio between the wage per hour of household i and the average hourly wage in the economy, $e_{i,t} = \frac{w_{i,t}}{W_t}$. Given $\{\varphi_{i,t}, e_{i,t}\}$, we can compute the wedge ω_t using the expression

$$\omega_t = \left[\frac{1}{N}\sum_{i=1}^{N} \varphi_{i,t}^{-\frac{\sigma}{\psi}} e_{i,t}^{\frac{1+\psi}{\psi}}\right]^{-\frac{1}{\psi}}. \tag{29}$$

If N is large enough, this expression is equivalent to the one in equation (20).

Finally, consider ω_t^{cm} defined in equation (26). To compute this object, we need the observed idiosyncratic productivity and the counterfactual behavior for the consumption shares in the economy with complete markets. Given our assumptions on households' preferences, consumption shares are not time-varying when markets are complete, but they can be potentially different across households because of initial heterogeneity in wealth. Therefore, to compute ω_t^{cm}, we need to know the initial distribution of consumption shares and its correlation with $e_{i,t}$ for every t. In our application we assume that the moments of the initial distribution are those of the first year in our sample. That is, we compute ω_t^{cm} as follows

$$\omega_t^{\mathrm{cm}} = \left[\frac{1}{N}\sum_{i=1}^{N} \varphi_{i,1}^{-\frac{\sigma}{\psi}} \times \frac{1}{N}\sum_{i=1}^{N} e_{i,t}^{\frac{1+\psi}{\psi}} + \mathrm{Cov}\left(\varphi_{i,1}^{-\frac{\sigma}{\psi}}, e_{i,1}^{\frac{1+\psi}{\psi}}\right)\right]^{-\frac{1}{\psi}}. \tag{30}$$

IV.B. Stochastic Process for the Wedges

As in Chari, Kehoe, and McGrattan (2007), we assume a Markov structure for the states, $\Pr(s_t|s^{t-1}) = \Pr(s_t|s_{t-1})$, and suppose that the equilibrium outcome is induced by a recursive competitive equilibrium. Under these assumptions, $\{\beta_{i,t+1}, \omega_{t+1}, \omega_{t+1}^{\mathrm{cm}}\}$ are functions of the aggregate state variables of the model. In what follows, we specify a class of stochastic process for the wedges based on a first-order approximation of these functions.

In a recursive competitive equilibrium, endogenous variables are functions of idiosyncratic and aggregate states. Let (z, X) be the current realization of the aggregate exogenous and endogenous states, with transition $X' = \Gamma(X, z)$, and let (v, x) be the exogenous and endogenous idiosyncratic states of the model, with transition $x' = \gamma(x, v, z, X)$.

To make things more concrete, consider a simple economy nested in the class of models of Section II. Suppose that households can only save and borrow in the risk-free nominal bond and face a borrowing limit $b_{i,t+1} \geqslant -\phi$, and assume that their idiosyncratic productivity e is an AR(1) process. In a recursive competitive equilibrium, the exogenous aggregate state is $z = [\theta, A, \varepsilon_m]$, the endogenous aggregate state is $X = [\Psi(e, b), i]$—with Ψ being the distribution of individual productivity and bond holdings, and i the lagged nominal interest rate—and the idiosyncratic state variables would be $v = e$ and $x = b$.

Due to the recursive structure, it is straightforward to derive the implied stochastic processes for the discount factor wedge. Denoting by $\varphi(v, x, z, X)$ the consumption share of a household i with individual characteristics (v, x) in the aggregate state (z, X), we can write the discount factor wedge for this household i as

$$\begin{aligned}\beta_{i,t+1} &= \sum_{v_{t+1}} \Pr(v_{t+1}|v_t, z_{t+1}) \left(\frac{\varphi(v_{t+1}, \gamma(v_t, x_t, z_t, X_t), z_{t+1}, \Gamma(z_t, X_t))}{\varphi(v_t, x_t, z_t, X_t)} \right)^{-\sigma} \\ &= f_{\beta_i}(X_t, z_t, z_{t+1}).\end{aligned}$$

Similarly, we can see that in a recursive competitive equilibrium $\omega_{t+1} = f_{\omega}(X_t, z_t, z_{t+1})$ and $\omega^{\text{cm}}_{t+1} = f_{\omega^{\text{cm}}}(X_t, z_t, z_{t+1})$.

Letting $\hat{y}_t$ be the log-deviation of variable y_t from its steady state, we can express the law of motion of the wedges $\mathbf{T}_{t+1} = [\hat{\beta}_{1,t+1}, \hat{\beta}_{2,t+1}, \ldots, \hat{\omega}_{t+1}, \hat{\omega}^{\text{cm}}_{t+1}]'$, up to a first-order approximation, as

$$\mathbf{T}_{t+1} = \mathbf{A} \times \hat{X}_t + \mathbf{B} \times \hat{z}_t + \mathbf{C} \times \hat{z}_{t+1}, \tag{31}$$

where the matrices $[\mathbf{A}, \mathbf{B}, \mathbf{C}]$ are functions of the primitives of the model.

While equation (31) is effectively a first-order approximation to the true law of motion for $\mathbf{T}_{t+1}$, there is a practical hurdle in using it in our application: certain elements of z_t and X_t may not be defined in the RA representation. To explain this issue, let us go back to the example discussed earlier. In that heterogeneous-agent economy, the distribution $\Psi(e, b)$ is a state variable; however, it does not directly appear in the RA representation.[17] In view of this issue, let us partition z_t and X_t as follows: $z_t = [z_t^{RA}, z_t^{HA}]$ and $X_t = [X_t^{RA}, X_t^{HA}]$, where (z_t^{RA}, X_t^{RA}) denote the aggregate states that are also defined in the equivalent RA economy and (z_t^{HA}, X_t^{HA}) are aggregates states in the heterogeneous-agent economy that do not directly appear in the RA representation. We then approximate equation (31) with

$$\mathbf{T}_{t+1} = \Phi(\mathbf{L}) \times \mathbf{T}_t + \mathbf{A} \times \hat{X}_t^{RA} + \mathbf{B} \times \hat{z}_t^{RA} + \mathbf{C} \times \hat{z}_{t+1}^{RA} + \varepsilon_{t+1}, \tag{32}$$

where $\mathbf{\Phi(L)}$ is a polynomial in the lag operator and ε_{t+1} are innovations with variance-covariance matrix $\mathbf{\Sigma}$. Essentially, our approach consists in proxying for the missing state variables in equation (31) with lagged values of $\mathbf{T}_{t+1}$.[18]

IV.C. Monte Carlo Analysis

In Section IV.A and IV.B we discussed how to measure the wedges using panel data, and we proposed a law of motion to capture their behavior. We made two main approximations in this process. First, the discount factor wedge is computed by averaging the changes in consumption shares of different households with similar characteristics rather than using different histories of the same individual. Second, we cannot condition on all the relevant state variables when considering the law of motion for the wedges because some of these state variables may not be defined in the RA representation. In Online Appendix C we study whether these approximations work well in practice by performing a Monte Carlo analysis on data simulated from benchmark incomplete-markets economies.

We consider two economies, the Krusell and Smith (1998) economy and the Guerrieri and Lorenzoni (2017) economy. In both cases, we solve the model and simulate 500 panel data sets comparable to the one we use in our application to the U.S. economy—10,000 households for 25 years. For each panel

17. That is, in the RA representation the distribution $\Psi(e, b)$ affects aggregate variables through its effect on the wedges.

18. These approximations are common in macroeconomics. For example, they are used when deriving the vector autoregressive representation of dynamic stochastic general-equilibrium models, see Fernández-Villaverde et al. (2007).

data set, we proceed in three steps. First, we use households' observations to construct the wedges following the same approach in Section IV.A. Second, we use the realization of the wedges to estimate the stochastic process in equation (32) with just one lag in $\mathbf{\Phi(L)}$. Third, given the estimated stochastic process for the wedges, we solve for the behavior of aggregate variables using the RA representation. Our test consists in comparing moments for aggregate variables computed using the RA representation with the actual moments from the heterogeneous-agent economy.

Online Appendix C describes in details these steps and provide the results of these experiments. There, we show that our approach to measure the wedges and to approximate their stochastic process works extremely well for both economies and that the moments computed using the RA representation are identical to the ones of the true underlying economy with heterogeneous agents.

V. An Application to the U.S. Economy

We now apply our framework to U.S. data. In Section V.A we use the Consumer Expenditure Survey (CE) to measure the wedges. In Section V.B we jointly estimate their stochastic process and the structural parameters of the RA representation. Section V.C uses the estimated model to assess the effect of imperfect risk sharing for the U.S. business cycle, and Section V.D performs an event study of the Great Recession. Section V.E concludes by discussing different economic mechanisms that can explain the fluctuations in the wedges that we document in the data.

V.A. Measuring the Wedges

We use the CE to collect information on income, expenditures, employment outcomes, wealth, and demographic characteristics for U.S. households between 1992 and 2017. Households in the CE report information on consumption expenditures for a maximum of four consecutive quarters, income and employment information are collected in the first and last interview, and wealth information in the last interview only.[19] Online Appendix D provides details on variable definitions and sample selection, presents

19. The CE asks questions about how assets and liabilities have changed in the preceding year, which allows us to backdate wealth information. See https://www.bls.gov/opub/mlr/2012/05/art3full.pdf.

summary statistics of the underlying micro data, and compares them to previous studies.

The model of Section II abstracts from important determinants of consumption and income, for example demographics. To improve the mapping between model and data, we use panel regressions to partial out the effects of these possible confounders. Denoting by $\hat{c}_{i,t}$ the log of consumption expenditures for household i, we estimate the following linear equation

$$\hat{c}_{i,t} = \alpha + \gamma' X_i + e_{i,t},$$

where X_i includes dummies for the sex, race, education, age of the head of household, and the state of residence. Our measure of residualized consumption is then $c_{i,t} = \exp\{\alpha + e_{i,t}\}$. We repeat this procedure for all variables used in the analysis. All monetary variables are converted into 2000 dollars using the CPI-U and are reported in per capita terms.

One concern with our analysis is that measurement errors may affect the computation of the wedges. We perform five steps to mitigate this concern. First, we winsorize the variables used in the construction of the wedges at the top and bottom 1% (year-by-year) to correct for reporting mistakes that could result in extreme outliers. Second, we follow Vissing-Jørgensen (2002) and use semiannual changes when computing $\frac{c_{i,t+1}}{c_{i,t}}$ to minimize time aggregation and category switching concerns. Third, we obtain the change in households' consumption shares, $\frac{\varphi_{i,t+1}}{\varphi_{i,t}}$, by scaling $\frac{c_{i,t+1}}{c_{i,t}}$ with corresponding semiannual changes in aggregate consumption $\frac{C_{t+1}}{C_t}$ from U.S. national accounts—arguably less subject to measurement errors.[20] Fourth, we de-mean the wedges, a step that helps deal with the presence of classical measurement errors.[21] Fifth, we explicitly model a measurement error for the wedges when estimating their stochastic process.

20. Our results are comparable when we instead compute aggregate consumption using the cross-sectional average of consumption expenditures in the CE; see Online Appendix D.5.

21. To understand why de-meaning helps, consider the construction of the discount factor wedge. Suppose that observed consumption is related to the true consumption as follows, $c_{i,t} = c_{i,t}^{\text{true}} \times \exp\{\eta_{i,t}\}$, where $\eta_{i,t} \sim \mathcal{N}(-\frac{\sigma_\eta^2}{2}, \sigma_\eta^2)$ is a classical measurement error. Simple calculations, then, show that $\beta_{g,t}$ defined in equation (28) converges to $\beta_{g,t}^{\text{true}} \times \exp\{\sigma_\eta^2\}$ when N_g is large. So de-meaning $\beta_{g,t}$ (in logs) removes the bias due to measurement error. A similar derivation can be done for the labor disutility wedge.

TABLE I
Summary Statistics for the Discount Factor Wedge

	Mean ($\hat{\beta}_{g,t}$)	Corr ($\hat{\beta}_{g,t}, \hat{\beta}_{g,t-1}$)	Stdev ($\hat{\beta}_{g,t}$)	Corr ($\hat{\beta}_{g,t}, \hat{Y}_t$)
Low income/low net worth	– 0.03	0.38	0.04	0.12
Low income/high net worth	– 0.02	0.63	0.05	0.25
High income/low net worth	0.00	0.53	0.04	– 0.15
High income/high net worth	0.01	0.33	0.04	0.12

Notes. For each t, we partition households into four groups depending on their income and net worth at date t as described in Section IV.A. We then compute for each household the inverse change in its consumption share between t and $t+1$, $(\frac{\varphi_{i,t+1}}{\varphi_{i,t}})^{-1}$. For each group we then compute $\beta_{g,t+1}$ using equation (28). We scale each $\beta_{g,t+1}$ by the sample average of the high-income/low net worth group, and report the normalized wedges in logs. See Online Appendix D for the definition of detrended output $\hat{Y}_t$.

To measure the wedges, we set the coefficient of relative risk aversion and the Frisch elasticity of labor supply to one, $\sigma = \psi = 1$, conventional values in the literature.

1. *The Discount Factor Wedge.* We measure $\beta_{g,t}$ following the approach described in Section IV.A. Each of the four groups of households contains approximately 875 observations per year. For each household in group g, we compute the semiannual change in consumption expenditures following Vissing-Jørgensen (2002) and scale it by the semiannual change in aggregate consumption over the same horizon. We square the resulting ratio and obtain an empirical analog to the yearly change in consumption shares, $\frac{\varphi_{i,t+1}}{\varphi_{i,t}}$. The discount factor wedge for group g is constructed by averaging $(\frac{\varphi_{i,t+1}}{\varphi_{i,t}})^{-1}$ across the households in that group, see equation (28).

Table I reports summary statistics of the discount factor wedge for the four groups. We report the wedges in log deviations from the sample average of the third group—the high-income/low net worth households. There are two features of the data that we wish to emphasize. First, high-income households have on average a higher discount factor wedge than low-income households. Focusing on households with low net worth, the discount factor wedge of the high-income group is 300 basis points higher than the discount factor wedge of the the low-income group. Second, the discount factor wedge for the high-income/low net worth group is countercyclical, with a correlation coefficient of −0.15 with U.S. detrended output, whereas it is procyclical for the other groups.

The fact that high-income/low net worth households have a relatively high and countercyclical discount factor wedge makes

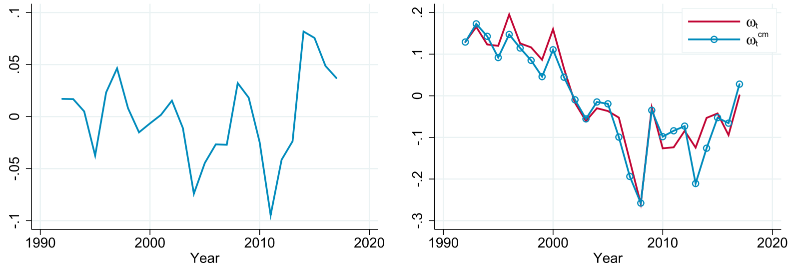

(A) Discount factor wedge

(B) Labor disutility wedge

FIGURE II

The Time Path of the Wedges

Authors' calculation based on the the CE. Panel A plots $\log(\beta_{g,t})$ defined in equation (28) for the group of households with high-income/low net worth. Panel B plots $\log(\omega_t)$ and $\log(\omega_t^{\text{cm}})$, defined respectively in equations (29) and (30). The series are normalized to have a sample average of zero.

it a natural choice to be a group of financially unconstrained households—standing in for households that attain the maximum in equation (21). Indeed, we have that

$$\max_i \mathbb{E}_t\left[\beta_{i,t+1}\frac{C_t}{C_{t+1}}\right] = \max_i \left\{\mathbb{E}_t[\beta_{i,t+1}]\mathbb{E}_t\left[\frac{C_t}{C_{t-1}}\right] + \text{Cov}_t\left(\beta_{i,t+1}, \frac{C_t}{C_{t+1}}\right)\right\}.$$

So groups that on average have a high and countercyclical discount factor wedge are likely to be those achieving the maximum in equation (21). This empirical finding is in line with standard incomplete-market models: households that experience positive income shocks and that have low net worth have the highest incentives to save, and thus are not financially constrained. Because of these reasons, we use $\beta_{g,t}$ of this group as the discount factor wedge in our analysis.

Figure II, Panel A plots the time series of the discount factor wedge. We can see that the discount factor wedge increases substantially during the Great Recession and remains elevated in the post-2010 period.

2. *The Labor Disutility Wedge.* We compute for each household in our panel the consumption share $\varphi_{i,t}$ and the ratio between

the household's hourly wage and the average hourly wage in the panel, $e_{i,t} = \frac{w_{i,t}}{W_t}$. We then use equations (29) and (30) to obtain the time path of ω_t and ω_t^{cm}.

Figure II, Panel B plots these two series. We can see for both series a clear downward trend in the early part of the sample. The downward trend is explained by the increase in the cross-sectional variance of wages for U.S. households during this period, a fact that is well established in the literature (Heathcote, Perri, and Violante 2010). As we explained in Section III.B, an increase in the cross-sectional dispersion of labor productivity leads high-productivity workers to increase their labor supply relative to low-productivity workers because of substitution effects. This mechanism operates irrespective of whether financial markets are complete, and the resulting change in the composition of the labor force that takes place in the heterogeneous-agent economy is captured in the RA representation by a decline in the labor disutility wedge. So both ω_t and ω_t^{cm} fall for a large part of the sample. We can also observe from the figure that the deviations between ω_t and ω_t^{cm} are typically small. This means that the wealth effects in labor supply that are due to uninsured idiosyncratic income risk—which our framework captures with a discrepancy between ω_t and ω_t^{cm}—are quite small.

V.B. Estimating the RA Representation

We estimate the stochastic process of the wedges and the structural parameters of the RA representation. The structural shocks in the RA representation are $z_t^{RA} = [\hat{\theta}_t, \hat{A}_t, \varepsilon_{m,t}]$. We assume that the aggregate preference and technology shocks in logs follow independent AR(1) processes,

$$\hat{\theta}_t = \rho_\theta \hat{\theta}_{t-1} + \varepsilon_{\theta,t} \qquad \varepsilon_{\theta,t} \sim \mathcal{N}\big(0, \sigma_\theta^2\big)$$

$$\hat{A}_t = \rho_a \hat{A}_{t-1} + \varepsilon_{A,t} \qquad \varepsilon_{A,t} \sim \mathcal{N}\big(0, \sigma_A^2\big),$$

and that monetary policy innovations are Gaussian, $\varepsilon_{m,t} \sim \mathcal{N}(0, \sigma_m^2)$. The endogenous state variables are the nominal interest rate and the wedges, $\mathbf{T}_t = [\hat{\beta}_t, \hat{\omega}_t, \hat{\omega}_t^{\mathrm{cm}}]'$. The law of motion for $\mathbf{T}_t$ is given by equation (32). We restrict $\mathbf{\Phi}(L)$ to have a one-lag structure and be block diagonal, so that ω_t^{cm} depends only on its own lags and does not load on the other equations. The other parameters in equation (32) are left unrestricted.

We estimate the parameters of the stochastic processes jointly with the other structural parameters governing preferences, technology, and the behavior of the monetary authority. We fix a subset of these parameters to conventional values in the literature. Consistent with the measurement of the wedges, we set $\sigma = \psi = 1$. We let $\mu = 1.2$ and set χ to $\frac{1}{\mu}$, so that consumption and output are normalized to 1 in a deterministic steady state of the model. Finally, we set the target inflation rate to 2%, and $\beta = 0.99$, values that guarantee that the model roughly matches the average inflation and nominal interest rate in our sample.

The remaining model parameters, which we collect in the vector Θ, are estimated with Bayesian methods using annual data on output, inflation, nominal interest rates, and the wedges. We map the log of output, $\hat{Y}_t$, to detrended log real GDP, the inflation rate π_t to the annual percent change in the Consumer Price Index, and nominal interest rates i_t to the annual effective Federal Funds rate; see Online Appendix D for definitions, data sources, and detrending methodology. The sample period is 1992–2017. We denote by $\mathbf{Y}_t = [\hat{Y}_t, i_t, \pi_t, \mathbf{T}_t]$ the observables at time t and by $\mathbf{S}_t = [i_{t-1}, \hat{\theta}_t, \hat{A}_t, \varepsilon_{m,t}, \mathbf{T}_t]$ the state vector. The RA representation of Proposition 1 defines implicitly a law of motion for these vectors,

$$\begin{aligned} \mathbf{Y}_t &= \mathbf{g}(\mathbf{S}_t; \Theta) + \eta_t \\ \mathbf{S}_t &= \mathbf{f}(\mathbf{S}_{t-1}, \varepsilon_t; \Theta), \end{aligned} \tag{33}$$

where $\mathbf{g}(.)$ and $\mathbf{f}(.)$ represent the policy functions of the RA representation, ε_t collects the innovations to the stochastic variables of the model, and η_t are Gaussian measurement errors. We introduce measurement errors only for the wedges and fix their variance to 10% of the unconditional variance of these series. For estimation, the policy functions are approximated using a first-order perturbation. When performing the counterfactuals, however, we solve the model with global solution methods that allow for the possibility of a binding zero lower bound constraint on nominal interest rates; see Online Appendix E.[22]

The posterior distribution of Θ is characterized using a random walk metropolis Hastings algorithm; see An and Schorfheide (2007). Online Appendix F provides details on the estimation

22. We estimate a log-linearized version of the model because the numerical solution is much faster and numerically more stable than the global approximation that we use for the counterfactuals. See Aruoba, Cuba-Borda, and Schorfheide (2018) for a similar approach.

algorithm and reports parameters' estimates and indicators of model fit. Online Appendix Tables A-4 and A-5 report prior and posterior statistics for the model parameters. Our estimates for the parameters of the Taylor rule and the price adjustment costs are in line with previous studies. In addition, we estimate some spillovers between the model state variables and the wedges. Recessionary shocks at time t—such as a negative technology shock or a positive aggregate discount factor shock—forecast an increase in the discount factor wedge at time $t+1$. This result is consistent with heterogeneous-agent economies where households' precautionary saving motives are more prevalent in recessions. Online Appendix Figure A-5 reports model-implied distributions for the mean, standard deviation, autocorrelation, and cross-correlation of output, inflation, and nominal interest rates and compare these to the corresponding sample moments. We can see from the figure that the estimated model fits reasonably well the behavior of these series, as sample statistics from the data lie within the corresponding model-implied distribution.

V.C. Imperfect Risk Sharing and the Business Cycle

In this and the next subsection, we compare the estimated RA representation to an economy that is identical to the RA representation with the exception that $\hat{\beta}_t = 0$ and $\hat{\omega}_t = \hat{\omega}_t^{\text{cm}}$. As we showed in Proposition 2, if we set the wedges to those values in the RA representation we would recover the path of aggregate variables in the presence of complete financial markets. So we refer to this economy as the complete-markets counterfactual (CM). By comparing these economies, we can assess the role of imperfect risk sharing for macroeconomic aggregates. The analysis focuses on comparing the business cycle properties of the two economies and not their long-run average behavior.[23]

We set the model parameters at the posterior mean, solve numerically for the RA representation and the CM, and compute long simulations for both economies. Table II reports key statistics for output, inflation, and nominal interest rates. If shocks

23. Our approach is not designed to study this latter question because, by construction, the CM and the RA representation have the same steady state. This happens because we normalize the wedges to have a mean of 1. In Online Appendix F.4 we show that this restriction does not affect our key results, as the business cycle properties of the RA representation and the CM counterfactual are almost invariant to the steady-state value of the wedges.

TABLE II
IMPERFECT RISK SHARING AND THE BUSINESS CYCLE

	RA	CM	CM, $\omega_t^{\text{cm}} = \omega_t$	RA, $\sigma_\beta = 0$
Stdev($\hat{Y}_t$)	4.05	3.79	3.82	3.99
Stdev(π_t)	1.39	0.96	0.94	1.36
Stdev(i_t)	2.20	1.90	1.84	2.17
Corr($\hat{Y}_t, \pi_t$)	0.45	0.12	0.22	0.42
Corr($\hat{Y}_t, i_t$)	0.17	−0.36	−0.11	0.13
Corr(π_t, i_t)	0.51	0.33	0.31	0.49

Notes. Each column reports statistics computed from simulations of four different economies: (i) the RA representation; (ii) the counterfactual economy with complete markets; (iii) the counterfactual economy with complete markets but $\omega_t^{\text{cm}} = \omega_t$; (iv) the RA representation with $\sigma_\beta = 0$. Statistics for each version of the model are computed on a long simulation ($T = 50{,}000$).

and frictions at the micro level were an important source of business cycle fluctuations, we should expect the CM to display significantly less volatile output than what we find in the RA representation. We can see from Table II that the standard deviation of output in CM is approximately 7% smaller relative to that of the RA representation. In addition, by comparing the cross-correlation patterns in the economies, we can see that imperfect risk sharing acts mostly as a "demand" shock on the economy: output, inflation, and nominal interest rates are much more positively correlated in the RA representation than in the economy with complete financial markets.

To understand these results, it is useful to note that there are two key differences between the RA representation and CM: the latter does not have a discount factor wedge, and it has a different labor disutility wedge. In an effort to understand which feature drives the results, we report in the third column in Table II moments computed from the CM economy but condition the realization of the labor disutility wedge to be the same as in the RA representation, $\omega_t^{\text{cm}} = \omega_t$. By comparing the second and third columns, we can see that these versions of the CM economy feature almost identical business cycle properties. Therefore, the differences between RA and CM are mostly driven by the movements in the discount factor wedge that are present in the former but not in the latter. This helps explain why the correlation between output, inflation, and nominal interest rates drops in CM, as shocks to the discount factor in the canonical three-equations New Keynesian model induce positive comovement between these variables.

These results suggest that movements in the discount factor wedge are key for understanding the aggregate implications of imperfect risk sharing. An interesting question is whether these fluctuations are mostly induced by the structural shocks, [A_t, θ_t, $\varepsilon_{m,t}$], or whether they are due to the innovations $e_{\beta,t}$. The fourth column of Table II helps us answer this question. There, we report statistics from the RA economy when the variance of $e_{\beta,t}$ is set to zero, $\sigma_\beta = 0$, so all the movements in the discount factor wedge are due to spillovers from the structural shocks. We can see that the sample moments in this economy are comparable to those of the benchmark RA representation, implying that most of the aggregate effects of imperfect risk sharing arise because incomplete markets amplify the effects of structural shocks.

The results of this section suggest that the main channel through which micro heterogeneity matters for aggregate variables is equivalent to movements in the discount factor. In the standard New Keynesian model, the aggregate implications of shocks to the discount factor critically depend on the response of the monetary authority. For a fixed nominal interest rate, an increase in the discount factor induces the stand-in household to cut consumption. If the monetary authority responds by lowering nominal interest rates, however, the increase in patience has more limited effects on consumption and output. Given our estimates of the Taylor rule, the response of the monetary authority is strong enough to offset most of the output effects of the measured changes in the discount factor wedge.

Importantly, this result depends on the response of the monetary authority. It is well known in the literature that changes in the discount factor can have substantial output effects when the zero lower bound constraint on nominal interest rates binds, as in that case the monetary authority cannot cut nominal interest rates further. These events are somewhat rare in the estimated model. Thus, while the results of Table II indicate limited output effects of imperfect risk sharing on average over the business cycle, they do not rule out that these frictions may play a more important role in periods during which the monetary authority is constrained by the zero lower bound. In the next subsection, we explore this possibility with an event study of the U.S. Great Recession.

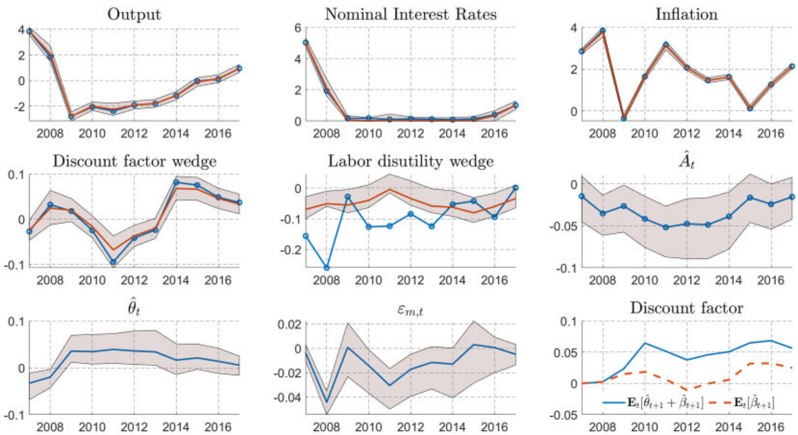

FIGURE III

Applying the Particle Filter to the RA Representation

The circled line reports the data used in the experiment. The solid lines report the posterior mean of the filtered series for $\mathbf{Y}_t$ and $\mathbf{S}_t$, while the shaded area reports 90% credible sets (equal tail probability). Inflation and nominal interest rates are reported in percentage points, and output is reported in percentage point deviations from a quadratic time trend. The remaining variables are in log deviation from their steady-state value.

V.D. Imperfect Risk Sharing and the Great Recession

We use the estimated model to measure the macroeconomic effects of imperfect risk sharing during the Great Recession. To this end, we proceed in two steps. In the first step, we apply the particle filter to the state-space system of equations (33) and estimate the realization of the structural shocks. In the second step, we feed the structural shocks into the CM economy to obtain the counterfactual paths for output, inflation, and nominal interest rates under complete financial markets. The difference between what we observe in the data and these counterfactual paths isolates the macroeconomic effects of imperfect risk sharing during the Great Recession. Online Appendix F provides a detailed description of these steps.

Starting with the first step of this procedure, Figure III reports the data (circled lines) along with the posterior mean (solid line) and 90% credible set for their model counterpart. The figure also reports the estimates for the three structural shocks. By construction, the model tracks very closely output, inflation, and nominal interest rates during the event. To replicate these

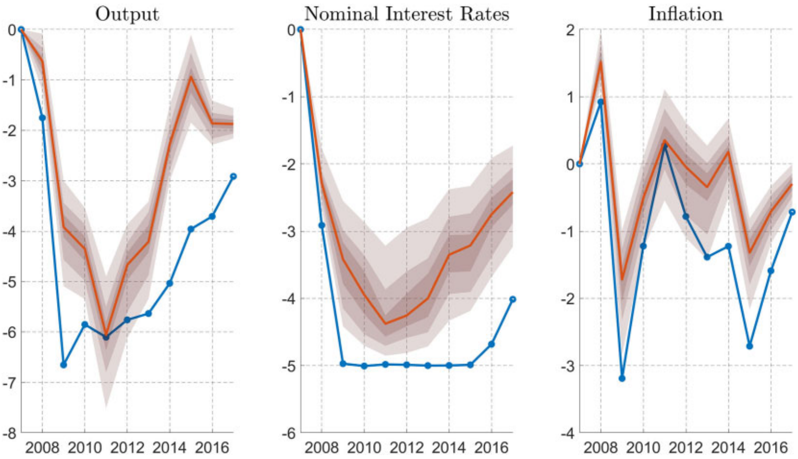

FIGURE IV

Imperfect Risk Sharing and the U.S. Great Recession

The circled line reports output, nominal interest rates, and inflation during the 2007–2017 period. Inflation and nominal interest rates are reported in percentage points, and output is reported in percentage points deviations from a quadratic time trend. We normalize these variables to zero in 2007. The solid line reports the posterior mean of the same variable in the counterfactual economy with complete financial markets while the shaded area reports 64% credible sets (equal tail probability).

paths, the model infers a substantial increase in the discount factor of the stand-in household: $\mathbb{E}_t[\hat{\beta}_{t+1} + \hat{\theta}_{t+1}]$ increases by 6 percentage points during the event. This is a well-known result from the literature: the canonical three-equations New Keynesian model needs an increase in the discount factor to fit the fall in output, inflation, and nominal interest rates observed after 2008. Interestingly, a substantial fraction of this increase is due to movements in $\mathbb{E}_t[\hat{\beta}_{t+1}]$, roughly 3 of the 6 percentage points by the end of the episode.

Equipped with the path for the structural shocks, we can construct the trajectories for output, nominal interest rates, and inflation that would prevail in an economy with complete financial markets. We solve numerically for the policy functions of the CM economy and construct the counterfactual paths by feeding these policy functions with the path for $\{\hat{\theta}_t, \hat{A}_t, \varepsilon_{m,t}\}$ we estimated using the particle filter and with $\hat{\omega}_t^{\text{cm}}$.

Figure IV compares the trajectories for output, inflation, and nominal interest rates in this counterfactual (solid lines) with the actual trajectories in U.S. data (circled lines) during 2007–2017. From 2007 to 2009, output in the United States fell by 6.6%. In the counterfactual economy, instead, output falls by only 3.9%. In addition, the figure shows that the recovery from the Great Recession would have been faster with complete financial markets relative to what we have observed in the data. Thus, our analysis supports the view that these financial market frictions were an important driver of the depth and persistence of the U.S. Great Recession.

Why do we observe these differences between the baseline and the counterfactual economy? From Figure III we can see that $\mathbb{E}_t[\hat{\beta}_{t+1}]$ contributes substantially to the increase in the overall discount rate during the Great Recession, and because interest rates are at their lower bound from 2009 onward, these developments lead to a substantial decline in aggregate demand. In the CM economy, instead, $\hat{\beta}_t$ is constant and does not increase during this period, implying a higher trajectory for output, inflation, and nominal interest rates.

V.E. Inspecting the Mechanisms

Although our results indicate that shocks and frictions limiting risk sharing across households contributed significantly to the depth and persistence of the Great Recession, so far we have been silent about the specific economic mechanisms responsible for this result. Here we use additional information to understand which economic mechanisms can account for our results and, specifically, for the rise in the discount factor wedge.

A first useful exercise consists in understanding which moments of the distribution of households' consumption shares are responsible for the observed movements in the discount factor wedge. From equation (28), we can decompose the log of $\beta_{g,t+1}$ for households in group g as follows

$$\log\left(\beta_{g,t+1}\right) = \underbrace{\Delta\log(C_{t+1}) - \Delta\log(C_{g,t+1})}_{\beta^{\text{Avg}}_{g,t+1}} + \underbrace{\log\left(\frac{1}{N_g}\sum_{i=1}^{N_g}\left[\frac{C_{g,t+1}}{C_{g,t}}\frac{c_{i,t}}{c_{i,t+1}}\right]\right)}_{\beta^{\text{Jen}}_{g,t+1}}, \tag{34}$$

where $C_{g,t}$ is the average consumption at time t of households in group g. Mechanically, there are two factors that can explain the increase in the discount factor wedge during the Great Recession. First, a persistent fall in the average consumption of households in the high-income/low net worth group relative to aggregate consumption, the term $\beta_{g,t}^{\text{Avg}}$ in equation (34). Second, a increase in the cross-sectional dispersion of consumption within group g, an effect captured by the "Jensen" term $\beta_{g,t+1}^{\text{Jen}}$. These two effects are intuitive: fixing the interest rate, a household that expects low and/or volatile consumption in the future has more incentives to save today, an effect that is captured in the RA representation by a higher discount factor wedge.

This decomposition is useful for model discrimination because two models that generate similar dynamics for $\beta_{g,t}$ may have different implications for the two terms in the right-hand side of equation (34). For example, take the "two-agent" New Keynesian model studied in Bilbiie (2008) and Galí, López-Salido, and Vallés (2007). As showed in Debortoli and Galí (2017), these models feature similar macroeconomic dynamics relative to New Keynesian models with richer heterogeneity, such as Kaplan, Moll, and Violante (2018). This means that those models imply a similar behavior for the discount factor wedge $\beta_{g,t}$. However, the two-agent New Keynesian model does not feature idiosyncratic consumption risk for unconstrained households, which implies that $\beta_t^{\text{Jen}} = 1$ and all movements in the discount factor wedge are induced by changes in β_t^{Avg}. Models with richer heterogeneity can instead generate time variation in β_t^{Jen}.

Figure V reports $\log(\beta_{g,t})$ and $\beta_{g,t}^{\text{Jen}}$ for the high-income/low net worth group, with $\beta_{g,t}^{\text{Avg}}$ being the difference between these two series. We can see that the dynamics of the discount factor wedge during the Great Recession are driven almost entirely by the Jensen component.[24] Therefore, to understand the rise in the discount factor wedge during this period, we need to focus on economic mechanisms that can generate an increase in the consumption risk of unconstrained households.

The existing literature has emphasized two main mechanisms that can lead to this pattern, which we discussed in the two examples of Section III.C. In the first example, we showed that an increase in idiosyncratic labor income risk can lead to higher consumption risk and more precautionary savings. In the

24. This is true also for the other group of households that are likely to be financially unconstrained, the high-income/high net worth households.

FIGURE V

Decomposition of equation (34)

This figure plots the decomposition described in equation (34) for high-income/low net worth households. The solid line reports $\log(\beta_{g,t})$ and the circled line reports β_t^{Jen}. The variables are expressed in deviations from their 2007 value.

second example, we showed that similar effects can arise after a persistent tightening of households' borrowing constraints. Both forces can rationalize an increase of the Jensen term in equation (34) and be consistent with the findings of Figure V.

To distinguish between these forces, it is useful to consider the following reduced-form model of households' consumption growth:

$$\Delta \hat{c}_{i,t} = \tau_t + \lambda_t \Delta \hat{y}_{i,t} + e_{i,t}, \tag{35}$$

where $\hat{y}_{i,t}$ is the log of households' disposable income at date t. In this framework, the conditional variance of consumption growth depends on the volatility of income and on the sensitivity of consumption to income changes, λ_t. It is straightforward to see that the two examples described in Section III.C have different implications for these two components. In the first example, we have that $\lambda_t = 1$ for all t; therefore, the volatility of consumption varies over time only because of changes in the volatility of income. In the second example, instead, income volatility is constant over time while shocks to households' borrowing constraint affect the sensitivity

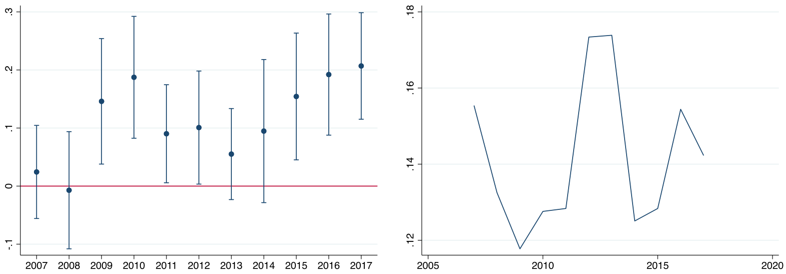

(A) Estimates of λ_t (B) $\text{Var}(\Delta\hat{y}_{i,t})$

FIGURE VI

Drivers of Households' Consumption Volatility, 2007–2017

For each t, we select households that at $t-1$ are in the group of high-income/low net worth households and estimate equation (35) by ordinary least squares. Panel A plots the point estimate for λ_t and 90% robust confidence interval. Panel B plots the cross-sectional variance of $\Delta\hat{y}_{i,t}$ for each t.

of their consumption to income shocks; therefore, consumption volatility changes over time only because of movements in λ_t.

We estimate equation (35) for the high-income/low net worth households and report the results in Figure VI. Panel A plots the estimated coefficient λ_t over time, and Panel B plots the variance of $\Delta\hat{y}_{i,t}$ for this group of households. From the figure we can see that λ_t tracks well the dynamics of $\beta^{\text{Jen}}_{g,t}$ reported in Figure V. The cross-sectional variance of income for these households, instead, does not display a substantial increase after 2007.[25]

Overall, these results suggest that the increase in the discount factor wedge during the Great Recession is mostly the result of an increase in consumption volatility of unconstrained households, and this was mostly driven by an increase in the sensitivity of their consumption to income changes. When

25. This result is consistent with previous research that examined the cyclical behavior of income changes. While Storesletten, Telmer, and Yaron (2004) found evidence of countercyclical income volatility, their findings are not comparable to those of Figure VI because we are conditioning on high-income/low net worth households while Storesletten, Telmer, and Yaron (2004) use their entire sample in their analysis. Consistent with our findings, Heathcote, Perri, and Violante (2010) find little cyclical variation in the volatility of earnings growth for households between the 50th and 90th percentiles of the earnings distribution. These authors document significant cyclical variation in income volatility for poorer households due to a higher incidence of unemployment during recessions, a result confirmed by Guvenen, Ozkan, and Song (2014).

interpreted through the lens of the examples of Section III.C, this evidence favors economic models that emphasize a disruption of households' ability to insure against fluctuations in their income during the Great Recession, as in Guerrieri and Lorenzoni (2017). It is worth noting, however, that this interpretation is somewhat special to these examples. In general, changes in the income process may affect the sensitivity of consumption to income—λ_t in equation (35)—and potentially account for the dynamics reported in Figure VI. For example, suppose that the Great Recession was accompanied by an increase in the variance of the permanent component of income and a reduction of its transitory component.[26] Such a shift could explain why in Figure VI we detect an increase in λ_t in the absence of significant movements in $\text{Var}(\Delta\tilde{y}_{i,t})$. This is because permanent shocks are harder to smooth than transitory ones in baseline incomplete-markets model, so consumption is more sensitive to the former than the latter.

VI. Conclusion

We have developed an approach to assess the macroeconomic implications of imperfect risk sharing implied by a class of New Keynesian models with heterogeneous agents. In this class of models, households' inability to insure idiosyncratic risk implies time variation in their consumption shares. Leveraging this insight, we use households' consumption choices to directly measure the degree of imperfect risk sharing for the U.S. economy. We documented a deterioration of risk sharing during the Great Recession, as the cross-sectional dispersion of households' consumption shares increases during this period. We proposed a methodology to quantify the aggregate implications of these movements. Through the lens of a prototypical New Keynesian model, these deviations from perfect risk sharing contribute little to the business cycle on average but explain a substantial part of the depth and persistence of the Great Recession.

This article clarifies that assumptions about the nature of the idiosyncratic risk faced by households and about the private and public risk-sharing mechanisms available to them matter for aggregate fluctuations only through their effect on two wedges.

26. This hypothesis is similar to the mechanism in Blundell, Pistaferri, and Preston (2008) to account for the changes in the consumption distribution relative to the income distribution over the 1970s to 1990s.

These wedges depend on the joint distribution of households' consumption shares and relative wages, and they can be computed using panel data. We believe that they are useful empirical targets for the analysis of heterogeneous-agent economies.

Duke University, United States
Stanford University, United States
University of Pennsylvania, United States

Supplementary Material

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online.

Data Availability

The data underlying this article are available in the Harvard Dataverse, https://doi.org/10.7910/DVN/Y0F3IX (Berger, Bocola, and Dovis 2023).

References

Acharya, Sushant, and Keshav Dogra, "Understanding HANK: Insights from a PRANK," Federal Reserve Bank of New York Staff Report 835, 2018.

Alvarez, Fernando, and Urban J. Jermann, "Efficiency, Equilibrium, and Asset Pricing with Risk of Default," *Econometrica*, 68 (2000), 775–797. https://doi.org/10.1111/1468-0262.00137.

An, Sungbae, and Frank Schorfheide, "Bayesian Analysis of DSGE Models," *Econometric Reviews*, 26 (2007), 113–172. https://doi.org/10.1080/07474930701220071.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare, "New Trade Models, Same Old Gains?," *American Economic Review*, 102 (2012), 94–130. https://doi.org/10.1257/aer.102.1.94.

Aruoba, S. Borağan, Pablo Cuba-Borda, and Frank Schorfheide, "Macroeconomic Dynamics Near the ZLB: A Tale of Two Countries," *Review of Economic Studies*, 85 (2018), 87–118. https://doi.org/10.1093/restud/rdx027.

Auclert, Adrien, "Monetary Policy and the Redistribution Channel," *American Economic Review*, 109 (2019), 2333–2367. https://doi.org/10.1257/aer.20160137.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub, "The Intertemporal Keynesian Cross," Stanford University Technical report, January 2017.

———, "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," NBER Working Paper No 26647, 2020.

Barro, Robert J., and Robert G. King, "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles," *Quarterly Journal of Economics*, 99 (1984), 817–839. https://doi.org/10.2307/1883127.

Bayer, Christian, Benjamin Born, and Ralph Luetticke, "Shocks, Frictions, and Inequality in US Business Cycles," CEPR Discussion Paper no. DP14364, 2020.

Bayer, Christian, Ralph Lütticke, Lien Pham-Dao, and Volker Tjaden, "Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk," *Econometrica*, 87 (2019), 255–290. https://doi.org/10.3982/ECTA13601.

Berger, David, Luigi Bocola, and Alessandro Dovis, “Replication Data for: ‘Imperfect Risk Sharing and the Business Cycle’,” (2023), Harvard Dataverse. https://doi.org/10.7910/DVN/Y0F3IX.

Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J. Sargent, “Inequality, Business Cycles, and Monetary-Fiscal Policy,” *Econometrica*, 89 (2021), 2559–2599. https://doi.org/10.3982/ECTA16414.

Bilbiie, Florin O., “Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic,” *Journal of Economic Theory*, 140 (2008), 162–196. https://doi.org/10.1016/j.jet.2007.07.008.

———, “The New Keynesian Cross,” *Journal of Monetary Economics*, 114 (2020), 90–108. https://doi.org/10.1016/j.jmoneco.2019.03.003.

Bilbiie, Florin, Giorgio E. Primiceri, and Andrea Tambalotti, “Inequality and Business Cycles,” Unpublished manuscript, 2022.

Bisin, Alberto, Gian Luca Clementi, and Piero Gottardi, “Capital Structure and Hedging Demand with Incomplete Markets,” NBER Working Paper 20345, 2014. https://doi.org/10.3386/w20345.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, “Consumption Inequality and Partial Insurance,” *American Economic Review*, 98 (2008), 1887–1921. https://doi.org/10.1257/aer.98.5.1887.

Brav, Alon, George M. Constantinides, and Christopher C. Geczy, “Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence,” *Journal of Political Economy*, 110 (2002), 793–824. https://doi.org/10.1086/340776.

Broer, Tobias, Niels-Jakob H. Hansen, Per Krusell, and Erik Öberg, “The New Keynesian Transmission Mechanism: A Heterogeneous-Agent Perspective,” *Review of Economic Studies*, 87 (2020), 77–101. https://doi.org/10.1093/restud/rdy060.

Buera, Francisco J., and Benjamin Moll, “Aggregate Implications of a Credit Crunch: The Importance of Heterogeneity,” *American Economic Journal: Macroeconomics*, 7 (2015), 1–42. https://doi.org/10.1257/mac.20130212.

Challe, Edouard, Julien Matheron, Xavier Ragot, and Juan F. Rubio-Ramirez, “Precautionary Saving and Aggregate Demand,” *Quantitative Economics*, 8 (2017), 435–478. https://doi.org/10.3982/QE714.

Chari, Varadarajan V., Patrick J. Kehoe, and Ellen R. McGrattan, “Business Cycle Accounting,” *Econometrica*, 75 (2007), 781–836. https://doi.org/10.1111/j.1468-0262.2007.00768.x.

Chien, YiLi, Harold Cole, and Hanno Lustig, “A Multiplier Approach to Understanding the Macro Implications of Household Finance,” *Review of Economic Studies*, 78 (2011), 199–234. https://doi.org/10.1093/restud/rdq008.

———, “Is the Volatility of the Market Price of Risk due to Intermittent Portfolio Rebalancing?,” *American Economic Review*, 102 (2012), 2859–2896. https://doi.org/10.1257/aer.102.6.2859.

Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo, “When Is the Government Spending Multiplier Large?,” *Journal of Political Economy*, 119 (2011), 78–121. https://doi.org/10.1086/659312.

Constantinides, George M., “Intertemporal Asset Pricing with Heterogeneous Consumers and without Demand Aggregation,” *Journal of Business*, 55 (1982), 253–267. https://doi.org/10.1086/296163.

Debortoli, Davide, and Jordi Galí, “Monetary Policy with Heterogeneous Agents: Insights from TANK Models,” Unpublished manuscript, 2017.

———, “Idiosyncratic Income Risk and Aggregate Fluctuations,” NBER Working Paper No 29704, 2022. https://doi.org/10.3386/w29704.

Den Haan, Wouter J., Pontus Rendahl, and Markus Riegler, “Unemployment (Fears) and Deflationary Spirals,” *Journal of the European Economic Association*, 16 (2018), 1281–1349. https://doi.org/10.1093/jeea/jvx040.

Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, Thomas J. Sargent, and Mark W. Watson, “ABCs (and Ds) of Understanding VARs,” *American Economic Review*, 97 (2007), 1021–1026. https://doi.org/10.1257/aer.97.3.1021.

Galí, Jordi, J. David López-Salido, and Javier Vallés, "Understanding the Effects of Government Spending on Consumption," *Journal of the European Economic Association*, 5 (2007), 227–270. https://doi.org/10.1162/JEEA.2007.5.1.227.

Gertler, Mark, and Peter Karadi, "A Model of Unconventional Monetary Policy," *Journal of Monetary Economics*, 58 (2011), 17–34. https://doi.org/10.1016/j.jmoneco.2010.10.004.

Gertler, Mark, and Nobuhiro Kiyotaki, "Financial Intermediation and Credit Policy in Business Cycle Analysis," in *Handbook of Monetary Economics,* vol. 3, B. M. Friedman and M. Woodford, eds. (Amsterdam: Elsevier, 2010), 547–599. https://doi.org/10.1016/B978-0-444-53238-1.00011-9.

Gornemann, Nils, Keith Kuester, and Makoto Nakajima, "Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy," International Finance Discussion Papers No. 1167, Board of Governors of the Federal Reserve System, 2016.

Guerrieri, Veronica, and Guido Lorenzoni, "Credit Crises, Precautionary Savings, and the Liquidity Trap," *Quarterly Journal of Economics*, 132 (2017), 1427–1467. https://doi.org/10.1093/qje/qjx005.

Guvenen, Fatih, Serdar Ozkan, and Jae Song, "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122 (2014), 621–660. https://doi.org/10.1086/675535.

Hagedorn, Marcus, Iourii Manovskii, and Kurt Mitman, "The Fiscal Multiplier," NBER Working Paper No 25571, 2019. https://doi.org/10.3386/w25571.

Heathcote, Jonathan, and Fabrizio Perri, "Wealth and Volatility," *Review of Economic Studies*, 85 (2018), 2173–2213. https://doi.org/10.1093/restud/rdx074.

Heathcote, Jonathan, Fabrizio Perri, and Gianluca Violante, "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States 1967–2006," *Review of Economic Dynamics*, 13 (2010), 15–51. https://doi.org/10.1016/j.red.2009.10.010.

Jones, Callum, Virgiliu Midrigan, and Thomas Philippon, "Household Leverage and the Recession," *Econometrica*, 90 (2022), 2471–2505. https://doi.org/10.3982/ECTA16455.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante, "Monetary Policy According to HANK," *American Economic Review*, 32 (2018), 697–743. https://doi.org/10.1257/aer.20160042.

Kaplan, Greg, and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82 (2014), 1199–1239. https://doi.org/10.3982/ECTA10528.

———, "Microeconomic Heterogeneity and Macroeconomic Shocks," *Journal of Economic Perspectives*, 32 (2018), 167–194. https://doi.org/10.1257/jep.32.3.167.

———, "The Marginal Propensity to Consume in Heterogeneous Agent Models," *Annual Review of Economics*, 14 (2022), 747–775. https://doi.org/10.1146/annurev-economics-080217-053444.

Kehoe, Timothy J., and David K. Levine, "Debt-Constrained Asset Markets," *Review of Economic Studies*, 60 (1993), 865–888. https://doi.org/10.2307/2298103.

Kocherlakota, Narayana, and Luigi Pistaferri, "Asset Pricing Implications of Pareto Optimality with Private Information," *Journal of Political Economy*, 117 (2009), 555–590. https://doi.org/10.1086/599761.

Krueger, Dirk, and Hanno Lustig, "When Is Market Incompleteness Irrelevant for the Price of Aggregate Risk (and When Is it Not)?," *Journal of Economic Theory*, 145 (2010), 1–41. https://doi.org/10.1016/j.jet.2009.10.005.

Krueger, Dirk, Hanno Lustig, and Fabrizio Perri, "Evaluating Asset Pricing Models with Limited Commitment Using Household Consumption Data," *Journal of the European Economic Association*, 6 (2008), 715–726. https://doi.org/10.1162/JEEA.2008.6.2-3.715.

Krueger, Dirk, Kurt Mitman, and Fabrizio Perri, "Macroeconomics and Household Heterogeneity," in *Handbook of Macroeconomics,* vol. 2, J. B. Taylor and H. Uhlig, eds. (Amsterdam: Elsevier, 2016), 843–921. https://doi.org/10.1016/bs.hesmac.2016.04.003.

Krusell, Per, and Anthony A. Smith Jr., "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106 (1998), 867–896. https://doi.org/10.1086/250034.
Makowski, Louis, "Competition and Unanimity Revisited," *American Economic Review*, 73 (1983a), 329–339. http://www.jstor.org/stable/1808116.
———, "Competitive Stock Markets," *Review of Economic Studies*, 50 (1983b), 305–330. https://doi.org/10.2307/2297418.
Maliar, Lilia, and Serguei Maliar, "Heterogeneity in Capital and Skills in a Neoclassical Stochastic Growth Model," *Journal of Economic Dynamics and Control*, 25 (2001), 1367–1397. https://doi.org/10.1016/S0165-1889(99)00058-5.
———, "The Representative Consumer in the Neoclassical Growth Model with Idiosyncratic Shocks," *Review of Economic Dynamics*, 6 (2003), 362–380. https://doi.org/10.1016/S1094-2025(03)00003-6.
McKay, Alisdair, "Time-Varying Idiosyncratic Risk and Aggregate Consumption Dynamics," *Journal of Monetary Economics*, 88 (2017), 1–14. https://doi.org/10.1016/j.jmoneco.2017.05.002.
McKay, Alisdair, Emi Nakamura, and Jón Steinsson, "The Power of Forward Guidance Revisited," *American Economic Review*, 106 (2016), 3133–3158. https://doi.org/10.1257/aer.20150063.
McKay, Alisdair, and Ricardo Reis, "The Role of Automatic Stabilizers in the U.S. Business Cycle," *Econometrica*, 84 (2016), 141–194. https://doi.org/10.3982/ECTA11574.
Mendoza, Enrique G., "Sudden Stops, Financial Crises, and Leverage," *American Economic Review*, 100 (2010), 1941–1966. https://doi.org/10.1257/aer.100.5.1941.
Nakajima, Tomoyuki, "A Business Cycle Model with Variable Capacity Utilization and Demand Disturbances," *European Economic Review*, 49 (2005), 1331–1360. https://doi.org/10.1016/j.euroecorev.2003.12.002.
Ravn, Morten O., and Vincent Sterk, "Job Uncertainty and Deep Recessions," *Journal of Monetary Economics*, 90 (2017), 125–141. https://doi.org/10.1016/j.jmoneco.2017.07.003.
———, "Macroeconomic Fluctuations with HANK & SAM: An Analytical Approach," *Journal of the European Economic Association*, 19 (2021), 1162–1202. https://doi.org/10.1093/jeea/jvaa028.
Schaab, Andreas, "Micro and Macro Uncertainty," Unpublished manuscript, 2020.
Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112 (2004), 695–717. https://doi.org/10.1086/383105.
Vissing-Jørgensen, Annette, "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy*, 110 (2002), 825–853. https://doi.org/10.1086/340782.
Werning, Iván, "Incomplete Markets and Aggregate Demand," NBER Working Paper No 21448, 2015. https://doi.org/10.3386/w21448.