

Imperfect Risk-Sharing and the Business Cycle*

David Berger[†]

Luigi Bocola[‡]

Alessandro Dovis[§]

March 2019

Preliminary and Incomplete

Abstract

Is imperfect risk-sharing across households important for aggregate fluctuations? This paper develops a framework to answer this question. We first propose an accounting procedure for households-level data on consumption, wages and hours worked. In our prototype model with complete financial markets, households' choices regarding labor supply and assets' holdings are distorted by individual-specific wedges. We measure the wedges using micro data and show how to combine them with a class of heterogeneous agents New Keynesian models. The models in this class have an equivalent representation featuring a stand-in household with state-dependent preferences. We derive a mapping between these preference "shocks" and the micro wedges, and use this equivalent representative-agent economy to perform counterfactuals. We find that deviations from perfect risk-sharing was an important determinant of aggregate demand during the U.S. Great Recession.

Keywords: Business Cycles, Heterogeneous Agents, Risk-Sharing

JEL codes: F34, E44, G12, G15

*First draft: February 2019. We thank Anmol Bhandari, Loukas Karabarbounis, Greg Kaplan, Patrick Kehoe, Dirk Krueger, Guido Lorenzoni and seminar participants at SED 2019, Atlanta Fed, Chicago Fed, UCLA, Wharton, Stanford, Harvard, UC Davis, Saint Louis Fed, Minneapolis Fed, AEA 2019, Princeton, AMCM conference, CREI, LSE, Bocconi and the National Bank of Belgium.

[†]Northwestern University and NBER

[‡]Stanford University and NBER

[§]University of Pennsylvania and NBER

1 Introduction

A classic question in macroeconomics is to what extent household heterogeneity and deviations from perfect risk-sharing are important for aggregate fluctuations. For a long time, most business cycle theories relied on the representative-agent paradigm, an approach that was partly justified by the influential result in [Krusell and Smith \(1998\)](#) that distributional issues play a limited role for macroeconomic dynamics in standard real business cycle models. More recently, a new line of research is reevaluating this question. Studies have shown that in environments where risk-sharing is limited, households' self-insurance motives can be quite relevant for aggregate demand and can thus play an important role for business cycle if output is partly demand-determined, see [Krueger, Mitman, and Perri \(2016\)](#) and [Kaplan and Violante \(2018\)](#) for surveys of this literature. The answer to this question does depend, however, on features of the structural model that are hard to discipline empirically, such as the precise set of financial assets and risk-sharing mechanisms that are available to households, as well as the timing and distribution of fiscal transfers.¹

In this paper, we seek to provide an answer to this question that is robust to these considerations. Rather than specifying the details of which risk-sharing mechanisms are available and how they are modeled, we use households' consumption and labor choices to measure the degree of imperfect risk-sharing present in U.S. data, and we provide a framework to assess its macroeconomic implications. Our approach consists of two steps.

In the first step, we develop an accounting procedure for household-level data. Any path for household's consumption, labor supply, and wages can be rationalized as the outcome of a *prototype* model with complete financial markets and time-varying *wedges*. These wedges capture deviations between the households' choices observed in the data and those that would be observed under perfect risk-sharing, and they can be measured using household-level data and the equilibrium conditions of the prototype model. In the second step, we characterize the moments of the joint distribution of the wedges that are relevant for macroeconomic dynamics. Specifically, and building on the work of [Nakajima \(2005\)](#), [Krueger and Lustig \(2010\)](#) and [Werning \(2015\)](#), we show that canonical business cycle models with incomplete financial markets have an equivalent representative-agent formulation where the stand-in household has state-dependent preferences. These preference "shocks" characterize the aggregate implications of frictions at the household-level, and they are functions of the

¹A 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. In addition, researchers have shown that many other modeling choices, which are inconsequential in representative agent economy, matter in heterogeneous agent economies. These include the timing and distribution of the fiscal transfers ([Kaplan, Moll, and Violante, 2018](#)), how profits get distributed across households ([Broer, Hansen, Krusell, and Oberg, 2018](#)), and the cyclicalities of idiosyncratic uncertainty and access to liquidity ([Werning, 2015](#); [Acharya and Dogra, 2018](#)).

wedges defined in the first step.

Our main contribution is to combine this equivalent representative-agent economy with the measured wedges to evaluate the implications of imperfect risk-sharing for the business cycle. In our application, we find that imperfect risk-sharing contributed significantly to the decline in nominal interest rates and real economic activity observed during the U.S. Great Recession.

We begin by describing the prototype model. We consider the decision problem of a household that chooses consumption, labor supply, and holdings of Arrow securities. The household's wage can differ from the average wage in the economy, and we label *efficiency wedge* the ratio between these two. In addition, we introduce individual specific taxes on labor supply and on purchases of Arrow securities. The former introduce a wedge between the households' marginal rate of substitution between consumption and leisure and the wage, so we label them *labor wedges*. The latter allow for the marginal utility of consumption between any two individuals to vary over time, so we label them *risk-sharing wedges*.

The wedges can vary across individuals and over time, and they guarantee that the prototype model rationalizes the path of consumption, hours worked, and wages observed in the data. Moreover, any structural model uniquely maps into a particular process for these wedges. As an example, we show that the allocation of a [Huggett \(1993\)](#) model where households face idiosyncratic income risk and can trade only a non-contingent bond is replicated in the prototype model by an appropriate sequence of efficiency and risk-sharing wedges: the former capture cross-sectional variation in wage per hour while the latter implies that individual consumption responds to idiosyncratic labor productivity shocks.²

We measure the wedges using the equilibrium conditions of the prototype model and two households' surveys routinely employed in the literature, the Consumer Expenditure Survey (CEX) and the Panel Study on Income Dynamics (PSID). By and large, these two datasets give a consistent picture of the evolution of the wedges. For example, we document a positive cross-sectional relation between the labor and the efficiency wedge. In the time-series, the average labor wedge spikes in recessions while, with the exception of the Great Recession, we observe little time-variation in the cross-sectional moments of the risk-sharing wedge. Because the wedges are highly dimensional, however, it is important to identify the moments that are relevant for macroeconomic dynamics. The second step of our procedure accomplishes this task. Specifically, we show that few statistics of the distribution of these wedges are sufficient to measure the aggregate implications of imperfect risk-sharing in

²Under CRRA preferences, the prototype model with zero risk-sharing wedges predicts that individual consumption is a constant fraction of aggregate consumption. The taxes on Arrow securities are thus necessary to generate a systematic relation between income and consumption at the household level, which is a key prediction of standard incomplete market models.

canonical business cycle models.

We illustrate this point by considering a class of New Keynesian models with idiosyncratic income risk and incomplete markets. The models in this class share the same specification for households' preferences, technology, market structure, and the conduct of monetary policy. They can differ in the nature of idiosyncratic risk faced by households, the set of assets they can trade, and their financial constraints. Each of these potential specifications imposes a particular set of restrictions on the distribution of the wedges. However, the equilibrium law of motion for aggregate variables for all the model in this class can be equivalently represented as that of an economy with a representative household that features state-dependent preferences – that is, time-varying and stochastic rate of time preference and disutility of labor. Importantly, these preference shifters are simple functions of the micro wedges we defined.

In the equivalent representative-agent economy, the discount factor that the stand-in household applies to a particular aggregate state of the world is the product of the “true” rate of time-preference and the average tax on Arrow securities that households in the heterogeneous agents economy pay in that state of the world. This result can be explained as follows. A high average tax on Arrow securities means that households have hard time insuring their idiosyncratic risk in that state of the world. In the equivalent representative-agent economy, these frictions impeding risk-sharing are captured by a higher discount rate of the stand-in household. The disutility of labor is, instead, a convolution of all the micro wedges, and it captures compositional changes in the labor force. As an example, suppose that in the heterogeneous-agent economy households with high consumption are also characterized by high labor productivity. More cross-sectional dispersion in consumption, holding the average constant, induces high productivity households to work less and low productivity households to work more because of wealth effects, a compositional change that reduces worked hours in efficiency units. In the equivalent representative-agent economy, these effects are captured by an increase in the disutility of labor of the stand-in household – a reduction in labor supply.

We construct the time-path of these preference shocks using the micro wedges that we measured with the CEX. An important finding is that the discount factor, while mostly constant throughout the sample, displays a sizable increase during the Great Recession. It is well known in the literature that changes in the discount factor can induce sizable business cycle fluctuations in representative-agent New Keynesian models. As shown for example in [Christiano, Eichenbaum, and Rebelo \(2011\)](#), these models can generate large output drops when the zero lower bound on nominal interest rate binds, and a common device used in the literature to reach the zero lower bound is an increase in the discount factor. An important question is whether the movements in the discount factor that are

due to imperfect risk-sharing are large enough to be a quantitatively important source of business cycle fluctuations.

To address this question, we estimate the parameters of the representative-agent economy and of the process governing the preference shocks and construct the path for aggregate output, inflation and nominal interest rates that would have emerged in an economy with complete financial markets – that is, an economy with no risk-sharing wedges. The difference between the observed and the counterfactual path isolates the effects that imperfect risk-sharing has on these aggregate variables. Our main finding is that imperfect risk-sharing across households had sizable macroeconomic effects during the Great Recession, accounting for roughly half of the output decline observed in 2008 and 2009.

Related Literature. Our research contributes to a growing literature that introduces heterogeneous agents and incomplete financial markets in New Keynesian models of the business cycle. Researchers have used these environments to study how frictions impeding risk-sharing across households affect the transmission mechanism of monetary and fiscal policy (Kaplan, Moll, and Violante, 2018; Auclert, 2017; McKay, Nakamura, and Steinsson, 2016; McKay and Reis, 2016), and more generally the business cycle. For example, and in the context of the Great Recession, Guerrieri and Lorenzoni (2017) study the aggregate implications of the tightening in individual borrowing constraints, while Heathcote and Perri (2018), Bayer, Lüticke, Pham-Dao, and Tjaden (2019) and Ravn and Sterk (2017) show that a rise in idiosyncratic income risk and the associated precautionary saving motives can induce a decline in aggregate demand and economic activity.

All these papers consider specific departures from perfect risk-sharing by imposing a given asset structure, income process and set of financial restrictions. We instead take a more agnostic approach about the amount of risk-sharing available to households and infer it from their observed choices: our approach is robust to misspecification of these model elements and, by construction, it is consistent with a broad set of facts about households' behavior. We think that these two approaches are complementary. We identify a set of cross-sectional moments that are informative about the macroeconomic effects of imperfect risk-sharing in this class of models and, consistent with some of the above-mentioned paper, we document that their behavior points toward an important role for micro-level frictions in explaining the decline in aggregate demand during the Great Recession.³ However, our approach is silent about the set of underlying frictions and shocks that can replicate the observed patterns of

³The literature emphasizes the role of the distribution of marginal propensities to consume (MPCs) as a critical statistic to discipline these structural models. Auclert, Rognlie, and Straub (2018) show that the distribution of MPCs at different time horizons – what they term intertemporal MPCs – is a sufficient statistic for the response of output to fiscal shocks in a class of New Keynesian models. In Section 4 we discuss the relation between the cross-sectional moments we identify and the intertemporal MPCs.

the wedges. Identifying these frictions is important because our approach cannot be used for policy evaluation, and so a fully specified structural model is needed.⁴

Our approach is similar to the one in a large literature that uses data on household consumption, labor supply, and earnings to measure the degree of risk-sharing in the data without explicitly specifying the mechanisms through which households share risk. See for example [Blundell, Pistaferri, and Preston \(2008\)](#) and the survey in [Jappelli and Pistaferri \(2010\)](#). The paper that is closer to our approach is [Heathcote, Storesletten, and Violante \(2014\)](#) who use households' optimality conditions and PSID and CEX data to measure the extent of risk-sharing present in the U.S. economy.⁵ The contribution of our paper relative to this literature is to study how the measured degree of partial risk-sharing affects aggregate dynamics.

At the methodological level, our approach builds on several papers that infer distortions from residuals in first-order conditions. [Chari, Kehoe, and McGrattan \(2007\)](#) measure wedges by combining the canonical real business cycle model with aggregate U.S. data, while [Hsieh and Klenow \(2009\)](#) measure the degree of misallocation of factor of productions through the lens of a frictionless heterogeneous firms model. Our contribution is to apply the logic of these accounting procedures for household-level data, and to derive simple summary statistics of the micro data that are sufficient to characterize the impact of imperfect risk-sharing for aggregate fluctuations.⁶ As in those studies, the measurement depends on the prototype model used: misspecification of households' preferences and technologies, for example, are captured in our framework as wedges. The idea that wedges can capture deviation from perfect risk-sharing has been used also in [Fitzgerald \(2012\)](#) and [Ohanian, Restrepo-Echavarria, and Wright \(2018\)](#) to study international capital flows. They do not focus on the role of imperfect risk-sharing in the propagation of shocks.

Finally, our paper is related to the literature that evaluates asset pricing models where aggregation does not hold using households' consumption data. See for example [Brav, Constantinides, and Geczy \(2002\)](#), [Vissing-Jørgensen \(2002\)](#), [Krueger, Lustig, and Perri \(2008\)](#), and [Kocherlakota and Pistaferri \(2009\)](#). The goal of 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 in the equivalent representative-agent economy in our approach. Clearly the scope of our analysis differs from these papers.

⁴See, for example, [Moll, Townsend, and Zhorin \(2017\)](#) and [Karaivanov and Townsend \(2014\)](#) for this type of analysis in different contexts.

⁵We discuss in more details the relationship between our accounting procedure and the framework in [Heathcote, Storesletten, and Violante \(2014\)](#) in Section 2.

⁶[Boerma and Karabarbounis \(2019\)](#) combines cross-sectional data with a model of home production, idiosyncratic income risk and incomplete markets, but use it to infer the degree of inequality across households rather than the macroeconomic implications of imperfect risk-sharing.

Layout. The paper is organized as follows. Section 2 presents the prototype model and defines the micro wedges while in Section 3 we use the CEX and PSID to measure the wedges and uncover key cross-sectional patterns. Section 4 introduces the class of heterogeneous agents economies at the center of our application, derives the equivalent representative-agent representation and explains how to combine it with the micro wedges to perform counterfactuals. Section 5 and 6 implements these counterfactuals for the US economy. Section 7 concludes.

2 Prototype model

This section introduces the benchmark *prototype* model and shows how individual specific wedges can be used to rationalize households' level data on consumption, hours worked and wages. We also discuss the restrictions on these wedges implied by various models with underlying frictions in labor and financial markets.

2.1 Prototype model

Time is discrete and indexed by $t = 0, 1, \dots$. The economy is populated by a continuum of households divided into a finite number of types $i \in I$. Let μ_i be the measure of type i agent in the economy. 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 $z^t = (z_0, z_1, \dots, z_t)$ be a history of realized aggregate states up to period t and $v^t = (v_0, v_1, \dots, v_t)$ be a history of idiosyncratic states up to period t . We also let $s_t = (z_t, v_t)$ and $s^t = (z^t, v^t)$. Let $\Pr_i(s^t | s^{t-1})$ be the probability of a history s^t . We assume that $\Pr_i(s^t | s^{t-1}) = \Pr_i(v^t | v^{t-1}, z^t) \Pr(z^t | z^{t-1})$ so we allow for the possibility that the aggregate state affects the cross-sectional distribution of the idiosyncratic state and that the agent's type affects the probability of drawing a given v_t .

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

$$\sum_{t=0}^{\infty} \sum_{s^t} \beta^t \Pr_i(s^t | s_0) U(c_i(s^t), l_i(s^t)), \quad (1)$$

We further assume that the period utility is given by

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

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

Each household is endowed with some initial wealth a_0 and with a unit of labor every period.⁷ Households choose how much to work, consume and save, taking as given the wage and the price of financial assets. We assume that the real wage after history s^t is given by $W(s^t) = \theta(v_t) W(z^t)$ where $W(z^t)$ is the average wage after history z^t . We refer to $\theta(v_t)$ as the *efficiency wedge* because it introduces a wedge between individual and average wages, and we assume that in all possible aggregate histories the cross-sectional average of $\theta(v_t)$ is constant and equal to 1, $\sum_i \mu_i \sum_{v_t} \Pr_i(v_t|z^t) \theta(v_t) = 1$ for all z^t .

Furthermore, we assume that asset markets are complete, and households can trade Arrow securities contingent on the aggregate and idiosyncratic state. We let $Q_i(s^t, s_{t+1})$ be the price of an Arrow security for an individual of type i with history s^t that pays one unit of the final consumption good next period if s_{t+1} is realized. This insurance is provided by competitive financial intermediaries. Absence of arbitrage opportunities requires that

$$Q_i(z^t, v^t, z_{t+1}, v_{t+1}) = \Pr_i(v_{t+1}|z^{t+1}, v^t) Q(z^t, z_{t+1}). \quad (3)$$

That is, the price of an individual Arrow security that pays in state (z^{t+1}, v^{t+1}) is the aggregate state price, $Q(z^{t+1})$, times the probability that the individual has idiosyncratic state v_{t+1} given the current history and the aggregate state is z_{t+1} .

Labor income is taxed with a proportional tax $\tau_{l,i}(s^t)$. The tax can depend on both the aggregate and idiosyncratic state of the economy and the type of the agent. We refer to $\tau_{l,i}(s^t)$ as the *labor wedge* because it introduces a wedge between the household's marginal rate of substitution between consumption and leisure and the household's wage.⁸ We also introduce a proportional tax on the purchase of Arrow securities, $\tau_{a,i}(s^{t+1})$. We will refer to $\tau_{a,i}(s^{t+1})$ as the *risk sharing wedge*. The presence of this wedge allows for the ratio of marginal utility of consumption for two individuals to vary over time.

The problem for an individual consumer is then to choose consumption, labor supply, and holdings of Arrow securities, $a_i(s^t, s_{t+1})$, to maximize his lifetime utility (1) subject to the budget constraint

$$c_i(s^t) + \sum_{s_{t+1}} \left(1 + \tau_{a,i}(s^{t+1})\right) Q_i(s^t, s_{t+1}) a_i(s^t, s_{t+1}) \leq (1 - \tau_{l,i}(s^t)) W(s^t) l_i(s^t) + a_i(s^t) + T_i(s^t),$$

⁷To ease notation, the initial level of wealth is included in v_0 and we do not index allocation by a_0 .

⁸Importantly, the labor wedge in our paper differs from the one in [Chari, Kehoe, and McGrattan \(2007\)](#). In their paper, the labor wedge is defined as the wedge between the marginal rate of substitution between consumption and leisure for the household and the marginal product of labor. We replace the marginal product of labor with the wage. See [Karabarbounis \(2014\)](#) and [Bils, Klenow, and Malin \(2018\)](#) for an empirical analysis of the wedge between households' marginal rate of substitution and the wage.

given lump-sum transfers $T(s^t)$. The solution is an allocation $\{c_i(s^t), l_i(s^t), a_i(s^{t+1})\}$ that satisfies the optimality conditions

$$\frac{\beta \Pr_i(s^{t+1}|s^t) U_{c,i}(s^{t+1})}{U_{c,i}(s^t)} = \left(1 + \tau_{a,i}(s^{t+1})\right) Q_i(s^t, s_{t+1}) \quad (4)$$

$$\frac{-U_{l,i}(s^t)}{U_{c,i}(s^t)} = (1 - \tau_{l,i}(s^t)) W(s^t), \quad (5)$$

and the budget constraint.

We can now define a (partial) equilibrium outcome in this economy. Given taxes and transfers $\{\tau_{a,i}(s^t), \tau_{l,i}(s^t), T_i(s^t)\}$, efficiency wedges $\{\theta(v_t)\}$ and aggregate prices $\{W(z^t), Q(z^t)\}$, an equilibrium is a set of allocations $\{c_i(s^t), l_i(s^t), a_i(s^{t+1})\}$ such that consumption, labor and asset holdings are optimal for the household and the price of Arrow securities $Q_i(s^t, s_{t+1})$ satisfy the no-arbitrage condition (3).

2.2 Recovering the wedges

We now show how data on households consumption, hours worked and wages can be used to measure the realization of the individual wedges of the prototype model. This is possible because, given any sequence $\{c_i(s^t), l_i(s^t), W_i(s^t)\}$, there exists a unique sequence of wedges that rationalizes it as an equilibrium outcome of the prototype economy.

Specifically, suppose we observe a panel for individual consumption, hours' worked and wages. The realization of the efficiency wedge can be constructed by taking the ratio between individual and average wages,

$$\theta(v_t) = \frac{W(z^t, v^t)}{W(z^t)}. \quad (6)$$

Similarly, we can specialize the labor supply condition (5) to our preferences, and obtain an expression for the labor wedge as a function of observables and model parameters

$$\tau_{l,i}(s^t) = 1 - \chi \frac{l_i(s^t)^\psi c_i(s^t)^\sigma}{W(s^t)}. \quad (7)$$

There is a layer of indeterminacy when it comes to the risk sharing wedge. From the Euler equation (4), we cannot separately identify the risk sharing wedge and $Q_i(z^t, z_{t+1})$. In what follows we set

$$Q(z^t, z_{t+1}) = \beta \Pr(z^{t+1}|z^t) \frac{U_c(z^{t+1})}{U_c(z^t)},$$

where $U_c(z^t)$ is shorthand notation for $U_c(C(z^t))$ and $C(z^t)$ is aggregate consumption, $C(z^t) = \sum_i \mu_i \sum_{v^t} \Pr_i(v^t|z^t) c_i(z^t, v^t)$. As we discuss later in the paper, this normalization

does not affect the results in our main application.

We can then combine this choice for $Q(z^t, z_{t+1})$ with equations (3) and (4) to obtain an expression for the risk sharing wedge as a function of observable variables,

$$\tau_{a,i}(s^{t+1}) = \left(\frac{C(z^{t+1})/C(z^t)}{c_i(s^{t+1})/c_i(s^t)} \right)^\sigma - 1. \quad (8)$$

The following proposition summarizes this result

Proposition 1. *Any sequence $\{c_i(s^t), l_i(s^t), W_i(s^t)\}$ can be supported as an equilibrium of the prototype model with wedges $\{\tau_{a,i}(s^t), \tau_{l,i}(s^t), \theta(v_t)\}$ given by equation (6), (7) and (8), and transfers given by*

$$\sum_{t=0}^{\infty} \sum_{s^t} \beta^t Pr_i(s^t) \{U_{c,i}(s^t) [c_i(s^t) - T_i(s^t)] + U_{l,i}(s^t) l_i(s^t)\} = U_{c,i}(s_0) a_0.$$

To understand the role of the wedges, it is instructive to study the allocation that would prevail if taxes were equal to zero for all histories s^t .

Consider the risk sharing wedge first. In an economy with perfect risk sharing, the ratio of the marginal utility between any two individuals is constant for all histories: for all v^{t+1} and \tilde{v}^{t+1} it must be that $U_{c,i}(z^{t+1}, v^{t+1}) / U_{c,i}(z^t, v^t) = U_{c,i}(z^{t+1}, \tilde{v}^{t+1}) / U_{c,i}(z^t, \tilde{v}^t)$. Given the preferences in equation (2), this would imply that individual consumption is a constant fraction of aggregate consumption across individual and aggregate histories,

$$c_i(z^t, v^t) = \varphi_i(v_0) C(z^t).$$

In other words, when $\tau_{a,i}(s^{t+1}) = 0$ for all s^{t+1} , individual consumption responds only to aggregate shocks and not to idiosyncratic ones. Clearly, such prediction is not borne in the data. The risk sharing wedge is what allows the prototype model to generate time-varying consumption shares. In particular, letting $\varphi_i(s^t) = c_i(s^t) / C(z^t)$ be the consumption share of an individual of type i with history s^t , we have from equation (8) that

$$\varphi_i(s^{t+1}) = \varphi_i(s^t) \left(1 + \tau_{a,i}(s^{t+1})\right)^{-1/\sigma} = \prod_{s^j \preceq s^{t+1}} \left(1 + \tau_{a,i}(s^j)\right)^{-1/\sigma} \varphi_i(v_0). \quad (9)$$

Consider the individual labor wedge next. If labor taxes were set to zero, our prototype

economy would imply the following behavior for individual hours worked

$$l_i(s^t) = \frac{[\theta(v_t) / \varphi_i(s^t)]^{1/\psi}}{\sum_i \mu_i \sum_{v^t} \Pr_i(s^t|z^t) [\theta(v_t) / \varphi_i(s^t)]^{1/\psi}} L(z^t),$$

where $L(z^t)$ are aggregate hours worked. That is, individuals with high efficiency/low consumption would work on average more. Any deviation from this optimal labor supply is absorbed by a labor wedge. By appropriately taxing labor supply, our procedure can rationalize any pattern for hours worked in the data.

An implication of Proposition 1 is that any *detailed economy*, that is any structural model, has unique implications for the wedges of the prototype model. This mirrors the equivalence results in [Chari, Kehoe, and McGrattan \(2007\)](#), and it implies that the wedges could be used to discriminate structural models: that is, one could verify whether the predictions of a given model are consistent with the data by comparing the wedges implied by the model with those obtained from micro data.

In Appendix A, we present several examples of detailed economies and show how they map into the wedges of the prototype model. First, we show that standard incomplete market models, along the lines of [Huggett \(1993\)](#), generate time-varying risk sharing and efficiency wedges at the micro level, and they also imply a positive cross-sectional correlation between these two, as positive (negative) idiosyncratic income shocks imply an increase (decrease) in individuals' consumption share. Second, we show that a model with wage rigidities as in [Erceg, Henderson, and Levin \(2000\)](#) generates a distribution of labor wedges across the population, where the heterogeneity arises because only a fraction of wages is reset every period. Finally, wedges might not only reflect "frictions" in labor and financial markets, but also misspecification of the prototype model. We make this point by showing that a simple endowment economy in which agents have heterogeneous risk aversion can generate variation in consumption shares even when financial markets are frictionless: through the lens of the prototype model, this variation in consumption shares would be captured by non-zero risk sharing wedges.

2.3 Discussion

(To be added: preference heterogeneity; relation to [Heathcote, Storesletten, and Violante \(2014\)](#))

3 Measuring the wedges

Having introduced the prototype model and defined the wedges, we now measure them using two households' level surveys routinely employed in the literature, the Consumer Expenditure Survey (CEX) and the Panel Study of Income Dynamics (PSID). Section 3.1 describes the data while Section 3.2 discusses the behavior of the wedges.

3.1 Data description

The CEX and the PSID collect information on income, expenditures, employment outcomes, wealth and demographic characteristics for a panel of U.S. households. We describe each data set below.

The Consumer Expenditure Survey (CEX) is a rotating panel of U.S. households, selected to be representative of the population, collected at a quarterly frequency. Households report consumption expenditures for a maximum of four consecutive quarters, income and employment information is collected in the first and last interview, and wealth information in the last interview only. The Panel Study of Income Dynamics (PSID) is a panel of U.S. households, selected to be representative of the U.S. population, collected at a bi-annual frequency. While the PSID has always collected information on income, hours and wealth, in 1999 the PSID started to collect comprehensive information about consumption expenditures as well. We use CEX data over the period 1984-2012 and PSID data from 1999-2015.

Our baseline sample includes all households where the head of the household is between the ages of 22 and 60. We only use households who participate in all four interviews in the CEX, as our main income measure and most savings questions are only asked in the final interview. For the PSID, we require that a household appears for at least three interviews. Further details about sample selection in each dataset are discussed in Appendix B.

In order to construct the wedges, we need information on consumption expenditures, total hours worked and wage per hour. Our measure of consumption is dollar spending in non-durables and services by the household. Total hours worked include hours worked by the head of household and the spouse over the entire year in all jobs. We obtain wage per hour by scaling an indicator of labor income at the household's level by total hours worked. Labor income is a pre-tax measure, and it includes wages and salaries, bonuses, overtime, tips plus income from a business. We also compute total household disposable income, which is equal to total household income including transfers net of taxes. In addition to the variables that are necessary for the calculation of the wedges, we obtain socio-demographics indicators about the households (education, sex, family size, etc.), information on the household's

assets and liabilities and disposable income.⁹ Specifically, we will use in our analysis an indicator of liquid assets (savings and checking accounts) and an indicator of net-worth at the households' level. We use the CPI-U to express all monetary variables in constant 2000 dollars. To eliminate outliers and mitigate any impact of time-varying top-coding, we drop observations in the top and bottom one percent of the consumption, hours, labor income, wage per hour, disposable income, liquid assets and net-worth distribution.

The prototype model abstracts from important features of the micro data, such as life-cycle dynamics. In order to have a clear mapping between model and data, we need to purify our variables of these factors. We do so through panel regressions. Let \tilde{c}_{it} the the log of consumption expenditures at the household level and \tilde{y}_{it} the log of labor income. We estimate, in both the CEX and PSID panel, the following linear equation

$$\tilde{c}_{it} = \alpha + \gamma' X_{it} + \gamma_y \tilde{y}_{it} + e_{it},$$

where X_{it} includes dummies for the sex, education and age of the head of household, state of residence, and the quarter in which the interview is conducted.¹⁰ After estimating this regression, we predict consumption only using labor income and the residual,

$$\tilde{c}_{it}^p = \alpha + \gamma_y \tilde{y}_{it} + e_{it}.$$

We repeat this procedure for all variables used in the analysis.¹¹ After estimating these relationships, we divide all variables in levels, with the exception of wage per hour, by the number of family members in order to obtain per-capita figures.

Appendix B presents summary statistics of the underlying micro data and some comparisons with previous studies in the literature. Households' characteristics in our sample are comparable, both for the PSID and for the CEX, with the ones reported in [Heathcote and Perri \(2018\)](#). In line with their findings, we also verify that the behavior of aggregate consumption expenditures, labor income and hours worked implied by the two datasets tracks reasonably well their corresponding values in National Accounts. We finally compute a set of cross-sectional statistics and show that their behavior over time closely mirror the one reported in previous papers in a large body of work on consumption and income inequality ([Blundell, Pistaferri, and Preston, 2008](#); [Krueger and Perri, 2006](#); [Aguiar and Bils, 2015](#); [Attanasio and Pistaferri, 2016](#)).

⁹Disposable income differ from our measure of labor income because it is after tax, and it includes government transfers and income from other sources.

¹⁰We only include quarter dummies in the consumption regression because all other variables are measured annually.

¹¹We do not include \tilde{y}_{it} when predicting labor income.

3.2 Wedges in the data

We can now use the observations $\{c_{it}, l_{it}, w_{it}\}$ and the expressions in (6), (7) and (8) to construct the wedges. For that purpose, we need to choose the parameters governing households' preferences. In our baseline calculations, we set the degree of constant relative risk aversion and the Frish elasticity of labor supply to 1. We also choose χ so that the average labor wedge in the first year of the sample equals 30%.

Households in the PSID are interviewed once every two years. The risk-sharing wedge, then, is computed using the changes in consumption between two consecutive interviews (8 quarters changes). The CEX has a quarterly frequency, and each household is interviewed four times. Thus, we have at most four observations for consumption and two observations for hours worked and income for each household. To minimize time aggregation issues, for each year we focus on households that have their last interview in the third and fourth quarter, and compute the risk-sharing wedge by using the change in consumption from the first quarter of that year. To make the risk-sharing wedge comparable across datasets, we report consumption changes at quarterly frequency.

Figure 1 reports the histograms of the three wedges pooled across all years in the overlapping sample. This is meant to give a snapshot of the cross-sectional variation that we uncover in our data sets. The top three panels of the figure report the wedges computed using the CEX, while the bottom three panels report comparable figures computed using the PSID.

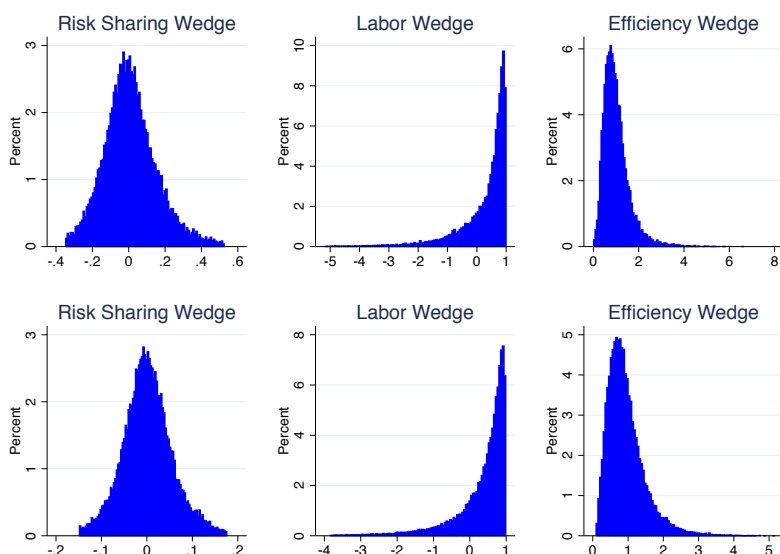
We can verify from the figure two main facts. First, there is a substantial degree of dispersion in all three individual wedges in both data sets. Second, comparing across the two rows, we can see that, with the exception of the risk-sharing wedge, the results are remarkably similar across the two data sets – qualitatively *and* quantitatively.¹²

Some benchmarks are useful for interpreting the cross-sectional dispersion in the wedges. Given log-preferences, we know that an individual risk sharing wedge equals zero if the consumption growth of that individual between time $t - 1$ and t was equal to aggregate consumption growth over the same period. Similarly, a risk sharing wedge equals to 0.3 implies a fall in consumption expenditure between time $t - 1$ and t of 24%, assuming no growth in the aggregate. Over 50% of households have risk sharing wedges between $[-0.3, 0.3]$ in our sample. Turning to the labor wedge, an unemployed person has a labor wedge equal to 1, a 100% tax on labor supply.¹³ Negative labor wedges are subsidies to labor. These

¹²The difference in the risk-sharing wedges across the two datasets is likely due to the fact that consumption changes are computed at different frequencies.

¹³There is an indeterminacy in our wedges when it comes to unemployment, as we do not observe the wage for an unemployed individual. So, we could in principle rationalize unemployment with $\theta_{it} = 0$ or with $\tau_{l,it} = 1$. We choose to follow the latter normalization.

Figure 1: The distribution of individual wedges



Top Panel is computed using data from the CEX; bottom panel is computed using data from the PSID.

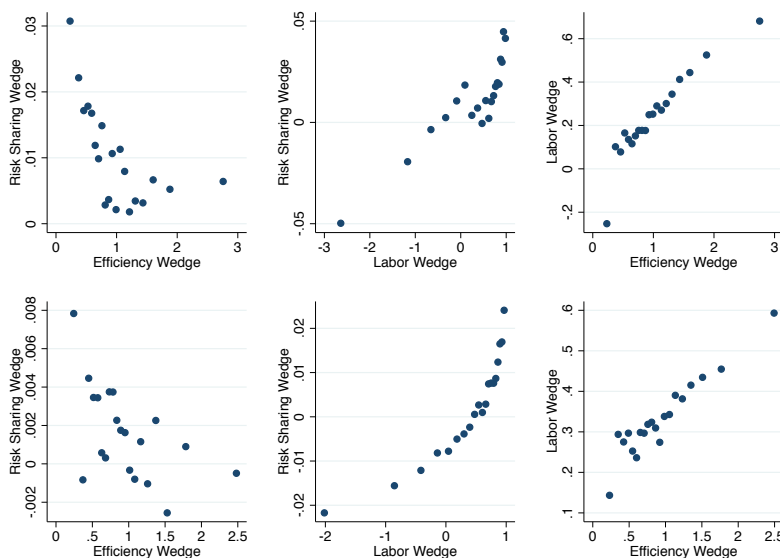
occurs, for example, when households that have high consumption/low efficiency and also have high hours worked. The efficiency wedge is the easiest to interpret: an efficiency wedge equal to 2 implies a wage per hour for that household that is twice as large as the average wage. Examining figure 1, we can see that most households earn less than the mean wage in the data.

Figure 2 shows bin scatter plots that display the pairwise cross-correlations between the three wedges. There are clear cross-sectional relationships between each pair-wise combination of wedges, which are consistent across the two data sets. In the left panel we plot the relation between the efficiency wedge and the risk sharing wedge. In a world with perfect risk sharing, this relation would be an horizontal line. In the data, instead, we see that households with low idiosyncratic efficiency wedge at time t typically high risk-sharing wedges – that is, below average consumption growth between time $t - 1$ and time t . We can also see that households with a higher labor wedges have also high efficiency wedges. The positive relation between the efficiency and the labor wedge means that the prototype model needs labor subsidies for households with “low efficiency”, and taxes on households with “high efficiency” to rationalize the data. This reflects the fact that wage per-hour is more dispersed across households than hours worked.

Figure 3 reports the cross-sectional average of the labor wedge and of the risk-sharing wedge across the two datasets.¹⁴ The labor wedge is highly cyclical in both datasets: it is

¹⁴We did not include the efficiency wedge because, by construction, the sample average equals one in every period.

Figure 2: Correlation across wedges



Top Panel is computed using data from the CEX; bottom panel is computed using data from the PSID.

high during the 2000-2001 recession and in the Great Recession of 2007-2009 in both dataset, in line with the results in [Karabarbounis \(2014\)](#) for the “household component” of the aggregate labor wedge. The average risk-sharing wedge increases in the Great Recession in both datasets.

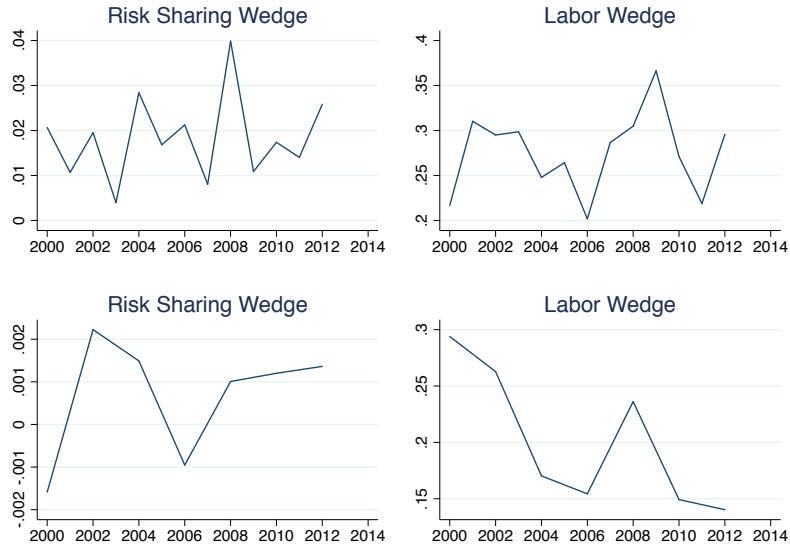
4 New Keynesian models with heterogeneous agents

So far we have defined the micro wedges and showed how to measure them using survey data. We now show how the wedges can be combined with equilibrium models to evaluate the implications of imperfect risk sharing for the business cycle.

We start in Section 4.1 by detailing a *class* of Heterogeneous Agents New Keynesian (HANK) models. The models in this class share the same specification for households’ preferences, technology, market structure and the conduct of monetary policy. However, they can differ in the nature and cyclical nature of idiosyncratic risk faced by households, the set of assets they can trade and their financial constraints, the timing and distribution of the fiscal transfers, and how profits are distributed across households. In other words, our formulation leaves fairly unrestricted the “micro side” of the models in this class, encompassing in this fashion several variants studied in the literature.

In Section 4.2 we discuss the determination of macroeconomic aggregates in this class of models. Specifically, we show that their law of motion is equivalent to that of a representative

Figure 3: Cross-sectional average of risk-sharing and labor wedge



Top Panel is computed using data from the CEX; bottom panel is computed using data from the PSID.

agent economy with “taste shocks”— that is, with a stochastic and time-varying discount factor and disutility of labor for the representative household. We emphasize two important features of this representation: *i*) the taste shocks are simple functions of the micro wedges defined in Section 2 and *ii*) this mapping is invariant to the specific set of assets of financial constraints that households are facing. Section 4.3 illustrates this result in two simple HANK economies.

Finally, Section 4.4 explains, at a conceptual level, how to combine this representation and the micro wedges to perform counterfactuals designed to isolate the role of imperfect risk sharing for the business cycle.

4.1 Model

Environment. As before, we denote the realization of the aggregate state by z_t and by v_t the realization of the idiosyncratic state. The economy is populated by households, final good firm producers, intermediate good firms and the monetary authority. The households derive utility from consuming the final good and from leisure according to the utility function defined in (2).

The final good producers are perfectly competitive and they aggregate intermediate inputs to produce a final good according to the technology

$$Y(z^t) = \left(\int_0^1 y_j(z^t)^{\frac{1}{\mu}} dj \right)^{\mu}, \quad (10)$$

where μ is related to the (constant) elasticity of substitution across varieties, ε , by the following, $\mu = \varepsilon / (\varepsilon - 1)$.

Firms producing intermediate goods operate under monopolistic competition, and they purchase labor services from households in competitive labor markets in order to produce their variety,

$$y_j(z^t) = A(z_t)n_j(z^t), \quad (11)$$

where $A(z_t)$ is an aggregate technology shock, common across firms, and $n_j(z^t)$ is labor in *efficiency units* demanded by firm j . In equilibrium, the labor market clears and we have that

$$\int n_j(z^t) dj = \sum_i \mu_i \sum_{v^t} \Pr_i(v^t | z^t) e(v_t) l_i(v^t, z^t).$$

That is, 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 income risk for households. We denote by $W(z^t)$ the wage per efficiency unit of labor.

Households. Households enter the period with some financial assets and they work for intermediate good producers. We assume that their labor income is taxed at $\tau_{l,i}(s^t)$ and that households receive a lump-sum transfer $T_i(s^t)$. This is a simple way to introduce a deviation between the marginal rate of substitution between consumption and leisure and the wage. The households choose consumption, new financial positions and labor in order to maximize their expected life-time utility.

We model financial markets in a flexible way. First, we assume that households can trade a risk-free nominal bond. We denote by $b_i(s^t)$ the position taken today by a household of type i and by $1 + i(z^t)$ the nominal return on the bond. We also assume that the household can trade a set \mathcal{K} of possible assets, with the nominal return of a generic asset $k \in \mathcal{K}$ for a given history (s^t, s_{t+1}) given by $R_k(s^t, s_{t+1})$. This formulation allows for different types of financial assets: individual Arrow securities, shares of the intermediate good firms, complex financial derivatives, etc. We let $a_{k,i}(z^{t-1}, v^{t-1})$ be the holdings of assets k that a household of type i with history v^{t-1} has accumulated after an aggregate history z^{t-1} . Trades in these additional financial assets potentially require transaction costs $\mathcal{T}(\{a_{k,i}(s^{t-1})\}_{k \in \mathcal{K}}, \{a_{k,i}(s^t)\}_{k \in \mathcal{K}}, s^t)$.

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

$$H_i \left(b_i(s^t), \{a_{k,i}(s^t)\}_{k \in \mathcal{K}}, s^t \right) \geq 0 \quad (12)$$

for some function H_i . We refer to the set of constraints in (12) as *trading restrictions*.

The set of assets \mathcal{K} , the transaction costs in \mathcal{T} , and the trading restrictions in (12) are a flexible way of representing a large class of models with incomplete financial markets. Our formulation can encompass as special cases a 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 forms of restriction on asset trading in Chien, Cole, and Lustig (2011, 2012). Note, also, that the H_i function can depend on s^t , which implies that we are allowing for aggregate and idiosyncratic shocks to affect the financial constraints of households. The only restriction that we impose is that purchasing risk-free nominal bonds weakly relaxes these constraints, $H_b(b, \{a_k\}_{k \in \mathcal{K}}, s^t) \geq 0$ and it does not require a transaction cost. By doing so we are ruling out limited participation economies where agents must pay a fixed cost to have access to the risk-free nominal bond.

Given this structure, we can write the household's problem as

$$\max_{c_i, l_i, b_i, \{a_{k,i}\}_{k \in \mathcal{K}}} \sum_t \sum_{s^t} \beta^t \Pr_i(s^t | s_0) \left[\frac{c_i(s^t)^{1-\sigma} - 1}{1-\sigma} - \chi \frac{l_i(s^t)^{1+\psi}}{1+\psi} \right]$$

subject to the nominal budget constraint,

$$\begin{aligned} P(z^t) c_i(s^t) + \frac{b_i(s^t)}{1+i(s^t)} + \sum_{k \in \mathcal{K}} a_{k,i}(s^t) + \mathcal{T}(\{a_{k,i}(s^{t-1})\}_{k \in \mathcal{K}}, \{a_{k,i}(s^t)\}_{k \in \mathcal{K}}, s^t) &\leq \\ &\leq (1 - \tau_{i,l}(s^t)) W(z^t) e(v_i) l_i(v^t, z^t) + T_i(s^t) + b_i(s^{t-1}) + \sum_{k \in \mathcal{K}} R_k(s^{t-1}, s_t) a_{k,i}(s^{t-1}), \end{aligned}$$

and the trading restrictions in (12) given initial asset holdings.

Because of the assumption that $H_b(\cdot) \geq 0$, a necessary condition for optimality is

$$\frac{1}{1+i(z^t)} \geq \sum_{s_{t+1}} \left\{ \beta \frac{\Pr_i(s^{t+1} | s^t)}{\Pi(z^{t+1})} \left[\frac{c_i(s^t, s_{t+1})}{c_i(s^t)} \right]^{-\sigma} \right\}, \quad (13)$$

where $\Pi(z^{t+1}) = P(z^{t+1})/P(z^t)$ is the gross inflation rate. The condition must hold with equality if the trading restrictions (12) do not bind. For the rest of the paper, we assume that there always exist an agent for which the trading restrictions are not binding.¹⁵ Hence (13) holds with equality for the agent with the highest valuation for the risk-free bond.¹⁶

¹⁵This assumption implicitly imposes restrictions on the set of additional assets available, the trading restrictions, and the shocks. As an example, the assumption is automatically satisfied in a Huggett economy where the only asset available is the nominal risk-free bond in zero net supply and the trading restrictions are a debt limit of the form $b \geq -\phi$ with $\phi > 0$.

¹⁶To see this, simply note that the agent that attains the maximum in the right side of (13) is the one with the lowest multipliers on the trading restriction constraints.

Moreover, labor supply must satisfy

$$c_i(s^t)^{-\sigma}(1 - \tau_{l,i}(s^t))w(z^t)e(v_t) = \chi l_i(s^t)^\psi \quad (14)$$

where $w(z^t) = W(z^t)/P(z^t)$ is the real wage per efficiency unit.

Final good producers. The final good is produced by competitive firms that operate the production function in (10). From their decision problem, we can derive the demand function for the j -th variety

$$y_j(z^t) = \left(\frac{P_j(z^t)}{P(z^t)} \right)^{\mu/(1-\mu)} Y(z^t) \quad (15)$$

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

Intermediate good producers. The monopolist of variety j operates the technology (11). As in Rotemberg (1982), we assume that the firm faces quadratic costs to adjust its price,

$$\frac{\kappa}{2} \left[\frac{P_j(z^t)}{P_j(z^{t-1})} - 1 \right]^2. \quad (16)$$

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. We assume that the firm discounts future profits using the real state price

$$Q(z^{t+1}) = \max_{i,v^t} \left\{ \beta \Pr(z^{t+1}|z^t) \sum_{v_{t+1}} \Pr_i(v^{t+1}|z^{t+1}, v^t) \left[\frac{c_i(z^{t+1}, v^{t+1})}{c_i(z^t, v^t)} \right]^{-\sigma} \right\}. \quad (17)$$

That is, firms discount future profits using the marginal rate of substitution of the agent that values dividends in the aggregate state the most.¹⁷ The firm's problem can be written recursively as

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

subject to the production function (11) and the demand function (15).

The solution to the firm's problem together with symmetry across firms requires that the

¹⁷If all agents could trade Arrow securities contingent on the aggregate state then this would be the equilibrium state price.

following version of the New Keynesian Phillips curve holds in equilibrium

$$\tilde{\Pi}(z^t) = \frac{1}{\kappa(\mu - 1)} Y(z^t) \left[\mu \frac{w(z^t)}{A(z^t)} - 1 \right] + \sum_{z^{t+1}} Q(z^{t+1}|z^t) \tilde{\Pi}(z^{t+1}) \quad (18)$$

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

Monetary policy and market clearing. We assume that the monetary authority follows a standard Taylor rule

$$1 + i(z^t) = \max \left\{ (1 + \bar{i})^{1-\rho_i} [1 + i(z^{t-1})]^{\rho_i} \left(\frac{\Pi(z^t)}{\Pi^*} \right)^{\gamma_\pi} \left(\frac{Y(z^t)}{Y(z^{t-1})} \right)^{\gamma_y} \exp\{\varepsilon_m(z^t)\}, 1 \right\}, \quad (19)$$

where Π^* is the inflation target of the monetary authority, $(1 + \bar{i}) = \Pi^*/\beta$ is the nominal interest in a deterministic steady state of the model and $\varepsilon_m(z^t)$ is a monetary shock. Note that we allow for the possibility of a binding zero lower bound constraint.

In equilibrium, the labor market, goods markets, and financial markets clear. Specifically, equilibrium in financial markets requires that

$$\sum_i \mu_i \sum_{v^t} \Pr_i(v^t|z^t) b_i(z^t, v^t) = 0,$$

because bonds are in zero net supply, the value of inherited assets must equal the nominal value of the firm cum-dividend,

$$\sum_i \mu_i \sum_{v^t} \Pr_i(v^t|z^t) \sum_{k \in \mathcal{K}} R_k(s^{t-1}, s_t) a_{k,i}(s^{t-1}) = P(z^t) V(P(z^{t-1}), z^t),$$

and the total value of new asset positions must equal to the nominal value of the firm ex-dividend,

$$\sum_i \mu_i \sum_{v^t} \Pr_i(v^t|z^t) \sum_{k \in \mathcal{K}} a_{k,i}(s^t) = P(z^t) \sum_{z^{t+1}} Q(z^{t+1}|z^t) V(P(z^t), z^{t+1}).$$

4.2 Equilibrium representation

We now show that the class of HANK models considered here has a common representation for the law of motion of aggregate variables. Specifically, the equilibrium conditions in the model can be divided in a *macro block*, a standard representative agent New Keynesian model with taste shocks, and a *micro block* that captures the evolution of these taste shocks.

Furthermore, we show that the taste shocks are a function of the idiosyncratic wedges we defined in Proposition 1.

To this end, we define

$$\beta_i(v^t, z^{t+1}) \equiv \beta \sum_{v_{t+1}} \Pr_i(v_{t+1}|v^t, z^{t+1}) [1 + \tau_{a,i}(s^t, v_{t+1}, z_{t+1})] \quad (20)$$

$$\omega(z^t) \equiv \left[\sum_i \mu_i \sum_{v^t} \Pr_i(v^t|z^t) \varphi_i(z^t, v^t)^{\frac{-\sigma}{\psi}} \theta(v_t)^{\frac{1+\psi}{\psi}} (1 - \tau_{l,i}(s^t))^{\frac{1}{\psi}} \right]^{-\psi}. \quad (21)$$

Recall that $\varphi_i(z^t, v^t)$ is the consumption share of an individual of type i with history s^t and it is related to the risk sharing wedges according to equation (9). We then have the following proposition:

Proposition 2. *Given $\{\beta_i(v^t, z^{t+1}), \omega(z^t)\}$ defined in (20) and (21), the equilibrium aggregate consumption, output, inflation, and nominal interest rate, $\{C(z^t), Y(z^t), \Pi(z^t), i(s^t), Q(z^{t+1})\}$ must satisfy the aggregate Euler equation,*

$$\frac{1}{1+i(z^t)} = \max_{i, v^t} \sum_{z^{t+1}} \left\{ \Pr(z^{t+1}|z^t) \frac{\beta_i(v^t, z^{t+1})}{\Pi(z^{t+1})} \left(\frac{C(z^{t+1})}{C(z^t)} \right)^{-\sigma} \right\}, \quad (22)$$

the Phillips curve,

$$\tilde{\Pi}(z^t) = \frac{Y(z^t)}{\kappa(\mu-1)} \left[\mu \chi \frac{Y(z^t)^\psi}{A(z^t)^{1+\psi}} C(z^t)^\sigma \omega(z^t) - 1 \right] + \sum_{z^{t+1}} Q(z^{t+1}|z^t) \tilde{\Pi}(z^{t+1}) \quad (23)$$

the Taylor rule (19), the resource constraint

$$Y(z^t) = C(z^t) + \frac{\kappa}{2} [\Pi(z^t) - 1]^2, \quad (24)$$

and the real state prices

$$Q(z^{t+1}|z^t) = \max_{i, v^t} \left\{ \beta_i(v^t, z^{t+1}) \Pr(z^{t+1}|z^t) \left(\frac{C(z^{t+1})}{C(z^t)} \right)^\sigma \right\}. \quad (25)$$

The proof for this result is straightforward. The aggregate Euler equation (22) is obtained by substituting the definition of the risk-sharing wedges in the individual Euler equation (13), and noting that under our assumptions it holds with equality for the agent with the highest marginal valuation of the bond—the "max" in equation (22). The Phillips curve (23) can be derived by substituting for the wage in (18) using the individual labor supply

decisions. Indeed, multiplying both side of equation (14) by $e(v_t)/C(z^t)^{-\sigma/\psi}$ and averaging both sides across individuals we obtain

$$w(z^t)^{\frac{1}{\psi}} \left[\sum_i \mu_i \sum_{v^t} \Pr_i(v^t|z^t) \varphi_i(z^t, v^t)^{\frac{-\sigma}{\psi}} e(v_t)^{\frac{1+\psi}{\psi}} (1 - \tau_{l,i}(s^t))^{\frac{1}{\psi}} \right] = \chi^{\frac{1}{\psi}} \left[\sum_i \mu_i \sum_{v^t} \Pr_i(v^t|z^t) e(v_t) l_i(s^t) \right] C(z^t)^{\frac{\sigma}{\psi}}.$$

We can then use the production function (11) and the fact that $e(v_t)$ equals the efficiency wedge $\theta(v_t)$ to express the real wage as

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

and substitute it in equation (18) to obtain the Phillips curve (23). To obtain (24) we simply substitute the adjustment costs (16) in the resource constraint.

Equations (19), (22), (23), (24) and (25) are equivalent to those of a representative agent economy with “taste shocks”: that is, shocks to the rate of time preferences and to the disutility of labor. Thus, the effects that micro heterogeneity has on macroeconomic variables can be represented in this class of models as if the household in an equivalent representative agent economy has become more/less patient or more/less inclined to work. See [Nakajima \(2005\)](#), [Krueger and Lustig \(2010\)](#), and [Werning \(2015\)](#) for related results.

Proposition 2 has two main implications. The first implication is that $\beta_i(v^t, z^{t+1})$ and $\omega(z^t)$ defined in equations (20) and (21) summarize all the information from the “micro block” of the model that is needed to characterize the behavior of aggregate variables. That is, we do not need to know the primitives of the model regarding the set of assets traded by households, when or how fiscal transfers are distributed, or the nature of their financial constraints to characterize the behavior of macro aggregates, as long as we know how $\{\beta_i(v^t, z^{t+1}), \omega(z^t)\}$ evolve. Second, there is a mapping between the micro wedges that we have defined in Section 2 and $\{\beta_i(v^t, z^{t+1}), \omega(z^t)\}$, a relation that is invariant to the set of assets traded by households and the nature of their financial constraints.

4.3 Examples

Before moving to show how to combine the measured micro wedges with the representation in Proposition 2 to perform counterfactuals, we further illustrate the main result of this section with two examples. We consider two simple examples of HANK models, and derive for both the representation of Proposition 2. In particular, we show how the two simple structural models define a mapping from shocks to $\beta_i(v^t, z^t)$. Both examples are motivated by recent research suggesting that microeconomic frictions might have been important factors behind the Great Recession.

The first example isolates the implications of a rise in idiosyncratic income risk on precautionary saving motives of households, a mechanism studied by [Ravn and Sterk \(2017\)](#) and [Heathcote and Perri \(2018\)](#) among others. The second example shows how a tightening of credit constraints at the micro level can lead to a fall in aggregate demand, see [Eggertsson and Krugman \(2012\)](#) and [Guerrieri and Lorenzoni \(2017\)](#).

Example 1: Precautionary motives and aggregate demand. Let assume that there is only one type of agent and drop the subscript i from allocations to economize on notation. Let $\sigma = 1$ and the idiosyncratic productivity shocks evolve according to

$$\Delta \log[e(v_t)] = -\frac{\sigma_e(z_t)}{2} + \varepsilon_t \quad \varepsilon_t | z^t \sim \mathcal{N}(0, \sigma_e(z_t)).$$

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 z^t : when $\sigma_e(z^{t+1})$ is high, households face higher idiosyncratic risk.

To obtain analytical expressions for the $\beta(v^t, z^{t+1})$ and the $\omega(z^t)$ implied by this model, we set the taxes on labor to zero and we also assume that households can only trade the risk-free bond and that they face the borrowing limit $b(s^t) \geq 0$. Because households cannot trade stocks of the firms, we also assume that the government levies taxes on the intermediate good producers and transfers the profits to the households in proportion to the realization of idiosyncratic productivity, $e(v_t)T(z^t)$.

The tight borrowing limits, coupled with the fact that bonds are in zero net-supply, implies that households in equilibrium cannot save.¹⁸ Thus, every household is hand-to-mouth and consumes every period all his cash on hand

$$c(s^t) = e(v_t) [w(z^t)l(s^t) + T(z^t)].$$

Furthermore, we can verify from the labor supply condition (14) and $\sigma = 1$ that $l(s^t)$ is the same across individuals. So, it must also be that $c(s^t) = e(v_t)C(z^t)$ from the aggregate resource constraint.

Given the equilibrium consumption function, the risk sharing wedges are just functions of the income process,

$$\tau_a(s^t, s_{t+1}) = \frac{e(v_t)}{e(v_{t+1})} - 1,$$

and the consumption shares of an individual with history s^t is $\varphi(s^t) = \theta(v_t)$.

¹⁸The 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 example.

Substituting these expressions in equation (20) and (21) we can compute the implied $\beta(v^t, z^{t+1})$ and $\omega(z^t)$ in this specific HANK model:

$$\begin{aligned}\beta(v^t, z^{t+1}) &= \beta \sum_{v^{t+1}} \Pr(v^{t+1} | v^t, z^{t+1}) \exp \{-\Delta \log[e(v_{t+1})]\} \\ &= \beta \exp\{\sigma_e(z^{t+1})\}\end{aligned}$$

and $\omega(z^t) = 1$. Note that in this example $\beta(v^t, z^{t+1})$ does not depend on individual histories while $\omega(z^t)$ does not vary over time. These two expressions, coupled with equations (19), (22), (23) and (24), are enough to characterize the law of motion for aggregate variables in this specific example.

This representation is useful to understand how the interaction between idiosyncratic risk and incomplete financial markets can affect aggregate variables in HANK models. Suppose that households face today higher idiosyncratic risk, that is they expect higher $\sigma_e(z^{t+1})$. If financial markets were complete, this shocks would not have any effects on the allocation. Because of incomplete financial markets, however, households have a precautionary motive to save in the risk-free bond. This increase in the propensity to save at the micro level can be represented as an increase in the discount factor in the equivalent representative agent New Keynesian model.¹⁹

Example 2: Credit constraints and aggregate demand. Consider now an economy where the debt limit depends on aggregate conditions. For simplicity, assume that there are only two types of agents, $i = 1, 2$ of equal measure. When the aggregate state z^t realizes, one the types samples a high efficiency labor units, e_H , while households in the other type draw e_L . We further assume that households within types have no further idiosyncratic shocks to their efficiency of labor and that profits from the monopolistic competitive firms are distributed to households proportionally to $e_i(z_t)$.

Households can trade only a non-contingent bond in zero-net supply, subject to the debt limit $\phi(z^t)$. Thus, asset holdings must be such that

$$b_i(z^t) \geq -\phi(z^t)$$

Assuming that $\phi(z^t)$ is sufficiently small so that the debt limit is binding for the agents with a low realization of the individual productivity shock, e_L , the individual consumption

¹⁹Note that if there is no time variation in idiosyncratic risk then there is no variation in $\beta(v^t, z^{t+1})$ and the heterogeneous agent economy is equivalent to a representative agent economy with a different discount factor as shown in [Krueger and Lustig \(2010\)](#).

allocations are given by

$$c_i(z^t) = \begin{cases} e_L C(z^t) + \frac{b_i(z^{t-1})}{\Pi(z^t)} + \frac{\phi(z^t)}{1+i(z^t)} & \text{if } e_i(z_t) = e_L \\ e_H C(z^t) + \frac{b_i(z^{t-1})}{\Pi(z^t)} - \frac{\phi(z^t)}{1+i(z^t)} & \text{if } e_i(z_t) = e_H \end{cases}$$

where $b_i(z^{t-1})$ depends on the particular history:

$$b_i(z^{t-1}) = \begin{cases} \phi(z^{t-1}) & \text{if } e_i(z_{t-1}) = e_H \\ -\phi(z^{t-1}) & \text{if } e_i(z_{t-1}) = e_L \end{cases}$$

We can then express the $\beta_i(z^{t+1})$ for the type that attains the maximum in the aggregate Euler equation (22) – the agent with a high realization of the individual productivity in z^t – as

$$\beta_i(z^{t+1}) = \left(\frac{e_i(z_{t+1}) + \left(\frac{\phi(z_t)}{\Pi(z^{t+1})} - \frac{b_i(z^{t+1})}{1+i(z^{t+1})} \right) / C(z^{t+1})}{e_H(z_t) + \left(\frac{b_i(z^{t-1})}{\Pi(z^t)} - \frac{\phi(z_t)}{1+i(z^t)} \right) / C(z^t)} \right)^{-\sigma}.$$

From the expression above, it is evident how a tightening of the debt limit, a reduction in $\phi(z^t)$, increases $\beta_i(z^t, z_{t+1})$ for the marginal agent. Intuitively, a reduction in $\phi(z^t)$ means that agent with a low income shock can borrow less to smooth their consumption. In equilibrium, this implies that the agent with currently higher income must save less and consume more to clear the asset market. The increase in current consumption share make this agent more willing to save and thus the measured $\beta_i(z^t)$ increases for the marginal agent.

Further notice that the above expression for $\beta_i(z^t, z_{t+1})$ shows how the “micro-block” is not independent from the “macro-block” that determines the dynamics of aggregates as the $\beta_i(z^t, z_{t+1})$ depends on aggregate consumption, the inflation rate, and the policy rate. Conversely, from Proposition 2, the evolution of these aggregates is affected by $\beta_i(z^t)$. Despite this caveat, we can still use the representation to assess the role of imperfect risk sharing in accounting for aggregate fluctuations as we next show.

4.4 The methodology at a conceptual level

We now explain how to use the representation in Proposition 2 and the micro wedges to evaluate the macroeconomic implications of imperfect risk sharing over the business cycle. For now, we assume that the probability distribution of z^t and its realization in a particular event are known and that we know how z^t affects the distribution of idiosyncratic shocks v^t , $\Pr_i(v^t|z^t)$. Suppose also that we know how $\{A(\cdot), \varepsilon_m(\cdot), \beta_i(\cdot), \omega(\cdot)\}$ depends on the under-

lying state. Thus, given the realization of z^t , we can use the representation in Proposition 2 to obtain the underlying equilibrium path for aggregate variables— output, inflation and nominal interest rates.

Our approach consists in comparing these benchmark paths to those that would arise in an economy with complete financial markets. To construct this counterfactual, we can set the risk sharing wedges to zero, $\tau_{a,i}(s^t, s_{t+1}) = 0$ for all histories and for all i . When doing so, we obtain a different process for $\beta_i(v^{t-1}, z^t)$ and $\omega(z^t)$, given by

$$\beta_i^{\text{cm}}(v^{t-1}, z^t) = \beta \quad \omega^{\text{cm}}(z^t) = \left[\sum_i \mu_i \sum_{v^t} \Pr_i(v^t | z^t) \varphi_i(s_0) \frac{\bar{v}^{-\sigma}}{v^t} \theta(v_t)^{\frac{1+\psi}{\psi}} (1 - \tau_{l,i}(s^t))^{\frac{1}{\psi}} \right]^{-\psi} \quad (26)$$

Thus, the economy with complete financial markets features a different stochastic process for β and ω and, because of this difference, it features a different behavior for the aggregate variables. We label these the *complete markets* paths.

The comparison between the benchmark and the complete markets paths isolates the impact that imperfect risk sharing at the micro level has for macroeconomic aggregates over the particular history z^t . This experiment can help us quantifying the role of households' heterogeneity on a number of topics recently studied in the literature.

Our main application will be to quantify the importance of imperfect risk sharing during the Great Recession. As we mentioned earlier, several papers in the literature have suggested that the deep decline in real economic activity during the Great Recession was partly induced by an increase in households' propensity to save, either because of an increase in precautionary motives or because of a tightening of individual's borrowing constraints. If these mechanisms were important, we should observe the output trajectory in the benchmark to be substantially below its complete market counterfactual when feeding the history z^t that led to the Great Recession.

Importantly, our methodology does not require z^t to be an observed history. For example, we could use our framework to quantify how the transmission mechanism of monetary policy in HANK type models differs from their complete markets analog, a topic that has been recently studied by Auclert (2017) and Kaplan, Moll, and Violante (2018). To do so, we can perform these counterfactuals conditional on a history that moves ε_m at some time t while keeping all other states constant.

In this discussion we have assumed that we know $\Pr_i(s^t | s^{t-1})$, the realization of z^t in a particular event, and how it affects $\{A(\cdot), \varepsilon_m(\cdot), \beta_i(\cdot), \omega(\cdot)\}$. These assumptions were made to illustrate conceptually the nature of our counterfactuals. In practice, however, we need to estimate these stochastic processes, and we need a procedure to retrieve z^t from the data. In the following sections we discuss how to use the measured micro wedges to obtain the time

path for the taste shocks and how to implement in practice these counterfactuals.

5 Measuring the taste shocks

We now use the micro wedges measured in Section 3 and the expressions (20) and (21) to obtain empirical counterparts for β_{it} and ω_t .²⁰ Consistent with the measurement of the wedges, we set $\sigma = \psi = 1$.

From equation (20) we can see that β_{it} is the conditional expectations, across idiosyncratic histories, of the risk-sharing wedge between time $t - 1$ and t . Our approach to estimate this conditional expectation consists in grouping households with similar observable characteristics at time $t - 1$ and compute

$$\bar{\beta}_{it} \equiv \beta \frac{1}{N_i} \sum_{j=1}^{N_i} (1 + \tau_{a,it}) = \beta \frac{1}{N_i} \sum_{j=1}^{N_i} \left[\frac{C_t/C_{t-1}}{c_{jt}/c_{jt-1}} \right], \quad (27)$$

where N_i is the number of individuals in the group at time $t - 1$. The logic of our approach builds on two premises. The first is that, by grouping individuals along certain observable characteristics, we are effectively proxying for an individual history v^{t-1} . The second is that the size of the groups are large enough, so that $\bar{\beta}_{it}$ in equation (27) approximates the conditional expectation in equation (20).

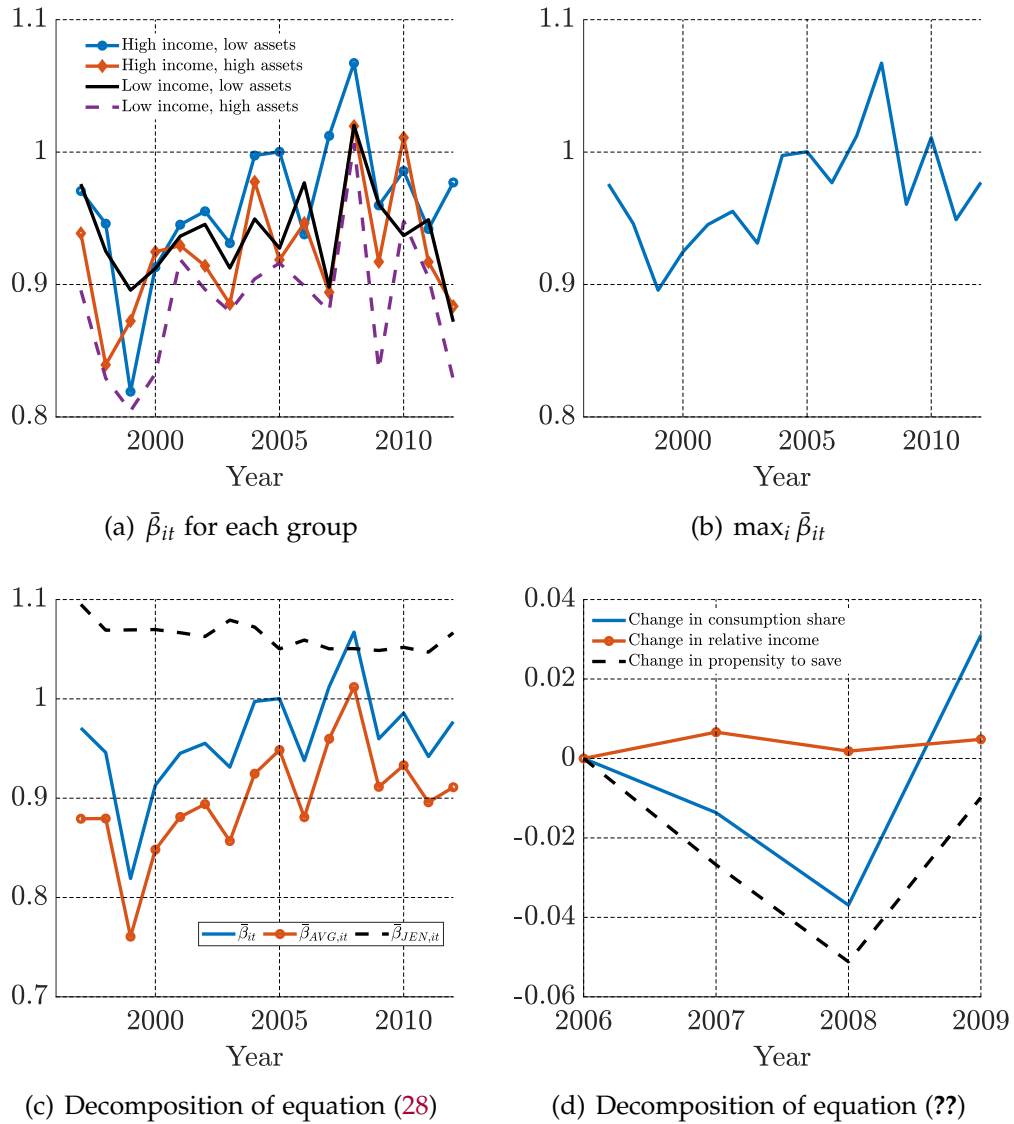
Our partitioning of households into different groups is guided by the logic of baseline incomplete markets models. In the basic version of these models, the current level of income and assets are sufficient statistics for individual histories. We follow this insight and group households according to whether their level of income at date $t - 1$ is above or below median income and, within each of these two groups, whether the level of their liquid assets is above or below the group median. Thus, for each $t - 1$, we end up with four different groups of households: *income rich/asset poor*, *income rich/asset rich*, *income poor/asset poor* and *income poor/asset rich*. For each group i , we use equation (27) to construct $\bar{\beta}_{it}$.²¹

Panel (a) in Figure 4 plots the time path of $\bar{\beta}_{it}$ in each group while panel (b) reports the year-by-year value of the maximum, which represents the relevant discount factor in equation (22) up to a first-order approximation. There are several features of the data that we want to emphasize. First, the asset poor groups have higher implicit discount factors relative to other groups in most years. Through the lens of our framework, this implies that the

²⁰We choose the CEX because we can construct wedges at annual frequency which is more appropriate for our application relative to using the bi-annual frequency of the PSID.

²¹In principle, one could consider finer partitions of the joint distribution of income and assets. However, given the sample size in the CEX, this would produce substantially noisier estimates of $\bar{\beta}_{it}$. Given our partitions, we have roughly 250 households per year within each group.

Figure 4: The evolution of $\max_i \bar{\beta}_{it}$



incentives to save of these groups are the relevant ones for determining the behavior of the risk-free rate in our sample.²² Second, the discount factor measured in our approach displays some important patterns: an upward trend starting in the early 2000s and a substantial increase during the Great Recession. That is, the behavior of the micro wedges is consistent with the idea that imperfect risk sharing was a driving force of the secular decline in real interest rates, and also with the economic mechanisms we have reviewed in Section 4.3.

Which feature of the distribution of individual risk-sharing wedges is responsible for such behavior in the data? To answer this question, we use risk-sharing wedges of the income

²²That is, individuals with low assets are those reaching the “max” in equation (22).

poor/asset poor as an illustrative example²³ and decompose their $\bar{\beta}_{it}$ as follows:

$$\bar{\beta}_{it} = \beta \underbrace{\left[\frac{C_t/C_{t-1}}{\frac{1}{N_i} \sum_{j=1}^{N_i} c_{jt}/c_{jt-1}} \right]}_{\bar{\beta}_{AVG,it}} \underbrace{\sum_{j=1}^{N_i} \left[\frac{\sum_{j=1}^{N_i} c_{jt}/c_{jt-1}}{c_{jt}/c_{jt-1}} \right]}_{\bar{\beta}_{JEN,it}}. \quad (28)$$

Mechanically, $\bar{\beta}_{it}$ can increase for two reasons. First, if on average the consumption growth of the group between date $t - 1$ and date t is smaller than consumption growth in the aggregate. This effect is captured by the term $\bar{\beta}_{AVG,it}$ in the above expression. Second, an increase in the cross-sectional dispersion of risk-sharing wedges raises $\bar{\beta}_{it}$ because of Jensen's inequality, effect that is captured by $\bar{\beta}_{JEN,it}$ in equation (28).²⁴ Panel (c) in Figure 4 presents this decomposition. We can verify that the bulk of the time-variation in the discount factor of this group is due to $\bar{\beta}_{AVG,it}$. That is, the sharp increase in the measured $\max_i \bar{\beta}_{it}$ observed during the Great Recession arises because, on average, income rich/asset poor households reduced their consumption significantly more than other households.

We now turn to the measurement of ω_t in equation (21). For each household in our panel, we compute the consumption share, $\varphi_{it} = c_{it}/C_t$ and combine it with the efficiency and labor wedge to construct $\varphi_{it}^{-\sigma/\psi} \theta_{it}^{(1+\psi)/\psi} (1 - \tau_{l,it})^{1/\psi}$. We then take the cross-sectional average for each t , and raise it to $-1/\psi$. Panel (a) in Figure 5 plots the time series for ω_{it} , normalized to be 1 in 1996. We can see that this statistic increases substantially in the 2001 and 2008 recession, a pattern that follows closely the time path of the average labor wedge reported in Appendix B.

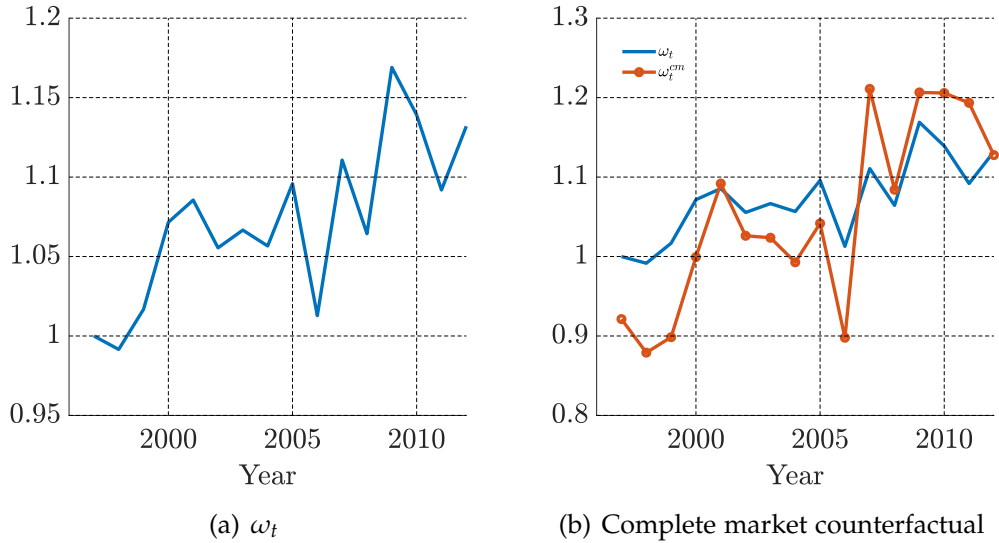
In panel (b) of the same figure we plot the counterfactual value of ω_t , reported in (26).²⁵ We can see that from 2007, and throughout the Great Recession, this statistic lies above the observed ω_t . Through the lens of our framework, this means that the impact of imperfect risk-sharing on macroeconomic aggregates during the Great Recession can be interpreted as a decrease in the marginal disutility of labor.

²³Results are quantitatively similar if one uses the income rich /asset poor group instead.

²⁴Note that in Example 1 where movements in $\bar{\beta}_{it}$ are driven by an increase in idiosyncratic income risk we have that all the variation in $\bar{\beta}_{it}$ is driven by $\bar{\beta}_{JEN,it}$ and $\bar{\beta}_{AVG,it}$ is constant over time and equal to 1. Conversely, in Example 2 where movements in $\bar{\beta}_{it}$ are driven by shocks to the borrowing limits with no idiosyncratic risk we have that all the variation in $\bar{\beta}_{it}$ is driven by $\bar{\beta}_{AVG,it}$ and $\bar{\beta}_{JEN,it}$ is constant over time and equal to 1.

²⁵To construct this counterfactual we assume that the initial distribution of consumption shares is the one observed in 1996, the first year in our dataset.

Figure 5: The evolution of ω_t



6 An application to the U.S. economy

We now discuss how to combine the measured taste shocks and the representation of Proposition 2 to measure the macroeconomic implications of imperfect risk-sharing. In Section 6.1 we discuss a Markovian implementation of the counterfactuals presented in Section 4.4. Section 6.2 reports the estimates of the structural parameters of the model and presents indicators of model fit. Section 6.3 presents the results of our counterfactual. We check the sensitivity of our results in Section 6.4.

6.1 Markovian implementation

In order to implement the counterfactuals discussed in Section 4.4 we need a procedure to estimate the stochastic processes $\{A(\cdot), \varepsilon_m(\cdot), \beta(\cdot), \omega(\cdot)\}$ and to retrieve z^t from the data. In this section we describe the restrictions that we place on these objects in order to implement our procedure.

We start by assuming a Markovian structure for the states, $\Pr_i(s_t | s^{t-1}) = \Pr_i(s_t | s_{t-1})$. Without loss of generality, we can then set $z_{1,t} = a_t = \log A_t$ and $z_{2,t} = \varepsilon_{m,t}$, and assume that the shocks are uncorrelated at all leads and lags. The vector z_t could potentially incorporate other aggregate shocks (for example, shocks affecting financial constraints), and we leave that unrestricted. Following much of the existing literature, we assume that the log of technology is an AR(1) process

$$a_t = \rho_a a_{t-1} + \varepsilon_{a,t}, \quad \varepsilon_{a,t} \sim \mathcal{N}(0, \sigma_a)$$

and $\varepsilon_{m,t} \sim \mathcal{N}(0, \sigma_m)$.

Making assumptions about $\beta_i(v^t, z^{t+1})$ and $\omega(z^t)$ is conceptually more problematic, because those are not fundamental shocks and do not necessarily inherit the Markov structure of s_t .²⁶ To explain the nature of the problem, suppose that households in the economy have only access to the risk-free nominal bond and they face a borrowing limit $b_i(s^t) \geq -\phi$. In a recursive competitive equilibrium of this framework, we know that the distribution of assets is an aggregate state and follows itself a Markov process. The implied process for $\beta_i(\cdot)$ and $\omega(\cdot)$, however, will be a VMA(∞), which is not a feasible process for estimation. In time series analysis, these processes are typically approximated with finite order ones. We follow this strategy and approximate the law of motion for these processes using a VAR(1). Specifically, let $\mathbf{T}_t = [\log(\beta_{it}), \log(\omega_t)]'$. We assume that \mathbf{T}_t follows the process

$$\mathbf{T}_t = (1 - \Phi)\mu + \Phi\mathbf{T}_{t-1} + \Gamma_a a_t + \Gamma_m \varepsilon_{m,t} + \eta_t, \quad (29)$$

where $\mu = [\mu_{\beta_i}, \mu_{\omega}]'$ represents the mean of the process, $\eta_t \sim \mathcal{N}(\mathbf{0}, \Sigma)$, and (Γ_a, Γ_m) representing loadings of the structural shocks on \mathbf{T}_t . The choice of a VAR(1) is arbitrary, and one could think about alternative time series models to approximate the VMA(∞) representation of \mathbf{T}_t , as for example factor models, VARs with time-varying parameters, etc.

There are two features of the system in (29) that we wish to emphasize. First, the process for β_{it} is individual-specific, as it depends on v^t . Given our strategy of grouping households, we should in principle include in \mathbf{T}_t the measured $\bar{\beta}_{it}$ for each group rather than just including $\max_i \bar{\beta}_{it}$. From a theoretical point of view, this might make a difference. Indeed, it could be that the group with the highest $\bar{\beta}_{it}$ in a given year is not the one with the highest incentives to save in equation (22), either because $\mathbb{E}_t[\beta_{it+1}]$ is not the highest or because there is another group with a $\bar{\beta}_{jt+1}$ more negatively related to aggregate consumption growth.²⁷ However, because for most periods in our sample the income-rich/asset poor households are the one achieving the max, and because their measured $\bar{\beta}_{it}$ is the most cyclical, we will only include $\max_i \bar{\beta}_{it}$ in \mathbf{T}_t . This simplification allows us to economize on the number of state variables when solving the equivalent representative-agent economy with nonlinear methods, which is important in our application because of the binding zero lower bound on nominal interest rates.

Second, we allow the structural shocks a_t and $\varepsilon_{m,t}$ to contemporaneously affect \mathbf{T}_t : this is because β and ω are in potentially functions of these aggregate shocks. This also explains why the errors in η_t can be contemporaneously correlated, a feature that captures the

²⁶A similar issue arises in the accounting procedure of Chari, Kehoe, and McGrattan (2007), see their discussion in Section 2.

²⁷Up to a first-order approximation, this second possibility can be ruled out because the Euler equation will necessarily hold for the household with the highest $\max_i \mathbb{E}_t[\beta_{it+1}]$.

possibility that other structural shocks might influence, at the same time, β and ω .

Given the stochastic processes for $\mathbf{X}_t = [a_t, \varepsilon_{m,t}, \mathbf{T}_t]'$ and the value of models' parameters, we can solve for the policy functions of aggregate variables as a function of the state $\mathbf{S}_t = [i_{t-1}, Y_{t-1}, \mathbf{X}_t]$ using the system of equations in Proposition 2. For this purpose, we employ the projection algorithm developed in [Gust, Herbst, López-Salido, and Smith \(2017\)](#) for studying a representative agent New Keynesian model with an occasionally binding zero lower bound constraint, see Appendix C for a detailed description of the algorithm.

After this step, it is conceptually straightforward to implement the counterfactuals described earlier. Specifically, we carry out three main steps.

In the first step, and using data on \mathbf{T}_t , nominal interest rates, output and inflation, we jointly estimate the parameters of the stochastic process \mathbf{X}_t and the structural parameters that govern preferences, technology, the importance of price adjustment costs and the behavior of the monetary authority.

In the second step, we apply the particle filter to the estimated model and retrieve the sequence of states that rationalizes the behavior of the observables in our sample. This step guarantees that, up to a measurement error, the model replicates the behavior of aggregate *and* micro-level data, as summarized by \mathbf{T}_t .²⁸ Moreover, it provides us with an estimate of the path of the latent structural shocks, $\{a_t, \varepsilon_{m,t}\}$.

The third and final step consists in obtaining the counterfactual complete markets paths for aggregate variables. To do so, we can compute the policy functions of the model with complete financial markets, which is defined by the system of equation in Proposition 2 and a law of motion for $\mathbf{X}_t^{\text{cm}} = [a_t, \varepsilon_{m,t}, \mathbf{T}_t^{\text{cm}}]$, with \mathbf{T}_t^{cm} constructed using the expressions in (26). We then obtain the counterfactual path for aggregate variables by feeding through these policy functions the realization of the structural shocks estimated earlier and \mathbf{T}_t^{cm} .

6.2 Estimation

The model is estimated at an annual frequency on the 1997-2012 period. The aggregate data used in estimation are observations on output, inflation and nominal interest rates. We map the log of output in the model, \hat{Y}_t , to the percentage deviations of log real GDP per capita from a linear deterministic trend. The inflation rate $\pi_t = \log(\Pi_t)$ is the annual percent change in the consumer price index, and i_t is mapped to the annual effective federal funds rate. The other observables that we use are $\max_i \bar{\beta}_{it}$ and ω_t .

The structural parameters are the ones governing preferences, $[\beta, \sigma, \nu, \chi]$, the importance

²⁸The presence of measurement errors is required to implement the particle filter and evaluate a likelihood function for the model, as we have less stochastic variables than observables in the state-space system.

of price adjustment costs κ , the elasticity of substitution across varieties μ , the behavior of the monetary authority $[\rho_i, \gamma_\pi, \gamma_{\Delta y}, \Pi^*]$ and the stochastic process of \mathbf{X}_t .

We fix a subset of these parameters at conventional values. Specifically, we set $\sigma = 1$ and $\nu = 1$. We also set $\mu = 1.2$, a conventional value in the literature. The parameters χ and $\exp\{\mu_\omega\}$ enter as a multiplication in the Phillips curve. To minimize the number of parameters to be estimated, we de-mean the $\log(\omega_t)$ series and set $\mu_\omega = 0$. The parameter χ is then chosen such that, given the other parameters, consumption equals 1 in a deterministic steady state of the model. In addition, we fix $\exp\{\mu_\beta\}$ to 0.98, the mean of $\max_i \bar{\beta}_{it}$ over the sample and we set the target for annual inflation to 2%.

The remaining parameters, $[\kappa, \Pi^*, \rho_i, \gamma_\pi, \gamma_{\Delta y}]$ and $[\rho_a, \sigma_a, \sigma_m, \mu_\omega, \Phi, \Sigma, \Gamma_a, \Gamma_m]$ are estimated with Bayesian methods. Let $\mathbf{Y}_t = [\hat{Y}_t, i_t, \pi_t, \mathbf{T}_t]$ be the vector of observable variables, and denote by \mathbf{Y}^T all the observations in our sample. Remember that $\mathbf{S}_t = [i_{t-1}, \hat{Y}_{t-1}, A_t, \varepsilon_{m,t}, \mathbf{T}_t]$ denotes the state vector. The model of Section 4 defines the non-linear state space model

$$\begin{aligned}\mathbf{Y}_t &= \mathbf{g}(\mathbf{S}_t; \theta) + e_t \\ \mathbf{S}_t &= \mathbf{f}(\mathbf{S}_{t-1}, \varepsilon_t; \theta).\end{aligned}\tag{30}$$

In the above notation, $\mathbf{g}(\cdot)$ and $\mathbf{f}(\cdot)$ represent the policy functions of the model, θ the vector of parameters to be estimated, ε_t collecting the innovations to the stochastic variables of the model and e_t collecting uncorrelated Gaussian errors that capture any deviation between the data \mathbf{Y}_t and $\mathbf{g}(\mathbf{S}_t; \theta)$. In estimation, we fix the variance of these measurement errors to 5% of the unconditional variance of their corresponding series.

Given this representation, we can apply standard filtering techniques to the above state-space system and evaluate the likelihood of the model, $\mathcal{L}(\theta|\mathbf{Y}^T)$. We can then combine this information with a prior for the structural parameters, $p(\theta)$, and apply a standard random walk Metropolis-Hastings algorithm to sample from the posterior distribution (An and Schorfheide, 2007)

$$p(\theta|\mathbf{Y}^T) \propto p(\theta)\mathcal{L}(\theta|\mathbf{Y}^T).$$

For tractability, in this preliminary version of the paper we solve for the policy functions with a first-order perturbation when estimating the model. The first-order perturbation solution is much faster and numerically more stable than the global approximation discussed in Appendix C, and it allows us to use the Kalman filter for the evaluation of the likelihood function. The main drawback is that we are not considering the possibility of a binding zero lower bound constraint on nominal interest rates when estimating the model's parameters.

Table 1 reports the results of the estimation. The structural parameters defining the behavior of the monetary authority and the degree of price stickiness are in line with previous

Table 1: Prior and posterior for estimated model parameters

Parameter	Distribution	Prior		Posterior	
		Mean	Standard deviation	Mean	90% Interval
$4 \times \kappa$	Gamma	85.00	15.00	90.79	[67.41, 115.75]
ρ_i	Beta	0.50	0.28	0.57	[0.29, 0.85]
γ_π	Normal	0.00	1.00	1.51	[0.85, 2.14]
$\gamma_{\Delta y}$	Normal	0.00	1.00	0.58	[0.21, 0.92]
ρ_a	Beta	0.50	0.28	0.52	[0.25, 0.79]
$\Phi_{\beta,\beta}$	Beta	0.50	0.28	0.76	[0.61, 0.92]
$\Phi_{\omega,\omega}$	Beta	0.50	0.28	0.73	[0.33, 0.99]
$100 \times \sigma_a$	InvGamma	1.00	1.00	7.23	[3.36, 10.39]
$100 \times \sigma_m$	InvGamma	1.00	1.00	1.91	[0.98, 2.81]
$100 \times \sigma_\beta$	InvGamma	1.00	1.00	1.88	[1.13, 2.60]
$100 \times \sigma_\omega$	InvGamma	1.00	1.00	3.80	[2.54, 5.05]

estimates reported in the literature. For instance, they are comparable to the estimates reported in the working paper version of [Gust et al. \(2017\)](#), who used aggregate quarterly data to estimate a similar version of our model – a representative agent three-equations New Keynesian with technology shocks, discount factor shocks and monetary policy shocks.

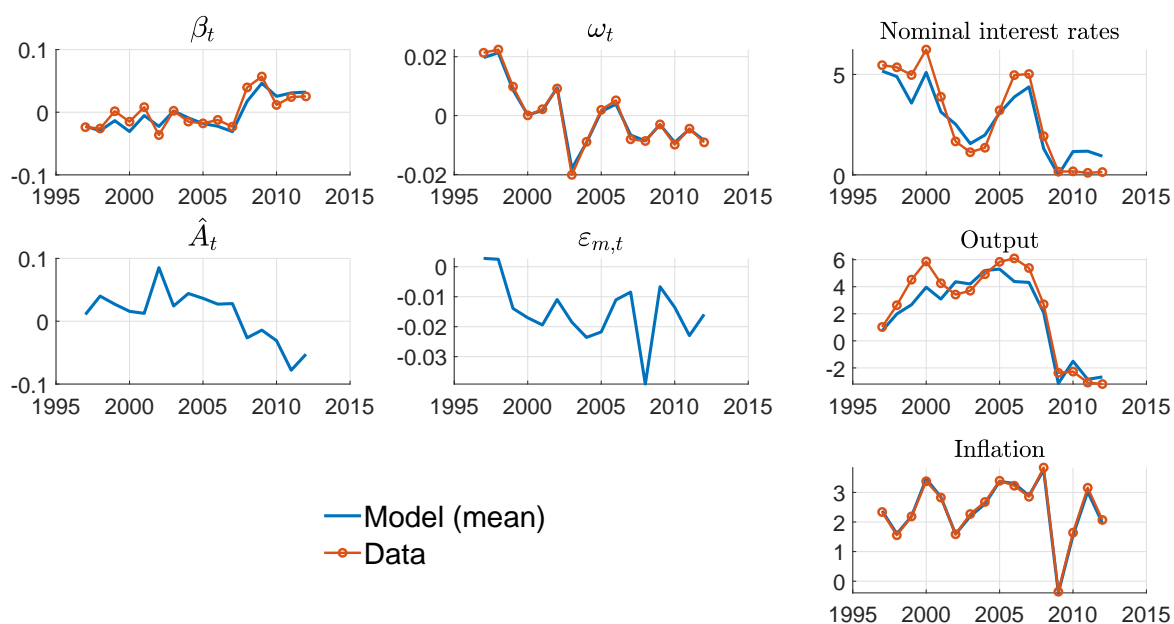
6.3 Assessing the macroeconomic consequences of imperfect risk sharing

(Need updating)

Having estimated the model, we can now perform the main counterfactual of the paper. The first step consists in using the state-space representation (30) along with our observables \mathbf{Y}^T to estimate the latent variables \mathbf{S}_t . Because the focus of our counterfactual is the Great Recession, we apply our nonlinear solution algorithm that explicitly accounts for a binding zero lower bound constraint on nominal interest rate. When doing so, we set the structural parameters at their posterior mean in Table 1. The state vector is estimated by applying the particle filter to our non-linear state space model. Because nominal interest rates, output and \mathbf{T}_t are observed in our procedure, the true latent states are the technology and the monetary policy shocks.

Figure 6 reports the estimated latent state, the observables implied by the model, and their data counterpart. Despite having fewer shocks than observables, the model does a remarkable job in fitting both the micro data, as summarized by the behavior of \mathbf{T}_t , and the aggregate data on inflation, output and nominal interest rates. We label the model implied trajectories obtained in this first step as the “baseline”.

Figure 6: Estimating the state vector



The second step consists in solving for the policy functions of the “complete markets” version of the model, that is the one in which $\beta_{i,t}$ and ω_t evolve according to (26).²⁹ After obtaining these policy functions, we construct the counterfactual path for the observables by feeding these policy functions with the path for $[\hat{A}_t, \varepsilon_{m,t}]$ reported in Figure 6 and the path for ω_t^c .

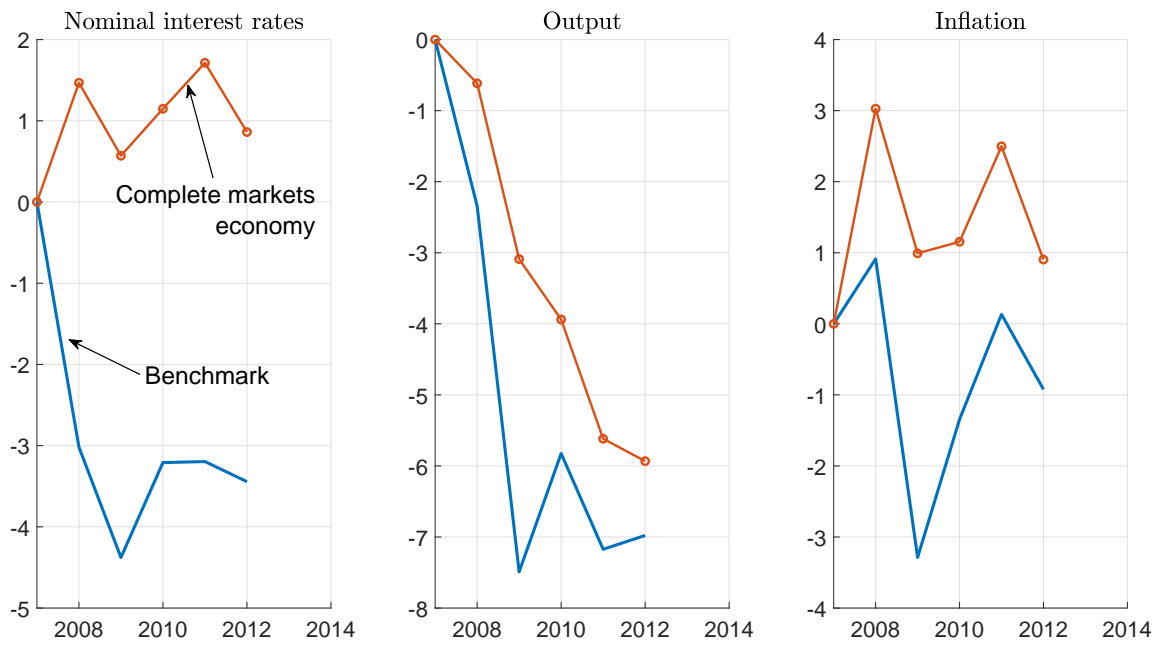
Figure 7 compares the trajectories for output, inflation and the nominal interest rate in this counterfactual with the ones in the baseline. As explained in Section 4.4, the difference between these two paths isolate the effects that imperfect risk sharing has for macroeconomic aggregates. Relative to the baseline, the counterfactual economy is associated to a fairly flat profile for nominal interest rate and inflation, and to a milder decline in output in the 2007-2009. Specifically, we find that imperfect risk sharing at the households’ level was responsible for the observed fall in nominal interest rates and inflation, and for roughly 50% of the severity of the recession.

6.4 Sensitivity analysis

(To do)

²⁹Essentially, the complete markets version has a constant discount factor for the representative households, and different parameters for the ω_t process. These parameters are estimated by fitting an AR(1) to ω_t^c .

Figure 7: Baseline vs. Complete Markets economy



7 Conclusion

References

- Acharya, Sushant and Keshav Dogra. 2018. "Understanding HANK: Insights from a PRANK." *Federal Reserve Bank of New York Staff Report* 835.
- Aguiar, Mark and Mark Bilal. 2015. "Has consumption inequality mirrored income inequality?" *American Economic Review* 105 (9):2725–56.
- Alvarez, Fernando and Urban J Jermann. 2000. "Efficiency, equilibrium, and asset pricing with risk of default." *Econometrica* 68 (4):775–797.
- An, Sungbae and Frank Schorfheide. 2007. "Bayesian analysis of DSGE models." *Econometric reviews* 26 (2-4):113–172.
- Attanasio, Orazio P and Luigi Pistaferri. 2016. "Consumption inequality." *Journal of Economic Perspectives* 30 (2):3–28.
- Auclert, Adrien. 2017. "Monetary policy and the redistribution channel." Tech. rep., National Bureau of Economic Research.
- Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2018. "The Intertemporal Keynesian Cross." Tech. rep., January 2017 working paper.
- Bayer, Christian, Ralph Lütticke, Lien Pham-Dao, and Volker Tjaden. 2019. "Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk." *Econometrica* 87 (1):255–290.
- Bils, Mark, Peter J Klenow, and Benjamin A Malin. 2018. "Resurrecting the role of the product market wedge in recessions." *American Economic Review* 108 (4-5):1118–46.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption inequality and partial insurance." *American Economic Review* 98 (5):1887–1921.
- Boerma, Job and Loukas Karabarbounis. 2019. "Inferring Inequality with Home Production." Tech. rep., University of Minnesota.
- Brav, Alon, George M Constantinides, and Christopher C Geczy. 2002. "Asset pricing with heterogeneous consumers and limited participation: Empirical evidence." *Journal of Political Economy* 110 (4):793–824.
- Broer, Tobias, Niels-Jakob H. Hansen, Per Krusell, and Erik Oberg. 2018. "The New Keynesian Transmission Mechanism: A Heterogeneous-Agent Perspective." *Review of Economic Studies* forthcoming.

- Chari, Varadarajan V, Patrick J Kehoe, and Ellen R McGrattan. 2007. "Business cycle accounting." *Econometrica* 75 (3):781–836.
- Chien, YiLi, Harold Cole, and Hanno Lustig. 2011. "A multiplier approach to understanding the macro implications of household finance." *The Review of Economic Studies* 78 (1):199–234.
- . 2012. "Is the volatility of the market price of risk due to intermittent portfolio rebalancing?" *American Economic Review* 102 (6):2859–96.
- Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo. 2011. "When is the government spending multiplier large?" *Journal of Political Economy* 119 (1):78–121.
- Eggertsson, Gauti B and Paul Krugman. 2012. "Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach." *The Quarterly Journal of Economics* 127 (3):1469–1513.
- Erceg, Christopher J, Dale W Henderson, and Andrew T Levin. 2000. "Optimal monetary policy with staggered wage and price contracts." *Journal of monetary Economics* 46 (2):281–313.
- Fitzgerald, Doireann. 2012. "Trade costs, asset market frictions, and risk sharing." *American Economic Review* 102 (6):2700–2733.
- Guerrieri, Veronica and Guido Lorenzoni. 2017. "Credit crises, precautionary savings, and the liquidity trap." *The Quarterly Journal of Economics* 132 (3):1427–1467.
- Gust, Christopher, Edward Herbst, David López-Salido, and Matthew E Smith. 2017. "The empirical implications of the interest-rate lower bound." *American Economic Review* 107 (7):1971–2006.
- Heathcote, Jonathan and Fabrizio Perri. 2018. "Wealth and volatility." *The Review of Economic Studies* 85 (4):2173–2213.
- Heathcote, Jonathan, Fabrizio Perri, and Gianluca Violante. 2010. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States 1967-2006." *The Review of Economic Dynamics* 13 (1):15–51.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante. 2014. "Consumption and labor supply with partial insurance: An analytical framework." *American Economic Review* 104 (7):2075–2126.
- Hsieh, Chang-Tai and Peter J Klenow. 2009. "Misallocation and manufacturing TFP in China and India." *The Quarterly journal of economics* 124 (4):1403–1448.

- Huggett, Mark. 1993. "The risk-free rate in heterogeneous-agent incomplete-insurance economies." *Journal of Economic Dynamics and Control* 17 (5-6):953–969.
- Jappelli, Tullio and Luigi Pistaferri. 2010. "The consumption response to income changes." .
- Kaplan, Greg, Benjamin Moll, and Giovanni L Violante. 2018. "Monetary policy according to HANK." *American Economic Review* 32 (3):697–743.
- Kaplan, Greg and Giovanni L Violante. 2014. "A model of the consumption response to fiscal stimulus payments." *Econometrica* 82 (4):1199–1239.
- . 2018. "Microeconomic Heterogeneity and Macroeconomic Shocks." *Journal of Economic Perspectives* 108 (3):167–194.
- Karabarbounis, Loukas. 2014. "The labor wedge: MRS vs. MPN." *Review of Economic Dynamics* 17 (2):206–223.
- Karaivanov, Alexander and Robert M Townsend. 2014. "Dynamic financial constraints: Distinguishing mechanism design from exogenously incomplete regimes." *Econometrica* 82 (3):887–959.
- Kocherlakota, Narayana and Luigi Pistaferri. 2009. "Asset pricing implications of Pareto optimality with private information." *Journal of Political Economy* 117 (3):555–590.
- Krueger, Dirk and Hanno Lustig. 2010. "When is market incompleteness irrelevant for the price of aggregate risk (and when is it not)?" *Journal of Economic Theory* 145 (1):1–41.
- Krueger, Dirk, Hanno Lustig, and Fabrizio Perri. 2008. "Evaluating asset pricing models with limited commitment using household consumption data." *Journal of the European Economic Association* 6 (2-3):715–726.
- Krueger, Dirk, Kurt Mitman, and Fabrizio Perri. 2016. "Macroeconomics and household heterogeneity." In *Handbook of Macroeconomics*, vol. 2. Elsevier, 843–921.
- Krueger, Dirk and Fabrizio Perri. 2006. "Does income inequality lead to consumption inequality? Evidence and theory." *The Review of Economic Studies* 73 (1):163–193.
- Krusell, Per and Anthony A Smith, Jr. 1998. "Income and wealth heterogeneity in the macroeconomy." *Journal of political Economy* 106 (5):867–896.
- McKay, Alisdair, Emi Nakamura, and Jón Steinsson. 2016. "The power of forward guidance revisited." *American Economic Review* 106 (10):3133–58.
- McKay, Alisdair and Ricardo Reis. 2016. "The role of automatic stabilizers in the US business cycle." *Econometrica* 84 (1):141–194.

- Moll, Benjamin, Robert M Townsend, and Victor Zhorin. 2017. "Economic development, flow of funds, and the equilibrium interaction of financial frictions." *Proceedings of the National Academy of Sciences* :201707055.
- Nakajima, Tomoyuki. 2005. "A business cycle model with variable capacity utilization and demand disturbances." *European Economic Review* 49 (5):1331–1360.
- Ohanian, Lee E, Paulina Restrepo-Echavarria, and Mark LJ Wright. 2018. "Bad Investments and Missed Opportunities? Postwar Capital Flows to Asia and Latin America." *American Economic Review* 108 (12):3541–82.
- Ravn, Morten O and Vincent Sterk. 2017. "Job uncertainty and deep recessions." *Journal of Monetary Economics* 90:125–141.
- Rotemberg, Julio J. 1982. "Sticky prices in the United States." *Journal of Political Economy* 90 (6):1187–1211.
- Vissing-Jørgensen, Annette. 2002. "Limited asset market participation and the elasticity of intertemporal substitution." *Journal of political Economy* 110 (4):825–853.
- Werning, Iván. 2015. "Incomplete markets and aggregate demand." Tech. rep., National Bureau of Economic Research.

ONLINE APPENDIX TO "ACCOUNTING FOR HETEROGENEITY"

BY DAVID BERGER, LUIGI BOCOLA AND ALESSANDRO DOVIS

A Detailed economies and micro wedges

In this Appendix we present three examples of detailed economies and highlight their implications for the wedges of the prototype model.

A.1 A bond economy

We consider an economy along the lines of [Huggett \(1993\)](#). We assume that there is only one type of household with preferences as in the prototype model.³⁰ The only asset available is an uncontingent bond in zero net supply, and we assume that agents face the debt limit

$$a(s^t) \geq -\phi. \tag{A.1}$$

The technology to produce the final good uses labor as the only inputs

$$Y(z^t) = A(z_t) \sum_{v^t} \Pr(v^t | v^{t-1}, z^t) e(v_t) l(v^t, z^t),$$

and firms are perfectly competitive.

An equilibrium is an allocation and a real risk free rate $R(z^t)$ such that: i) the allocation maximizes (1) subject to the budget constraint

$$c(s^t) + a(s^t) \leq W(s^t) l(s^t) + a(s^{t-1}) R(z^t),$$

and the debt limit (A.1); ii) the real wage is given by $W(s^t) = e(v_t) A(z_t)$; iii) the bond market clears in that $\sum_{v^t} \Pr(v^t | z^t) a(z^t, v^t) = 0$

The equilibrium for this economy is characterized by the household's optimality conditions: $U_c(s^t) e(v_t) = -U_l(s^t)$ and the Euler equation

$$U_c(s^t) \geq \beta R(z^t) \sum_{s^{t+1} \succeq s^t} \Pr(s^{t+1} | s^t) U_c(s^{t+1}) \tag{A.2}$$

with equality whenever the debt limit is not binding.

³⁰Since there is only one type of household we omit the subscript i to index allocations by type.

Comparing equation (A.2) with equation (4) in the main text, and abstracting from the debt limit, we notice that (4) holds state by state while (A.2) holds only in expectations. This difference is captured in the prototype model by non-zero risk sharing wedges.

To see this more clearly, let's focus on a special case with tight debt limits, $\phi \rightarrow 0$. Because agents cannot borrow and assets are in zero net supply, there is no trade in the bond market and the equilibrium allocation is equivalent to the one in financial autarky. In particular, allocations are static in that they depend only on the current realization of the exogenous state, $s = (z, v)$, and solve

$$\begin{aligned} e(v) A(z) c(z, v)^{-\sigma} &= \chi l(z, v)^\psi \\ c(z, v) &= e(v) A(z) l(z, v). \end{aligned}$$

Solving for this system we obtain the consumption and labor allocations

$$c(z^t, v^t) = \chi^{-\frac{1}{\sigma+\psi}} [e(v_t) A(z_t)]^{\frac{1+\psi}{\sigma+\psi}} \quad (\text{A.3})$$

$$l(z^t, v^t) = \chi^{\frac{\sigma}{\sigma+\psi} - \frac{1}{\psi}} [e(v_t) A(z_t)]^{1 - \frac{\sigma(1+\psi)}{\sigma+\psi}}. \quad (\text{A.4})$$

From equations (A.3) and (A.4) and the definition of the wedges we have the following result

Claim 1. *The allocation $\{c(s^t), l(s^t)\}$ and individual wages $\{W(s^t)\}$ generated by the bond economy with tight debt limits is an equilibrium allocation for the prototype economy with wedges $\{\tau_a(s^t), \tau_l(s^t), \theta(v_t)\}$ given by*

$$\theta(v_t) = e(v_t) \quad (\text{A.5})$$

$$\tau_a(s^{t+1}) = 1 - \left(\frac{\theta(v_{t+1}) / \Theta(z_{t+1})}{\theta(v_t) / \Theta(z_t)} \right)^{-\frac{\sigma(1+\psi)}{\sigma+\psi}} \quad (\text{A.6})$$

$$\tau_l(s^t) = 0, \quad (\text{A.7})$$

where

$$\Theta(z_{t+1}) \equiv \sum_{v_{t+1}} \text{Pr}(v_{t+1} | z_{t+1}) \theta(v_{t+1})^{\frac{(1+\psi)}{\sigma+\psi}}.$$

The model predicts a positive correlation between the risk sharing wedge, $\tau_a(s^{t+1})$, and the growth of the individual productivity, $\theta(s^{t+1}) / \theta(s^t)$. In the [Huggett \(1993\)](#) model, a decrease in income in period $t + 1$ after an idiosyncratic shock translates into a decline in households' consumption, see equation (A.3). In a world with frictionless financial markets, the household would insure against this idiosyncratic shock by accumulating Arrow securities that pay in this state of the world. Thus, to align the predictions of the [Huggett \(1993\)](#)

model to the prototype model, we need a tax on saving for that particular state— a negative $\tau_a (s^{t+1})$.

Since there are no distortions to the labor supply and the labor market, the model predicts no labor wedges.

A.2 Heterogeneous risk aversion

We now consider a pure exchange economy with two types of households: high and low risk aversion. The period utility function for an agent of type $i \in \{H, L\}$ is $U(c) = \frac{c^{1-\sigma_i}-1}{1-\sigma_i}$ for $\sigma_H > \sigma_L$. Risk preferences is the only dimension of heterogeneity and there is no idiosyncratic risk. The only risk that agent face is an aggregate endowment risk. Specifically, the economy can be either in a boom (B) or a recession (R). The aggregate endowment is $A(z_B)$ in a boom and $A(z_R) < A(z_B)$ in a recession. We assume that financial markets are complete and agents are born with no wealth.

The equilibrium consumption allocation in this model can be summarized by the consumption share of the high risk aversion agent in boom and recessions, $\varphi_H(z_B)$ and $\varphi_H(z_R)$ respectively. These consumption shares satisfy the standard risk sharing condition

$$\frac{\varphi_H(z_B)^{-\sigma_H} A(z_B)^{-\sigma_H}}{\varphi_H(z_R)^{-\sigma_H} A(z_R)^{-\sigma_H}} = \frac{(1 - \varphi_H(z_B))^{-\sigma_L} A(z_B)^{-\sigma_L}}{(1 - \varphi_H(z_R))^{-\sigma_L} A(z_R)^{-\sigma_L}}, \quad (\text{A.8})$$

where we have used the relation $\varphi_L(z) = 1 - \varphi_H(z)$. From equation (A.8) we can see that $\varphi_H(z_B) < \varphi_H(z_R)$, because the high risk aversion agent buys insurance from the high risk aversion guy against aggregate shocks in equilibrium.

The model does not have an equivalent for labor supply. If we interpret the pure exchange economy as one in which there is no disutility from working and households supply zero hours in every period, we have the following result:

Claim 2. *The allocation $\{c_i(z^t), l_i(z^t)\}$ generated by the pure exchange economy with heterogeneous risk aversion is an equilibrium allocation for the prototype economy with wedges $\{\tau_{a,i}(z^t), \tau_{l,i}(z^t)\}$ given by*

$$\tau_{a,i}(z^{t+1}) = 1 - \left(\frac{\varphi_i(z^t)}{\varphi_i(z^{t+1})} \right)^\sigma \quad (\text{A.9})$$

$$\tau_{l,i}(z^t, j) = 1, \quad (\text{A.10})$$

for $i = H, L$.

Thus our accounting procedure interprets the increase in the consumption share of the

high risk aversion type transiting from a boom to a recession as a subsidy on saving, $\tau_{a,H}(z^{t+1}) = 1 - \left(\frac{\varphi_H(z_B)}{\varphi_H(z_R)}\right)^\sigma > 0$, and the reduction in the consumption share of the high risk aversion type transiting from a recession into a boom as a tax on saving, $\tau_{a,H}(z^{t+1}, H) = 1 - \left(\frac{\varphi_H(z_R)}{\varphi_H(z_B)}\right)^{-\sigma} < 0$. The labor wedge arises solely because of misspecification due to the need to match the first order condition for labor supply in the misspecified model.

A.3 Sticky wages and labor wedges

We now consider an example economy where the distribution of labor wedges is not degenerate. The economy is populated by a continuum of type of households indexed by $i \in [0, 1]$. A typical household of type i has preferences given by (1) and it supplies labor of type i . The productivity of labor of type i is assumed to be constant and equal to 1, $\theta_i(v_t) = 1$, and constant. The final good in the economy is produced utilizing different types of labor according to

$$Y(z^t) = A^*(z_t) \left[\int_0^1 n_i(z^t)^v di \right]^{1/v} = A^*(z_t) N_e(z^t),$$

where $A^*(z_t)$ is productivity, $v \in (0, 1)$, and $N_e(z^t)$ is an aggregate labor index.

The final good sector is competitive. Given a distribution of wages $\{W_i(z^t)\}_i$ where $W_i(z^t)$ is the wage for labor input of type i , the demand for such type of labor is

$$n_i(z^t) = \left(\frac{W(z^t)}{W_i(z^t)} \right)^{1/(1-v)} N_e(z^t) \quad (\text{A.11})$$

where $W(z^t)$ is the (nominal) ideal wage index, $W(z^t) = \left[\int W_i(z^t)^{v/(v-1)} di \right]^{(v-1)/v}$. The final good price is $P(z^t) = W(z^t) / A^*(z^t)$.

Households of type i are organized in a union. The union set wages in a staggered fashion. In particular, they can change their wage with probability ζ in each period, iid over time. Given the prevailing wage in the union, union members commit to supply labor according to the labor demand (A.11). We assume that financial markets are complete in that households can insure against aggregate shock and idiosyncratic shocks (the ability to update the wage). Without loss of generality, we can think of the union choosing assets for the households and not only setting their wage. The problem for a union that has the opportunity to set its wage

after history z^t is

$$V^*(\omega, z^t) = \max_{c, l_i, W_i, a^r, a} U(c, l_j) + \beta \sum_{z^{t+1}} \Pr(z^{t+1}|z^t) \left[\zeta V^*(a^r(z^{t+1}), z^{t+1}) + (1 - \zeta) V(a(z^{t+1}), W_i, z^{t+1}) \right]$$

subject to the budget constraint,

$$P(z^t) c + \sum_{s^{t+1}} Q(z^{t+1}) \left[\zeta a^r(z^{t+1}) + (1 - \zeta) a(z^{t+1}) \right] \leq W_i l_i + \omega, \quad (\text{A.12})$$

and the demand $l_j = n_j(z^t)$ in (A.11) where ω is the financial wealth of the representative member of the union and $V(\omega, W_i, z^t)$ is the value if it cannot reset its price next period given by

$$V(\omega, w_j, z^t) = \max_{c, a^r, a} U(c, l_j) + \beta \sum_{z^{t+1}} \Pr(z^{t+1}|z^t) \left[\zeta V^*(a^r(z^{t+1}), z^{t+1}) + (1 - \zeta) V(a(z^{t+1}), w_j, z^{t+1}) \right]$$

subject to the budget constraint (A.12) and the demand $l_j = n_j(z^t)$ in (A.11).

The solution to this problem gives that the reset wage after an aggregate history z^t , $W^*(z^t)$, is common across types that get to choose a new wage and is given by

$$W^*(z^t) = \frac{1}{v} \frac{\sum_{r=0}^{\infty} \sum_{z^{t+r}} \beta^r (1 - \zeta)^r \Pr(z^{t+r}|z^t) \left[-U_{li}(z^{t+r}) \frac{N_e(z^{t+r})}{W(z^{t+r})^{-1/(1-v)}} \right]}{\sum_{r=0}^{\infty} \sum_{z^{t+r}} \beta^r (1 - \zeta)^r \Pr(z^{t+r}|z^t) \left[\frac{U_{ci}(z^{t+r})}{P(z^{t+r})} \frac{N_e(z^{t+r})}{W(z^{t+r})^{-1/(1-v)}} \right]}$$

and consumption is completely insured against the risk of not being able to adjust the wage so that for all i, j and z^t, z^r we have

$$\frac{U_{cj}(z^t)}{U_{cj}(z^r)} = \frac{U_{ci}(z^t)}{U_{ci}(z^r)}$$

To make things simple, we assume that the initial holdings of arrow securities are such that $c_i(z^t) = C(z^t)$ for all i .

In this economy, the idiosyncratic history v^t consists of the realization of the Poisson process that determines the ability of the union to reset its wage. In particular, since we are assuming that markets are complete (and thus households can insure against the chance of their union not being able to update the price), to determine the allocation of an agent in

period t and history z^t , the only relevant information we need is the last time their union reset the wage. Alternatively, if we know the the distribution of wage in z^t , $\{W_j(z^t)\}$,³¹ then we have the following result:

Claim 3. *The allocation $\{c_i(z^t, v^t), l_i(z^t, v^t), W_i(z^t, v^t)\}$ generated by the sticky wages economy is an equilibrium allocation for the prototype economy with wedges $\{\tau_{a,i}(z^t, v^t), \tau_{l,i}(z^t, v^t), \theta_i(z^t, v^t)\}$ given by*

$$\begin{aligned}\tau_{s,i}(s^{t+1}) &= 0 \\ \tau_{l,i}(s^t) &= 1 - \chi \frac{N_e(z^t)^\psi}{W(z^t)Y(z^t)^{-\sigma}} \left(\frac{W_i(s^t)}{W(z^t)} \right)^{-[\psi+(1-v)]/(1-v)} \\ \theta_i(s^t) &= \frac{W_i(s^t)}{\bar{W}(z^t)}\end{aligned}$$

where $\bar{W}(z^t)$ is the average wage.³²

To derive the individual labor wedge, we manipulate the first order condition for labor supply in the prototype model as follows:

$$\begin{aligned}1 - \tau_{l,i} &= \chi \frac{l_i(s^t)^\psi}{c_i(z^t)^{-\sigma} W_i(s^t)} \\ &= \chi \frac{\left[(W_i(s^t)/W(z^t))^{1/(1-v)} N_e(z^t) \right]^\psi}{Y(z^t)^{-\sigma} W_i(s^t)} \\ &= \chi \frac{\left[(W_i(s^t)/W(z^t))^{1/(1-v)} N_e(z^t) \right]^\psi}{W(z^t) Y(z^t)^{-\sigma} W_i(s^t)/W(z^t)} \\ &= \chi \frac{N_e(z^t)^\psi}{W(z^t) Y(z^t)^{-\sigma}} \left(\frac{W_i(s^t)}{W(z^t)} \right)^{-[\psi+(1-v)]/(1-v)}\end{aligned}$$

where the first line is the first order condition in the prototype model, in the second line we use the labor demand (A.11) to express labor supply $l_i(s^t)$ and that $c_i(s^t) = Y(z^t)$, in the third line we divide and multiply the denominator by $W(z^t)$, the last line is simple algebra.

Note that the individual labor wedge has an aggregate and an idiosyncratic component. The aggregate component, $\chi \frac{N_e(z^t)^\psi}{W(z^t)Y(z^t)^{-\sigma}}$, is driven by the fact that unions set their wage as a markup over the marginal rate of substitution, $1/v > 1$. The idiosyncratic component,

³¹For t sufficiently large, the distribution of wages in a given period with history z^t has support $\{W^*(z^{t-k})\}_{k \geq 0, z^{t-k} \leq z^t}$ with probabilities $\{\xi(1-\xi)^k\}_{k \geq 0}$.

³²Note that the average wage may differ from the ideal wage index $W(z^t)$.

$(W_j(z^t) / W(z^t))^{-[\psi+(1-\nu)]/(1-\nu)}$, is driven by wage dispersion.³³ Workers with a higher wage will supply less labor than workers with lower wages that are identical in all other respects: same productivity and same consumption (given complete asset markets and separable utility).

B Data

In this appendix we give more details about sample selection and variables definition. We also present some summary statistics of the raw data and show that our sample both aggregates reasonably and is consistent with recent work on consumption inequality.

B.1 Definition of variables and sample selection in the CEX

Consumption expenditures. Our measure of consumption expenditure is close to the NIPA definition of nondurable and services expenditures. It is constructed by aggregating up the following expenditure sub-categories: food, tobacco, domestic services, adult and child care, utilities, transportation, pet expenses, apparel, education, work-related and training, health-care, insurance, furniture rental and small textiles, housing related expenditures excluding rent.

Total hours worked. We compute total hours worked for the head of household by multiplying the number of weeks worked full or part time over the last year (*INCWEEK1*) multiplied by the numbers of hours usually worked per week (*INC_HRS1*). We obtain the same indicator for the spouse and add the two.

Labor income. We compute labor income as the sum of total household (CU) income from earnings before taxes (*FSALARYX*), plus the total income received from farm (*FFRMINCX*) and nonfarm business (*FNONFRMX*).

Disposable income. Prior to 2003, disposable income is the variable *FINCATAX*. From 2004, we use the variable *FINCATXM*.

Liquid assets. It includes the total amount the households held in savings accounts in financial institutions (*SAVACCTX*), checking and brokerage accounts (*CKBKACTX*).

³³Note $\psi + (1 - \nu) > 0$ since $\psi \geq 0$ and $\nu \in (0, 1)$.

Assets. It includes liquid assets plus the value of owned automobiles (*NETPURX*), residential housing (*PROPVALX*) and money owned to the household by individuals outside of the household (*MONYOWDX*).

Net worth. We compute net worth as the value of assets minus the current value of the households home mortgage (*QBLNCM3X*) plus auto debt (*PRINCIPX*).

Taking 2006 as the year of reference, the number of households reporting consumption and income in interview 5 is 7091. We next keep households whose head is in the age bracket 22-60, leaving us with 5058 households that reported income and consumption in 2006. Within this group, we keep households that are considered “full income responders” (4331), and drop any household that observed a change in family size between the first and the last interview (3826). We then drop observations on consumption, labor income, total hours, wage per hour, disposable income, liquid assets and net worth that fall below the 1st percentile or above the 99th percentile of the distribution of these variables, leaving us with 3424 households for 2006. As explained in the main text, we keep only households whose fifth interview falls in the third and fourth quarter of each year in order to minimize time aggregation issues. This leaves us with 995 households in 2006.

B.2 Definition of variables and sample selection in the PSID

Consumption expenditures. Our measure of consumption expenditure is close to the NIPA definition of nondurable and services expenditures. It is constructed by aggregating up the following expenditure sub-categories: food, adult and child care, utilities, transportation, education, work-related and training, healthcare, insurance, and housing related expenditures including rent.

Total hours worked. We compute total hours worked for the household as the sum of the total annual work hours on all jobs including overtime by the head (*ER40876*) and the spouse (*ER40887*).

Labor income. We compute labor income as the sum of total household wages (*ER40921* and *ER40933*) plus the labor earnings from farm (*ER36854*) and nonfarm businesses (*ER40900* and *ER403930*).

Disposable income We compute disposable income in two steps. First, we construct an estimate of the households income tax liability using the NBER’s *taxsim9* program. Second, we construct disposable income as the difference between total household labor income minus the household’s total tax liability net of tax credits.

Liquid assets. It includes the total amount the households held in CDs, savings and checking accounts in financial institutions (*ER37595*).

Assets. It includes liquid assets plus the value of owned automobiles (*ER37557*), residential housing (*ER37553*).

Net worth. We compute net worth as the value of assets minus the current value of the households home mortgages (*ER36042* and *ER36054*), and the value of all other debt (*ER37621*).

From the initial sample we drop households that report negative values for consumption expenditures, labor income, disposable income and hours worked at any interview. We take 2006 as our reference year and start with 8289 households before sample selection. We keep households whose head is in the age bracket 22-60, leaving us with 6669 households that reported income and consumption in 2006. Within this group, we keep households in both the main and immigrant samples who have non-zero weights, at least three observations per household and who head ages without measurement error. This leaves us with 5485 households reporting income and consumption in 2006. We then drop observations on consumption, labor income, total hours, wage per hour, disposable income, liquid assets and net worth that fall below the 1st percentile or above the 99th percentile of the distribution of these variables, leaving us with 5099 households for 2006.

B.3 Summary statistics

Table [A-1](#) reports selected households' characteristics for 2006. In the CEX, the average age for the head of household was 43.30 years, and roughly 30% of the households' head held a college degree. The average size of the household was 2.68. On average, households spent roughly 10000 dollars per person in non-durables and services, and the average income per person was 24000 dollars. Households worked 1240 hours per year per person on average, earning an average wage of 19.80 dollars per hour. The mean net worth for the household was 100000 dollars, with 7700 dollars in liquid assets. The figures for the PSID are very close to those in the CEX, with the exception of liquid assets, which is a bit higher in the PSID. As a comparison with previous papers, the average characteristics of the household in our sample are very close to those reported in [Heathcote and Perri \(2018\)](#), see Table 1 in their paper.

Table A-1: Average characteristics of households in 2006

	CEX	PSID
Age of head	43.30	43.48
Household size	2.68	2.63
Head with college (%)	29.07	34.92
Consumption expenditures per person	9697.88	8256.90
Labor income per person	23927.08	23994.33
Disposable income per person	24194.78	23344.35
Hours worked per person	1239.14	1272.23
Wage per hour	19.80	18.84
Household's net worth	96696.08	109934.80
Liquid assets	7703.08	12883.90

Notes: The sample size is 995 households for the CEX and 5099 for the PSID. All statistics are computed using sample weights. All monetary variables are expressed in 2000 U.S. dollars.

B.4 Aggregation

In this section we examine whether the dynamics of aggregate consumption, income, and total hours per capita in our cross sectional data capture the broad contours of national income and product accounts (NIPA) aggregates. The results are shown in figure A-1. Each graph is normalized to 1 in 2004.³⁴

The top left panel of figure A-1 shows the dynamics of average per capita expenditures in the PSID, the CEX and the equivalent measure in the NIPA. The top right panel shows average per capita disposable income in the PSID, the CEX and NIPA. The bottom panel shows average total hours worker per capita in the two surveys and compare it to the aggregate counterpart obtained from the BLS. While the fit is not perfect, it is clear that both datasets capture the broad contour of each aggregate series during the Great Recession.³⁵

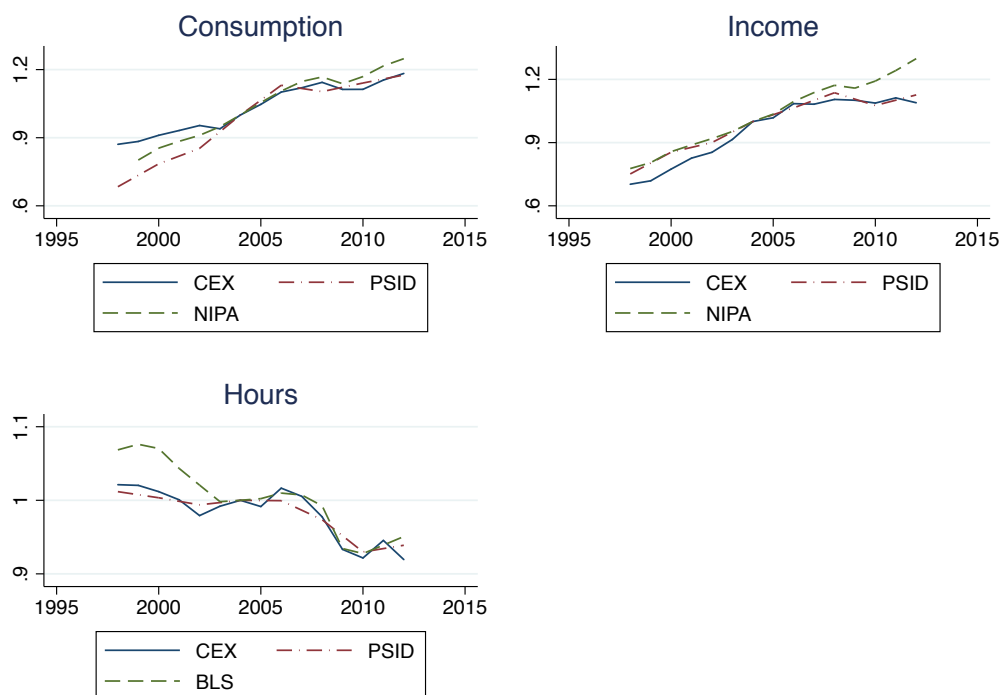
B.5 Trends in inequality

A large literature has documented that consumption inequality has increased in both the CEX and the PSID (Blundell, Pistaferri, and Preston (2008); Krueger and Perri (2006); Aguiar and Bils (2015); Attanasio and Pistaferri (2016)). Consistent with this literature, we find that the variance on log consumption has increased in both of our data sets. These results are displayed in figure A-2. There is clear visual evidence that consumption volatility has

³⁴These figures are constructed before any sample selection.

³⁵Our figure A-1 is very similar to the relevant panels in figure 13 of Heathcote and Perri (2018) giving us further confidence.

Figure A-1: Comparing aggregates across micro data sets

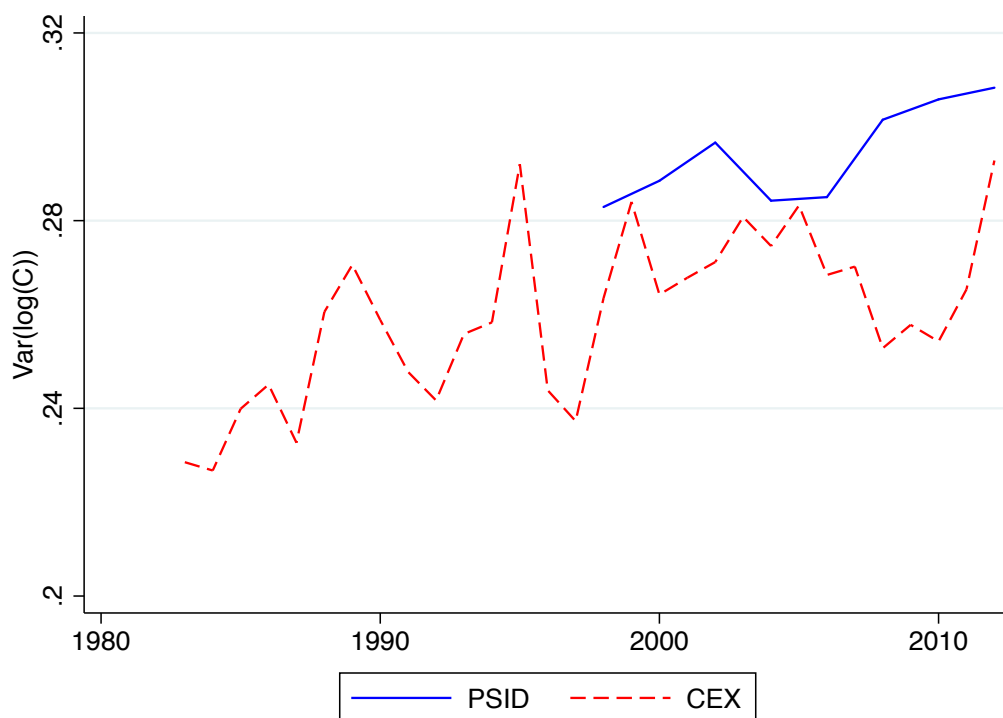


Top panel compares the behavior of nominal consumption and disposable income per capita in the PSID, CEX and NIPA. The bottom panel shows the behavior of total hours worker per capita in the three data sets. Each graph is normalized to 1 in 2004 because the levels vary somewhat across data sets.

increased. Moreover, the levels of consumption inequality that we find are very similar to previous work in the literature. In particular, we find that the variance of log consumption has increased from 0.23 to 0.28 in the CEX over the period 1985 to 2005, which almost the exact same increase in both levels and changes that [Heathcote, Perri, and Violante \(2010\)](#) find over the same time period (see figure 1 in the recent survey by [Attanasio and Pistaferri \(2016\)](#) for more details). Likewise, our results for the PSID are quantitatively similar to the results reported in [Attanasio and Pistaferri \(2016\)](#). Overall, this suggests that our sample selection procedure is reasonable.

C Numerical solution

Figure A-2: Consumption Inequality (CEX and PSID)



Dashed red line shows the time-trend in the CEX starting in 1983; the solid blue line shows the time-trend in the PSID starting in 1998.