

Idiosyncratic Income Risk and Aggregate Fluctuations*

Davide Debortoli
UPF, CREi and BSE

Jordi Galí
CREi, UPF and BSE

December 2022

Abstract

We study the role of idiosyncratic income risk for aggregate fluctuations within a simple heterogeneous households framework. We show that the presence of idiosyncratic income shocks affects the economy's response to an aggregate shock even in the absence of binding borrowing constraints and/or cyclical income risk. Their impact can be captured by a consumption-weighted average of the changes in consumption risk generated by an aggregate shock. We apply this framework to two example economies—an endowment economy and a New Keynesian economy—and show that under plausible calibrations the impact of idiosyncratic income risk on aggregate fluctuations is quantitatively small, since most of the changes in consumption risk are concentrated among poorer (low consumption) households.

JEL Classification Numbers: E21, E32, E50

Keywords: heterogeneous households, economic fluctuations, aggregate shocks, idiosyncratic shocks, monetary policy, HANK models.

*We thank Sushant Acharya, Susanto Basu, Benjamin Moll and participants at the Copenhagen Macro-Days Workshop, the NBER Summer Institute, and seminars at Boston College, Maryland, London Business School, Bank of Canada, Ecole Polytechnique/CREST, Univ. of Lausanne, Bocconi, Oxford, ECB, Norges Bank, PSE and CREI Faculty lunch for their comments and suggestions, and Kadir Ozen and Danila Smirnov for excellent research assistance. We acknowledge the financial support from the European Research Council under the European Union's Horizon 2020 research and innovation program (Galí, Advanced Grant 882332-HEMPEF), from La Caixa Research Grant on Socioeconomic Wellbeing (Debortoli), from the Spanish Ministry of Economy and Competitiveness through the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2019-000915-S) and through grants PID2020-116268GB-I00 and from the Generalitat de Catalunya, through CERCA and SGR Programme (2017-SGR-1393). The contribution by Davide Debortoli has been prepared under the Wim Duisenberg Fellowship sponsored by the ECB.

1 Introduction

Most efforts at modeling and understanding aggregate fluctuations over the past decades have relied on frameworks that assume an infinitely-lived representative household. While that assumption is obviously unrealistic, its widespread adoption reflects the view that both the finite lifetimes and the pervasive heterogeneity observed in the real world (in education, wealth, income, etc.) are not important factors behind aggregate fluctuations, and thus can be safely ignored when seeking to understand the nature and causes of that phenomenon, as well as its implications for policy.¹

But the dominance of the representative household paradigm in macroeconomics has been challenged in recent years by a number of researchers who have argued that such an assumption, while convenient on tractability grounds, is less innocuous than one may think, even when the focus is to understand aggregate fluctuations and macroeconomic policies. The emergence of Heterogeneous Agent New Keynesian (HANK) models in recent years is a reflection of this challenge. HANK models up to date have focused on household heterogeneity and its implications for aggregate consumption. They commonly assume the presence of idiosyncratic shocks to households' income, together with the existence of incomplete markets and borrowing constraints. Those features are combined with the kind of nominal rigidities and monetary non-neutralities that are the hallmark of New Keynesian (henceforth, NK) models. An important focus of that recent literature has been the role of heterogeneity in the transmission of monetary policy shocks.²

Rather than developing a richer HANK model that accounts for a broader set of facts or innovates over existing ones in some dimension, in the present paper we take a step back and use a basic model of individual and aggregate consumption with a specific goal in mind: to shed some light on the mechanisms through which the presence of exogenous idiosyncratic income risk may influence aggregate fluctuations in the absence of complete markets. Our framework features exogenous idiosyncratic income risk shocks as the only source of heterogeneity, in an environment where (i) the only asset available is a riskless one-period bond, (ii) borrowing constraints are not binding in equilibrium and (iii) idiosyncratic income risk is time-invariant. While the previous assumption are admittedly very strong and unrealistic, we study such an environment in order to isolate as

¹Early attempts to introduce heterogeneity into real business cycle models tended to support that view. See for instance Heathcote, Storesletten and Violante (2009), Guvenen (2011) and Krueger, Mitman and Perri (2016) for useful surveys of this earlier literature.

²See, among others, Auclert (2019), Kaplan, Moll and Violante (2018), Werning (2015), Acharya and Dogra (2020), Ravn and Sterk (2021) and McKay et al. (2016).

much as possible the *intrinsic* role of idiosyncratic income risk in shaping aggregate fluctuations, thus deliberately abstracting from other features of heterogeneous agent (HA) models that the literature has stressed as playing an important role in those models, namely: (i) differential liquidity across types of assets, (ii) hand-to-mouth behavior by a fraction of households (possibly due to binding borrowing constraints) and (iii) counter-cyclical income risk.

At the core of our analysis is an (approximate) Euler equation for (log) aggregate consumption which we derive by aggregating the corresponding Euler equations of individual households. That aggregation is possible given our assumption of non-binding borrowing constraints.

We show that the Euler equation for aggregate consumption in the HA economy differs from its representative agent (RA) counterpart by including a term that captures a precautionary savings motive resulting from individual consumption risk. This additional term, which we refer to as the *risk shifter*, takes the form of a *consumption-weighted average of individual consumption risk*, with the latter being measured by the conditional variance of one-period-ahead (log) consumption. In a representative agent model, where aggregate shocks are the only source of uncertainty, the risk shifter is of second order relative to variations in aggregate consumption and for that reason it is usually ignored. In contrast, in a HA economy, due to the presence of (potentially large) idiosyncratic income shocks in the background, variations in the risk shifter in response to aggregate shocks may be of the same order of magnitude as the latter, and hence potentially play a more important role.

A central result of our analysis is that the role of idiosyncratic income risk on aggregate fluctuations depends on how aggregate shocks affect the *distribution* of consumption risk across households. In fact, aggregate shocks alter households' ability to insure against their idiosyncratic income shocks, thus leading to fluctuations in consumption risk, even if the underlying income risk remains constant, as we assume. As an example, consider an aggregate shock, such as an increase in interest rates, that reduces the ability of households to insure against their individual income shocks, and for this reason leads to a widespread and persistent increase in *average* consumption risk. That effect, by itself, would tend to reduce aggregate consumption, due to a precautionary savings motive. But the change in average consumption risk is not enough to predict the impact of the shock on aggregate consumption: how it is distributed across households matters. Thus, to the extent that the increase in consumption risk is concentrated among poorer (i.e. low consumption) households, the impact on aggregate consumption will be smaller. This is what we refer to as the *distribution* channel.

After deriving and discussing the properties of the Euler equation for aggregate consumption we embed that equation into two fully fledged model economies. The first economy is an endowment economy where households are subject to endowment shocks, both idiosyncratic and aggregate. In that context, we study the mechanisms through which the presence of idiosyncratic income risk influences the response of the (real) interest rate to aggregate endowment shocks. The second economy is described by a baseline New Keynesian model with households subject to idiosyncratic productivity shocks. Our interest lies in studying the impact of those idiosyncratic shocks in shaping the response of aggregate output to aggregate shocks, such as monetary policy and technology shocks. The simplicity of both models and the fact that the presence of exogenous idiosyncratic income risk is the only departure from their RA counterparts allows us to isolate better the intrinsic role of that risk in shaping aggregate fluctuations, independently from its possible interaction with other features of the economy.

From a quantitative viewpoint, we find that idiosyncratic income risk has very small net effect on aggregate fluctuations in the two calibrated model economies that we analyze, mainly because of the neutralizing impact of the *distribution* channel mentioned above.

The rest of the paper is structured as follows. Section 2 reviews the related literature. Section 3 presents the model and the corresponding Euler Equation for aggregate consumption. Section 4 and 5 embed the previous framework into an endowment economy and a New Keynesian economy, respectively, highlighting the role of the distribution of consumption risk, both from a qualitative and a quantitative perspective. Section 6 concludes.

2 Related Literature

This paper belongs to a growing literature that studies the role of household heterogeneity in aggregate economic fluctuations. In that literature, the differences in the behavior of aggregate variables relative to a representative agent economy are a consequence of several features embedded in the proposed models which are absent from their RA counterparts. However, understanding which is the exact role played by each of these factors remains an open question. Our objective in the present paper is to shed light on the role of exogenous idiosyncratic income risk and the channels through which it affects aggregate fluctuations.³

³An exercise in a similar spirit, but focusing of firms' heterogeneity and the role of collateral constraints can be found in Cao and Nie (2017).

Several studies in the literature have developed *tractable* frameworks to isolate the channels through which heterogeneity operates. Following the original formulation of Campbell and Mankiw (1989), some studies in that literature (see e.g., Galí et. al. (2007), Bilbiie (2008), Debortoli and Galí (2018) and Broer et. al. (2020)) have focused on the role of binding constraints, by analyzing models with two types of agents (unconstrained and hand-to-mouth), but abstracting from the presence of idiosyncratic income risk within each type. Here we do the opposite, and focus instead on the role of idiosyncratic income risk, showing how the latter may give rise to amplification/dampening of aggregate shocks, even in the absence of binding borrowing constraints.

A number of authors (e.g. Werning (2015), McKay et al. (2016), Bilbiie (2021), Ravn and Sterk (2021)) have studied economies with idiosyncratic income risk, but under assumptions regarding the nature of borrowing constraints that imply a degenerate wealth distribution in equilibrium.⁴ As a result, individual consumption and income are equated in equilibrium, with the former inheriting the risk properties of the latter. That literature emphasizes the role played by the cyclical nature of *income risk* and *liquidity* for the transmission of aggregate shocks. Our work instead emphasizes the role of variations in *consumption risk*, above and beyond the presence of a time-varying income risk, and uncovers a novel channel related to how changes in consumption risk is distributed across households. Similarly to us, Werning (2015, Section 4) derives an Euler equation for aggregate consumption in a heterogeneous agent model with positive liquidity, where a "wedge" summarizes all the differences relatively to a representative agent framework. Our paper gives an economic interpretation to that wedge and relates it to the distribution of consumption risk across households.

In related work, Acharya and Dogra (2020) consider an heterogeneous household economy with CARA preferences and, like our paper, no binding borrowing constraints. Yet, in that economy, due to the assumption on preferences, all households face the same consumption risk (the marginal propensity to consume out of their cash-on-hand is identical across households), and heterogeneity mainly operates as a result of the cyclical nature of *income risk*. We instead consider a framework with more standard CRRA preferences, associated with a non-trivial relationship between individual consumption, income and wealth. In our setting, the cyclical behavior of *consumption risk* in response to aggregate shocks plays a crucial role for the transmission of the latter, regardless of whether the volatility of the underlying idiosyncratic risk is constant or not.

⁴For instance, economies with zero-liquidity, or with no (or limited) wealth inequality among unconstrained households. See also Challe and Ragot (2011) and Challe, Matheron, Ragot, and Rubio-Ramirez (2017) for tractable models where the wealth distribution has finite support.

Our paper is also related to several studies in the literature proposing some “sufficient statistics” to summarize the aggregate implications of household heterogeneity (see e.g. Auclert (2019), Auclert, Rognlie and Straub (2018) and Luetticke (2021) and the references therein). Those studies have emphasized the role of the cross-sectional distribution of variables like the marginal propensity to consume, income, portfolios, etc. Our contribution is to show that the role of idiosyncratic risk can be summarized by the cross-sectional distribution of changes in consumption risk.

Our work is related to several quantitative studies which analyzed real business cycle (RBC) models augmented with idiosyncratic shocks to households’ income and a borrowing constraint, following the seminal work of Krusell and Smith (1998). A main finding in that literature is that this class of models feature an “approximate aggregation” property, which means that the dynamics of aggregate variables can be accurately described using only the mean of the wealth distribution, but ignoring higher-order moments. The “approximate aggregation” result indicates that there exists a parsimonious representation of the equilibrium dynamics of a heterogeneous agent economy. That equilibrium, however, does not necessarily coincide with the corresponding economy without idiosyncratic shocks, thus still allowing the latter to play a significant role in influencing aggregate fluctuations.⁵ Krusell and Smith (1998) conjecture that the “approximate aggregation” result obtains when variations in the *marginal propensity to consume* in response to aggregate shocks are concentrated among low wealth households. We shed light on that conjecture, and show analytically that in order for idiosyncratic risk to have a small impact on aggregate variables, it must be the case that variations in *consumption risk* are concentrated among low consumption households.⁶

A recent paper by Berger et al. (2022) bears a close relation to ours in that the authors also derive an aggregate Euler equation incorporating a wedge that captures the departures from perfect risk sharing (and, hence, from equivalence with an RA model), while being agnostic about the precise nature of those departures.⁷ The focus of their paper lies in the use of micro data to measure the evolution of that wedge over time, and to estimate the contribution of its variations to aggregate output volatility (which they find to be small). In order to do so, they estimate a stochastic process for the wedge, which they then feed as an additional impulse to the equations describing the equilib-

⁵For instance, Krusell and Smith (1998) consider an example model with heterogeneity in discount factors that generates significant differences from the predictions of the corresponding RA model.

⁶As explained in more detail in Section 3.1, there is no simple mapping between the distribution of consumption risk and the distribution of the marginal propensity to consume.

⁷Thus, the wedge in Berger et al. (2022) could be due to hand-to-mouth behavior by a fraction of households, even in the absence of idiosyncratic income risk. See also Section 4 in Werning (2015).

rium evolution of aggregate variables. In the present paper, the wedge in the aggregate Euler equation arises from a specific departure from the RA model, namely, the presence of uninsurable income risk, and is given a clear interpretation (related to consumption risk). On the other hand, we solve for the equilibrium using as input of our model a calibrated process for individual income (rather than an estimated reduced form process for the measured aggregate wedge).

3 An Euler Equation for Aggregate Consumption

Throughout we assume a continuum of households indexed by $j \in [0, 1]$. Preferences are common to all households and given by $\mathbb{E}_0\{\sum_{t=0}^{\infty} \beta^t U(C_t(j))\}$ where $C_t(j)$ denotes household j 's consumption in period t , $\beta \equiv \exp\{-\rho\}$ is the discount factor and $U(C) = (1 - \sigma)^{-1}(C^{1-\sigma} - 1)$, with $\sigma \geq 0$. Households can borrow and lend at a (gross) riskless real rate $R_t \equiv \exp\{r_t\}$, subject to the natural debt limit. The Euler equation describing optimal consumption for an individual household is given by

$$1 = \beta R_t \mathbb{E}_t\{(C_{t+1}(j)/C_t(j))^{-\sigma}\} \quad (1)$$

which is assumed to hold for $t = 0, 1, 2$, and for all $j \in [0, 1]$. Our objective in this section is to derive an approximate Euler equation for (log) aggregate consumption. In our approximation we include all the terms of a Taylor expansion whose variations are of the same order –which we henceforth denote as $\mathcal{O}(|\varepsilon|)$ – as variations in *aggregate* consumption growth or the real interest rate.

As derived in Appendix A.1, up to order $\mathcal{O}(|\varepsilon|)$, equation (1) can be written as follows:

$$\mathbb{E}_t\left\{\frac{\Delta C_{t+1}(j)}{C_t(j)}\right\} \simeq \frac{1}{\sigma} \left(1 - \frac{1}{\beta R_t}\right) + \frac{\sigma + 1}{2} v_t(j) \quad (2)$$

where $v_t(j) \equiv \mathbb{E}_t\{(\Delta C_{t+1}(j)/C_t(j))^2\} \simeq \text{var}_t\{c_{t+1}(j)\}$, with $c_t(j) \equiv \log C_t(j)$.⁸ We can thus interpret $v_t(j)$ as a measure of risk regarding household j 's one period ahead (log) consumption, whose effect on expected consumption growth captured by (2) reflects the so called precautionary savings motive resulting from the convexity of marginal utility.⁹

⁸The use of approximations is only made to facilitate the economic interpretation of the resulting terms in the Euler equation. Appendix A.2 contains an analogous representation that does not rely on any approximation, and which is actually used in our quantitative exercises.

⁹Note that under our assumed utility function the coefficient of "relative prudence" –a measure of that convexity– is constant and given by $-(U''' / U'')C = \sigma + 1$. Appendix A.3 contains an analogous derivation for a general utility function, and also a special case for Constant-Absolute-Risk-Aversion (CARA) utility.

Due to the presence of (potentially large) idiosyncratic income shocks in the background, we allow variations in $v_t(j)$ to be of order $\mathcal{O}(|\varepsilon|)$. This is in contrast with the representative household case, for which $v_t \equiv \mathbb{E}_t\{(\Delta C_{t+1}/C_t)^2\} \sim \mathcal{O}(|\varepsilon|^2)$, which justifies the absence of v_t from the familiar first-order approximations of the consumption Euler equation found in the literature. Similarly, the equations below should be understood as holding up to an error term of order $\mathcal{O}(|\varepsilon|^2)$.

Next we derive the main result of the present section. Let $C_t \equiv \int C_t(j) dj$ denote aggregate consumption. Aggregating eq. (2) across households, we get that expected aggregate consumption growth is given by

$$\begin{aligned} \mathbb{E}_t \left\{ \frac{\Delta C_{t+1}}{C_t} \right\} &= \mathbb{E}_t \left\{ \int \frac{\Delta C_{t+1}(j)}{C_t} dj \right\} \\ &= \int \frac{C_t(j)}{C_t} \mathbb{E}_t \left\{ \frac{\Delta C_{t+1}(j)}{C_t(j)} \right\} dj \\ &= \frac{1}{\sigma} \left(1 - \frac{1}{\beta R_t} \right) + \frac{\sigma + 1}{2} v_t \end{aligned} \quad (3)$$

where

$$v_t \equiv \int \frac{C_t(j)}{C_t} v_t(j) dj \quad (4)$$

is a *consumption-weighted average of individual consumption risk*. The response of v_t to aggregate shocks will be shown to be key in understanding the role of idiosyncratic income risk in aggregate fluctuations. Henceforth we refer to v_t as the *risk shifter*.

Evaluating (3) at a stochastic steady state with constant aggregate consumption, we obtain the relation

$$0 = \frac{1}{\sigma} \left(1 - \frac{1}{\beta R} \right) + \frac{\sigma + 1}{2} v \quad (5)$$

where R and v denote the values of R_t and v_t at that steady state. Note that (5) captures an inverse equilibrium relation between risk and the real interest rate, working through precautionary savings, with $\beta R \leq 1$ and $\lim_{v \rightarrow 0} \beta R = 1$.

A first-order Taylor expansion of (3) around the stochastic steady state yields a linear Euler equation for (log) aggregate consumption $c_t \equiv \log C_t$:

$$c_t = \mathbb{E}_t \{ c_{t+1} \} - \frac{1}{\sigma} \hat{r}_t - \frac{\sigma + 1}{2} \hat{v}_t \quad (6)$$

where $\hat{r}_t \equiv \frac{1}{\beta R} \left(\frac{R_t - R}{R} \right)$ and $\hat{v}_t \equiv v_t - v$. Thus, we see how the presence of idiosyncratic income shocks calls for an additional term in an otherwise familiar log-linear Euler equation for aggregate consumption. The additional term, $-\frac{\sigma+1}{2} \hat{v}_t$, will generally vary

endogenously, thus amplifying or dampening the response of consumption to aggregate shocks, conditional on a given path for the real interest rate.¹⁰

In order to further understand how the risk shifter evolves over time we can decompose v_t as defined in equation (4) as follows:

$$v_t = \bar{v}_t + cov_j \left\{ \frac{C_t(j)}{C_t}, v_t(j) \right\} \quad (7)$$

where $\bar{v}_t \equiv \int v_t(j) dj$ is an (unweighted) average of individual consumption risk while the second term captures the cross-sectional covariance between consumption risk and relative consumption.

As shown formally in Appendix A.4, the dynamic response of the risk shifter to a generic aggregate shock ε_t , denoted by $\frac{dv_{t+k}}{d\varepsilon_t}$ for $k = 0, 1, 2, \dots$, can be written as

$$\frac{dv_{t+k}}{d\varepsilon_t} \simeq \frac{d\bar{v}_{t+k}}{d\varepsilon_t} + cov_j \left\{ c_{t+k}(j), \frac{dv_{t+k}(j)}{d\varepsilon_t} \right\} \quad (8)$$

The two terms on the right-hand side of (8) respectively capture the *average* and *distribution* channels of the effect of an aggregate shock on the risk shifter v_t .

A number of implications follow from the previous analysis. First, note that the presence of idiosyncratic income risk will have an impact on aggregate consumption fluctuations *only if* aggregate shocks have an effect on *individual consumption risk*, i.e. only if $\frac{dv_t(j)}{d\varepsilon_t}$ for a positive mass of households. Otherwise, both terms on the right-hand side of (8) would be equal to zero and the risk shifter would be unaffected by those shocks.¹¹

Secondly, the size of the response of the risk shifter depends crucially on the cross-sectional covariance between the response of individual consumption risk and the level of individual consumption, i.e. the second term on the right-hand side of (8). Thus, for any given increase in (unweighted) consumption risk in response to an aggregate shock, the change in the risk shifter (and hence the impact on aggregate consumption) will be larger the higher is the cross-sectional covariance between the change in individual consumption risk and the level of individual consumption. The intuition for the previous result is straightforward: a given change in consumption risk $\frac{\partial v_{t+k}(j)}{\partial \varepsilon_t}$ has an identical *percent* impact on the consumption of all households, independently of their initial level

¹⁰Or, alternatively, it will amplify or dampen the response of the real interest rate to an aggregate shock, conditional on a given path for aggregate consumption, as in the endowment economy considered below.

¹¹Of course, an exogenous generalized change in households' consumption risk (a "risk shock") will always have an impact on aggregate consumption.

of wealth, consumption, etc.; however, any given percent change in the consumption of an individual household has a larger impact on *aggregate* consumption (both in absolute and relative terms) the larger is the household's initial level of consumption. Thus, how any given change in consumption risk is distributed across households and, in particular, how it comoves with their level of consumption is an important factor in determining the variation in the risk shifter. In the limiting case, if consumption risk changes only for a subset of households with consumption close to zero, the impact on aggregate consumption would be negligible.

In the example economies considered below the change in consumption risk in response to an aggregate shock, $\frac{dv_{t+k}(j)}{d\varepsilon_t}$, tends to be larger, in absolute value, for low consumption households. As a result the distribution channel tends to dampen the impact of any change in average consumption risk, hence limiting the influence of idiosyncratic income risk on aggregate fluctuations.

3.1 Understanding Variations in Consumption Risk

The discussion above has made clear the importance of changes in consumption risk and their distribution in shaping aggregate fluctuations in economies where households face idiosyncratic income shocks. In the present section we try to dig further in order to shed some light on the sources of those changes.

We assume the existence of a consumption function for household j , given by:

$$c_t(j) = \mathcal{C}(s_t(j), S_t) \tag{9}$$

where $s_t(j)$ is a vector of household-specific state variables and S_t is a vector of aggregate state variables. The state variables contain all the information available at time t that is relevant to determine $c_t(j)$ (including the distribution of household-specific variables). The existence and properties of a consumption function like (9) can be established under standard assumptions.

Let $\zeta_t(j)$ and ε_t be the vectors of idiosyncratic and aggregate shocks (i.e., the mutually orthogonal, serially uncorrelated innovations in the individual and aggregate *exogenous* driving variables). We can write the innovation in household j 's consumption in period t as follows:

$$\begin{aligned} \tilde{\zeta}_t(j) &\equiv c_t(j) - \mathbb{E}_{t-1} \{c_t(j)\} \\ &= f_{t-1}^j(\zeta_t(j), \varepsilon_t) \end{aligned} \tag{10}$$

where $f_{t-1}^j(\cdot)$ is a function satisfying $f_{t-1}^j(0,0) = 0$. In what follows, and in order to keep the algebra simple, we assume $\zeta_t(j)$ and ε_t are scalars.

Under our assumptions, and using (10), we can approximate individual consumption risk $v_t(j) = \mathbb{E}_t\{\zeta_{t+1}(j)^2\}$ in period t as

$$v_t(j) \simeq \psi_t(j)^2 \sigma_\zeta^2 + \varphi_t(j)^2 \sigma_\varepsilon^2$$

where $\psi_t(j) \equiv \partial f_t^j(0,0)/\partial \zeta_{t+1}(j)$, and $\varphi_t(j) \equiv \partial f_t^j(0,0)/\partial \varepsilon_{t+1}$ are the (local) elasticities of individual consumption with respect to idiosyncratic and aggregate shocks, while $\sigma_\zeta^2 \equiv \mathbb{E}\{\zeta_t(j)^2\}$ for all $j \in [0,1]$ and $\sigma_\varepsilon^2 \equiv \mathbb{E}\{\varepsilon_t^2\}$ are, respectively, the variances of those shocks. Under our assumptions, variations in individual consumption risk driven by aggregate shocks are of second order relative to aggregate variables, i.e. $\varphi_t(j)^2 \sigma_\varepsilon^2 \sim \mathcal{O}(|\varepsilon|^2)$.¹²

Thus, for our purposes we can use the approximation

$$v_t(j) \simeq \psi_t(j)^2 \sigma_\zeta^2$$

which in turn implies the following expression for v_t :

$$v_t \simeq \sigma_\zeta^2 \int \frac{C_t(j)}{C_t} \psi_t(j)^2 dj \quad (11)$$

An implication of eq. (11) is that the risk shifter is proportional to the consumption-weighted average (across households) of the square elasticities of consumption with respect to the idiosyncratic shock.

As shown in Appendix A.4, we can then approximate the dynamic response of the risk shifter as follows:

$$\frac{dv_{t+k}}{d\varepsilon_t} \simeq \sigma_\zeta^2 \int \frac{C_{t+k}(j)}{C_{t+k}} \frac{d\psi_{t+k}(j)^2}{d\varepsilon_t} dj.$$

Thus, under our assumptions, the risk shifter will change in response to an aggregate shock only to the extent that it elicits a change in individual consumption elasticities. The ultimate impact on the risk shifter (and hence, aggregate consumption) will depend on how the change in individual consumption elasticities triggered by the shock is distributed across households. If that change is largely concentrated on low consumption households, the impact on the risk shifter will be muted. This is indeed what we find in the example economies analyzed below.

¹²To see this, note that if that term was first order, then $v_t(j)$ would be of first order even in the absence of idiosyncratic shocks, which would violate our working assumption.

An important message of our analysis is that the risk shifter will generally fluctuate in response to aggregate shocks regardless of the properties of the variance of the underlying idiosyncratic risk (σ_{ξ}^2). Throughout our analysis, we have maintained the assumption that the variance of idiosyncratic income shocks (σ_{ξ}^2) is constant over time —i.e. the idiosyncratic *income* risk is acyclical. Needless to say, the cyclicity of idiosyncratic income risk is a potentially important factor behind fluctuations in aggregate consumption, and one that has been emphasized already by several authors.¹³ Our objective here has been to point to the presence of an additional endogenous channel (above and beyond cyclical income risk) through which the very presence of idiosyncratic income risk may affect aggregate fluctuations, *independently of its cyclical properties*.

To stress the distinction between the two channels it is useful to consider the economy with heterogeneous agents and CARA preferences analyzed in Acharya and Dogra (2020). In that economy the sensitivity of consumption to idiosyncratic shocks is the same across households, independently of their level of wealth and consumption —i.e. due to CARA preferences all households have the same marginal propensity to consume—and it remains invariant to aggregate shocks. As a result, the presence of idiosyncratic income risk has an impact on aggregate consumption fluctuations only to the extent that it displays some cyclicity.

It is also important to notice that, while related, the elasticity of consumption to idiosyncratic income shocks $\psi_t(j)$ is not equivalent to the marginal propensity to consume $MPC_t(j)$. The latter is usually defined as the change in consumption implied by a one-unit unexpected increase in liquid wealth (such as winning a lottery prize). The MPC and the elasticity to idiosyncratic income shocks are tightly related only under i.i.d. idiosyncratic income shocks, since in that case individual consumption only depends on the sum of current income and wealth (“cash on hand”). At the other extreme, if idiosyncratic shocks were highly persistent, the consumption response to an idiosyncratic income shock would generally differ from the MPC, although it would be similar across households. This implies that there is no simple mapping between the two concepts.

¹³See Bayer et. al. (2019) and Ravn and Sterk (2020), among others, for examples of heterogenous household economies where cyclical idiosyncratic risk plays a central role.

4 Idiosyncratic Risk and Aggregate Fluctuations in an Endowment Economy

Consider an endowment economy populated by a continuum of households, indexed by $j \in [0, 1]$, with identical preferences given by $\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t(j))$, with $U(C) \equiv \frac{C^{1-\sigma}-1}{1-\sigma}$ where $C_t(j)$ is period t consumption of the single good by household j . The household's period budget constraint is given by:

$$C_t(j) + B_t(j) \leq B_{t-1}(j)R_{t-1} + Y_t(j)$$

$$Y_t(j) = Y_t \exp\{\zeta_t(j)\}$$

for $t = 0, 1, 2, \dots$, where $B_t(j)$ represents holdings of one-period bonds, which yield a gross riskless real return R_t , and are in zero net supply. The household endowment, $Y_t(j)$, has two components (in logs): an aggregate component $y_t \equiv \log Y_t$, which is common to all households, and follows an $AR(1)$ process with autocorrelation $\rho_y \in [0, 1)$; and an idiosyncratic component $\zeta_t(j) \in [\zeta_1, \dots, \zeta_K]$, which follows a stationary K -state Markov process, independent across households and satisfying $\mathbb{E}\{\exp\{\zeta_t(j)\}\} = 1$.¹⁴ Note that by setting $\zeta_t(j) = 0$ for all $j \in [0, 1]$ and all t , together with a uniform initial condition $B_{-1}(j) = 0$ for all $j \in [0, 1]$, the previous model collapses to one with a representative household.

In equilibrium, the bonds and goods markets must clear, which implies $\int_0^1 B_t(j) dj = 0$ and $\int_0^1 C_t(j) dj = Y_t$. We can use the Euler equation for (log) aggregate consumption (6) to derive an expression for the equilibrium real interest rate:

$$\hat{r}_t = -\sigma(1 - \rho_y)y_t - \frac{\sigma + 1}{2}\hat{v}_t \quad (12)$$

The first term on the right-hand side of (12) is the equilibrium real rate in the corresponding representative agent economy, and captures the well known effect on the interest rate of the desire to smooth consumption in the face of short-run output fluctuations. The impact of idiosyncratic risk on the interest rate is captured by the second term, which moves in proportion to the risk shifter \hat{v}_t . Thus, an increase in the latter variable tends to increase the demand for precautionary savings, leading to a reduction in the equilibrium interest rate.

In summary, equation (12) implies that the impact of idiosyncratic risk on the re-

¹⁴The previous normalization together with the law of large numbers guarantees that $Y_t = \int Y_t(j) dj$, for all t .

sponse of the real interest rate to an aggregate endowment shock is determined by the response of the risk shifter. In particular, the sign and size of that response determines the extent to which the effect of the aggregate endowment shock on the interest rate is amplified or dampened. Next we turn to a quantitative assessment of these effects in a calibrated version of the above economy.

4.1 Calibration and Solution Method

The baseline calibration of our endowment economy is summarized in Table 1. Each period is assumed to correspond to a quarter. We set the coefficient of risk aversion $\sigma = 1$, which corresponds to log utility. We set the discount factor $\beta = 0.9937$, which implies a real risk-free rate of 2 percent (in annual terms) in the steady state.

We calibrate the parameters of the K -state Markov process for idiosyncratic income using the Rouwenhorst method in order to match the volatility and persistence of an AR(1) process $\zeta_t(j) = \rho_\zeta \zeta_{t-1}(j) + \zeta_t(j)$, where $\zeta_t(j) \sim N(0, \sigma_\zeta \sqrt{1 - \rho_\zeta^2})$, with $\rho_\zeta = 0.966$ and $\sigma_\zeta = 0.5$ as in Auclert et. al. (2021).¹⁵ Finally, we set the autoregressive coefficient in the AR(1) process for the (log) aggregate endowment to $\rho_y = 0.9$.

Regarding the numerical solution method, we build a grid for individual assets of 500 points, equally distanced (in logs) between a lower bound (which corresponds to the natural debt limit, as discussed below) and an upper bound set to 300 times quarterly income. We impose a borrowing constraint of the form

$$B_t(j) \geq \underline{B} \tag{13}$$

for all t . We set $\underline{B} = -Y \exp\{\zeta_1\}/r$, which constitutes the “natural debt limit,” given aggregate output and interest rate at their steady state values (Y, r) . The desire to avoid zero consumption (given that $\lim_{c \rightarrow 0} U_c = +\infty$) guarantees that $B_t(j) > \underline{B}$ for all t when aggregate output and the interest rate are at their steady state levels. Given sufficiently small fluctuations in the previous two variables, the fraction of constrained households in equilibrium can be made arbitrarily close to zero.¹⁶

For given values of the real interest rate and the aggregate endowment, we solve for the households’ policy functions using the endogenous gridpoints method described in

¹⁵As a robustness check, Appendix B considers an alternative income process which combines a transitory and persistent component, and is a discrete-time (quarterly) version of the continuous-time process in Kaplan, Moll and Violante (2018).

¹⁶In our simulations, the fraction of constrained consumer is negligible (below 0.1 percent) both in steady state, and in response to aggregate shocks.

Carroll (2006), which are then used to calculate the implied equilibrium asset distribution. We solve for the steady state iterating on the value of the discount factor β so that the stationary asset distribution implied by the households' choices satisfies the market clearing condition $\int B_t(j) dj = 0$ at an (annualized) steady state real rate of 2 percent.

For the transition dynamics, we adopt the Sequence-Space Jacobian approach described in Auclert et. al. (2021). This amounts to finding the first-order approximation of the equilibrium responses to arbitrary sequences of anticipated shocks to the aggregate endowment (i.e. under perfect foresight) over a finite horizon (set to $T = 300$ quarters). Due to certainty equivalence, the resulting dynamics are equivalent to the ones that would be obtained solving the linearized rational expectations model, e.g. as in Reiter (2009) and Ahn et. al. (2018).¹⁷ Also, by construction, the approximate responses to positive and negative aggregate shocks are fully symmetric, and proportional to the size of the shocks. Most importantly, the assumption of perfect-foresight (or certainty equivalence) with respect to aggregate shocks implies that idiosyncratic income shocks are the only source of individual (and aggregate) uncertainty.

Finally, we note that in all our numerical exercises, and in order to accurately capture the quantitative role of idiosyncratic risk, we do not rely on the approximation described in Section 1, but instead on the the exact representation contained in Appendix A.2.

4.2 Findings

We focus our discussion on the dynamic response of the real interest rate to a positive aggregate endowment shock. Figure 1 shows the responses of the real interest rate and (log) aggregate output to a 1 percent positive shock in the latter variable. The response of the real interest rate (expressed in annual terms) is plotted on the left panel for both our baseline model with heterogeneity (red line with circles) and for the corresponding representative agent model (blue line with crosses). The real rate declines persistently in both models. Finally, the same Figure displays (green dashed line) the real rate response to the same shock under the assumption that the response of the risk shifter corresponds to that of average consumption risk, i.e. $\frac{\partial v_{t+k}}{\partial \varepsilon_t} = \frac{\partial \bar{v}_{t+k}}{\partial \varepsilon_t}$, thus implicitly turning off the distribution channel by setting $cov_j \left\{ c_{t+k}(j), \frac{\partial v_{t+k}(j)}{\partial \varepsilon_t} \right\} = 0$.

The overall effect of idiosyncratic risk on the response of the real interest rate is positive, i.e. it dampens the decline in the interest rate relatively to a representative agent model, but quantitatively small (less than 5 basis points at all horizons). That positive

¹⁷See also Boppart, Krusell and Mitman (2018) for a related perfect-foresight sequence-based approach.

impact is a consequence of a decline in the risk shifter. Note, however, that there are two distinct forces operating in opposite directions. On the one hand, the increase in aggregate output leads to a reduction in average risk \bar{v}_t , which lowers the demand for savings and tends to increase the interest rate. This is captured by the green dashed line, which lies considerably higher than the response implied by the representative agent model. On the other hand, the gap between the green dashed line and the red circled line captures the distribution channel, which nearly fully offsets the effect of average consumption risk, making the overall impact on the risk shifter (and, hence, of idiosyncratic risk) very small.¹⁸

As mentioned in section 3.1, the behavior of average consumption risk is related to the distribution of the change in the (square) elasticity of consumption with respect to the idiosyncratic shock. This is illustrated in Figure 2, which shows the steady state relationship between (log) consumption and the corresponding (square) elasticity of consumption $\psi_t^2(j)$.¹⁹ As the Figure makes clear, there is a negative relationship between these two variables, since households with higher consumption have more buffer to absorb unexpected changes in income, and thus their consumption is less sensitive to idiosyncratic shocks. Thus, an increase in aggregate income, which in and of itself causes an increase in consumption for most households, leads to a decline in the average elasticity of consumption. At the same time, the figure shows that the relationship between consumption and the (square) elasticity of consumption is convex. Intuitively, the elasticity of consumption varies substantially as households get closer to their natural debt limit, but roughly constant (and small) for households with high income and wealth, who behave almost as permanent-income consumers. This explains why an increase in aggregate income generates a significant reduction in consumption risk among low consumption households, but little change in the risk of higher consumption households, thus accounting for the offsetting distribution channel on the risk shifter. Intuitively, those households whose saving behavior is significantly affected by a reduction in consumption risk due to the positive aggregate endowment shock account for a small fraction of aggregate consumption and, hence, have a limited effect on aggregate savings and the real interest rate through this channel.

Figure 3 shows the results for a calibration with higher coefficient of risk aversion

¹⁸This result is consistent with earlier findings in the asset pricing literature, see e.g. Heaton and Lucas (1996) and Marcet and Singleton (1999), showing that household heterogeneity and market incompleteness have small effects on the volatility of returns.

¹⁹More precisely, the figure displays the range of (square) elasticities $\psi_t^2(j)$ as well as the corresponding median for each value of consumption. The existence of a range is due to the fact that, a given level consumption could be associated with different combinations of the two individual state variables, namely wealth and idiosyncratic shocks, giving rise to different elasticities.

($\sigma = 3$). In this case, as it can be seen in the left panel, the overall effects of idiosyncratic risk remain relatively small, even though a bit larger than in the baseline calibration. This is mainly because under this calibration the (square) elasticity of consumption is less convex (see Figure 4), and thus the offsetting distribution channel is weaker.

5 Idiosyncratic Risk and Aggregate Fluctuations in a New Keynesian Economy

Next we analyze the role of idiosyncratic risk and aggregate fluctuations in a version of the New Keynesian model. The economy is populated by a continuum of households, indexed by $j \in [0, 1]$, with identical preferences given by $\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t(j), \mathcal{N}_t(j))$. The term $C_t(j) \equiv \left(\int_0^1 C_t(i, j)^{1-\frac{1}{\epsilon}} di \right)^{\frac{\epsilon}{\epsilon-1}}$ is a consumption aggregator. $C_t(i, j)$ denoting the quantity of good i consumed by household j . $\mathcal{N}_t(j)$ denotes work hours. We assume $U(C, \mathcal{N}) = \left(\frac{C^{1-\sigma}-1}{1-\sigma} - \frac{\mathcal{N}^{1+\varphi}}{1+\varphi} \right)$.

Optimal allocation of expenditures requires that $C_t(i, j) = (P_t(i) / P_t)^{-\epsilon} C_t(j)$, where $P_t(i)$ is the price of good i and $P_t \equiv \left(\int_0^1 P_t(i)^{1-\epsilon} di \right)^{\frac{1}{1-\epsilon}}$ is the aggregate price index. This in turn implies that total expenditures are given by $\int_0^1 P_t(i) C_t(i, j) di = P_t C_t(j)$. The household's period budget constraint can thus be written as follows:

$$C_t(j) + B_t(j) \leq B_{t-1}(j)R_t + W_t \mathcal{N}_t(j) \exp\{\zeta_t(j)\} + D_t(j)$$

where $B_t(j)$ denotes holdings of real bonds (fully indexed to inflation) yielding a riskless real return R_t , W_t is the real wage (per efficiency unit of labor), $D_t(j)$ are real dividends, and $\zeta_t(j)$ is an idiosyncratic productivity shifter which follows a stationary K-state Markov process identical to the one assumed in the previous section, satisfying $\mathbb{E}\{\exp\{\zeta_t(j)\}\} = 1$.²⁰ Firms' shares are assumed to be nontradable and to be held in equal amounts by all households. As a result, dividends are distributed uniformly to all households, i.e. $D_t(j) = D_t$. As in the endowment economy analyzed in the previous section we assume that the borrowing constraint is not binding in equilibrium, so that an Euler equation like (1) holds for all households at all times.

The supply side of the economy is kept as simple as possible, and such that it remains insulated from the effects of idiosyncratic risk. This allows us to focus on the impact of

²⁰The assumption of a riskless real bond implies that we are abstracting from the redistributive effects due to inflation (Fisher's debt deflation channel). Changes in the real interest rate, however, still have differential income effects on households, depending on their individual net wealth positions.

the latter on aggregate demand (which coincides with aggregate consumption in our simple model), in the spirit of Werning (2015).

On the production side, we assume a continuum of firms, indexed by $i \in [0, 1]$. Each firm produces a differentiated good with the linear technology

$$Y_t(i) = A_t N_t(i) \quad (14)$$

where $N_t(i)$ is the quantity of labor (expressed in efficiency units) hired by firm i , and $A_t \equiv \exp\{a_t\}$ is an exogenous technology parameter common to all firms. Each firm sets the price of its good optimally each period, subject to a quadratic adjustment cost $\frac{\xi}{2} P_t Y_t \left(\frac{P_t(i)}{P_{t-1}(i)} - 1 \right)^2$ where $\xi > 0$, and a sequence of demand constraints $Y_t(i) = (P_t(i) / P_t)^{-\epsilon} Y_t$, where Y_t denotes aggregate output. Profit maximization, combined with the symmetric equilibrium conditions $P_t(i) = P_t$ and $Y_t(i) = Y_t$ for all $i \in [0, 1]$, implies:

$$\Pi_t (\Pi_t - 1) = \mathbb{E}_t \left\{ \Lambda_{t,t+1} \left(\frac{Y_{t+1}}{Y_t} \right) \Pi_{t+1} (\Pi_{t+1} - 1) \right\} + \frac{\epsilon}{\xi} \left(\frac{W_t (1 - \tau)}{A_t} - \frac{1}{\mathcal{M}_p} \right) \quad (15)$$

where $\Pi_t \equiv P_t / P_{t-1}$ is (gross) price inflation rate and $\mathcal{M}_p \equiv \epsilon / (\epsilon - 1) > 1$ is the desired (or flexible price) price markup. The term τ denotes a proportional labor subsidy, which is set to eliminate all the steady-state distortions due to monopolistic power in the goods and labor markets, and is financed with lump-sum taxes on firms.²¹ Aggregate profits are then given by $D_t = Y_t \Delta^p(\Pi_t) - W_t N_t$ where $\Delta^p(\Pi_t) \equiv 1 - (\xi/2) (\Pi_t - 1)^2$.

We assume a wage schedule

$$W_t = \mathcal{M}_w C_t^\sigma N_t^\varphi \quad (16)$$

where $C_t \equiv \int_0^1 C_t(j) dj$ and $N_t \equiv \int_0^1 N_t(i) di$ denote aggregate consumption and employment, respectively, and where $\mathcal{M}_w > 1$ is a constant (gross) average wage markup.

Combining equations (16)-(15), and taking a first-order approximation around the zero-inflation steady state gives the well known New Keynesian Phillips curve

$$\pi_t = \beta \mathbb{E}_t \{ \pi_{t+1} \} + \kappa \tilde{y}_t \quad (17)$$

where $\kappa \equiv (\sigma + \phi) (\epsilon - 1) / \xi$, and where $\tilde{y}_t \equiv y_t - y_t^n$ denotes the output gap, which is the difference between (log) output y_t and its natural (i.e. flexible price) counterpart $y_t^n \equiv a_t (1 + \phi) / (\sigma + \phi)$. Note that the latter is independent from monetary policy and,

²¹Formally, the subsidy is chosen such that $\mathcal{M}^p \mathcal{M}^w (1 - \tau) = 1$, where \mathcal{M}^w is a wage markup introduced below.

importantly, is unaffected by idiosyncratic risk.

Regarding monetary policy, we assume the central bank controls directly the real interest rate \hat{r}_t , which follows an exogenous $AR(1)$ process $\hat{r}_t = \rho_r \hat{r}_{t-1} + \varepsilon_{m,t}$, where $\mathbb{E}_t\{\varepsilon_{m,t+1}\} = 0$. This specification allows us to isolate the (direct) effects of idiosyncratic income on aggregate demand, abstracting from the potential (indirect) effects due to a different endogenous monetary policy response. In Appendix C, we also consider a case where the central bank follows a Taylor-type rule for the real interest rate, and show that our main qualitative findings remain unaltered.

In the symmetric equilibrium $Y_t(i) = Y_t$ and $C_t(i) = C_t$ for all $i \in [0, 1]$. Thus, market clearing in the goods market requires

$$C_t = Y_t \Delta^p(\Pi_t) \quad (18)$$

Market clearing in the bonds markets implies that $\int_0^1 B_t(j) dj = 0$ for all t . Aggregate employment is given by $N_t = Y_t/A_t$. We assume firms distribute their demand for work hours uniformly across households, i.e. $\mathcal{N}_t(j) = N_t$ for all $j \in [0, 1]$.²² Clearing of the labor market $N_t = \int_0^1 \mathcal{N}(j) \exp\{\zeta_t(j)\} dj$ is then guaranteed by the fact that $\int_0^1 \exp\{\zeta_t(j)\} dj = 1$.

Up to a first-order approximation and in a neighborhood of the zero inflation steady state (18) can be written as

$$c_t = y_t$$

Combining the previous condition with the Euler equation for aggregate consumption derived in Section 3 we obtain a version of the dynamic IS equation:

$$y_t = \mathbb{E}_t\{y_{t+1}\} - \frac{1}{\sigma} \hat{r}_t - \frac{\sigma + 1}{2} \hat{v}_t$$

Iterating forward the previous condition and imposing $\lim_{T \rightarrow \infty} \mathbb{E}_t\{y_{t+T}\} = 0$ (which is the steady state natural output, given our assumptions we obtain the following expression for (log) aggregate output

$$y_t = \underbrace{-\frac{1}{\sigma} \sum_{k=0}^{\infty} \mathbb{E}_t\{\hat{r}_{t+k}\}}_{\text{RA model}} - \underbrace{\frac{\sigma + 1}{2} \sum_{k=0}^{\infty} \mathbb{E}_t\{\hat{v}_{t+k}\}}_{\text{Risk component}} \quad (19)$$

²²Thus, we implicitly assume $W_t \exp\{\zeta_t(j)\} \geq C_t(j)^\sigma N_t^\varphi$ holds for all $j \in [0, 1]$ and all t , so that all households are willing to supply the work hours demanded by firms at a wage W_t (per efficiency unit).

The first term in the previous expression corresponds to equilibrium output in the RA version of the New Keynesian model. The second term reflects the impact of idiosyncratic risk on equilibrium output, which is decreasing in current and anticipated risk shifter — through its effects on precautionary savings. As discussed in section 3, the response of the risk shifter to an aggregate shock is given by a consumption-weighted average of the responses of individual consumption risk. Formally, letting $\hat{y}_t^H \equiv -\frac{\sigma+1}{2} \sum_{k=0}^{\infty} \mathbb{E}_t \{\hat{v}_{t+k}\}$ denote the component of aggregate output fluctuations associated with changes in the risk shifter, we can write:

$$\begin{aligned}
\frac{dy_{t+k}^H}{d\varepsilon_t} &= -\frac{\sigma+1}{2} \sum_{k=0}^{\infty} \frac{dv_{t+k}}{d\varepsilon_t} \\
&\simeq -\frac{\sigma+1}{2} \sum_{k=0}^{\infty} \int \frac{C_{t+k}(j)}{C_{t+k}} \frac{dv_{t+k}(j)}{d\varepsilon_t} dj \\
&\simeq -\frac{\sigma+1}{2} \sum_{k=0}^{\infty} \left[\frac{d\bar{v}_{t+k}}{d\varepsilon_t} + cov_j \left\{ c_{t+k}(j), \frac{dv_{t+k}(j)}{d\varepsilon_t} \right\} \right].
\end{aligned} \tag{20}$$

In the numerical simulations shown below for a calibrated version of our model, the dynamic response of consumption risk to an aggregate shock is larger for low consumption households. As a result, the impact of the shock on average consumption risk is muted by the distribution channel, leading to a small aggregate impact.

5.1 Calibration

We set $\beta = 0.9937$ and $\sigma = 1$ as in the endowment economy analyzed above, and consider the same calibration for the idiosyncratic shock $\zeta_t(j)$. In addition, we set the (inverse) Frisch elasticity of substitution to unity ($\varphi = 1$). Also, we set the elasticity of substitution among good varieties $\epsilon = 11$, which implies an average price markup of about 10 percent and the price adjustment cost parameter ζ so that the resulting slope of the Phillips Curve is $\kappa = 0.10$, in line with available estimates. Regarding the persistence of aggregate shocks, we assume that $\rho_a = 0.9$ and $\rho_r = 0.5$. We adopt the same numerical solution method described in Section 4.1.

5.2 Findings

We now analyze how idiosyncratic risk affects the response of our New Keynesian economy to monetary policy and technology shocks. For concreteness we focus on the response of aggregate output, and assume that the monetary policy rule takes the form of

an exogenous process for the real rate, as introduced above. In Appendix C we show results are similar when considering a standard Taylor rule.

Figure 5 shows the response of aggregate output to a 25 basis point expansionary monetary shock, which leads to a 100 basis point reduction in the (annualized) real interest rate. The figure displays that response for three economies: our baseline model with idiosyncratic income risk (red line with circles), and economy with idiosyncratic risk but no distribution channel (green dashed line) and a representative agent economy (blue line with crosses).

Note that the presence of idiosyncratic risk tends to amplify the output effects of the monetary policy shock. The effects are stronger on impact, and more persistent. However, from a quantitative viewpoint, the magnitude of this amplification seems very small—less than 0.05 percentage points at all horizons. That small effect arises despite the quantitatively large change in the average risk component, as captured by the green dashed line. The reason for the difference between the latter effect and the total effect of consumption risk lies in the offsetting impact of the distribution channel: the decrease in risk is concentrated on low consumption households, which tends to mute the overall impact on aggregate consumption and output.

Finally, Figure 6 shows the dynamic responses to a positive technology shock. Again the difference between the models with and without heterogeneity in terms of the responses of output and inflation is quantitatively negligible, due to the offsetting distribution channel.²³

6 Concluding Remarks

The objective of the present paper was to study the role of idiosyncratic income risk for aggregate fluctuations within a simple heterogeneous household framework with no binding borrowing constraints. We derive analytically an approximate Euler equation for (log) aggregate consumption, which helps us shed some light on the differential behavior of such an economy relative to its representative agent counterpart. In particular, we show that those differences are related to how changes in consumption risk are distributed among households, as captured by a consumption-weighted average of changes in consumption risk.

²³Note that output remains unchanged in response to the technology shock. This is due to the constancy of the real rate implied by our baseline monetary policy rule. See Appendix C for corresponding results under a standard Taylor rule.

Our findings raise several issues that are relevant to current efforts to introduce heterogeneity in models of aggregate fluctuations.

Firstly, an implication of our findings is that idiosyncratic risk may have to be combined with other ingredients to have a significant impact on aggregate fluctuations. The assumption of financial frictions in the form of binding borrowing constraints is a prominent candidate to play that role. From that viewpoint, our findings can be interpreted as providing a rationale for the widespread adoption of that assumption in the recent literature, in addition to its arguable realism. On the other hand, our findings may also be read as suggesting that one may want to ignore altogether idiosyncratic risk when introducing heterogeneity in macro models, focusing instead on the presence of a binding borrowing constraint. This is the approach adopted in models with a constant fraction of hand-to-mouth households (as exemplified by the TANK models of Galí et al. (2007), Bilbiie (2008, 2021), and Broer et al. (2019)). In a companion paper (Debortoli and Galí (2018)), we analyze the extent to which the predictions of richer heterogeneous agent models, with non-trivial interactions between idiosyncratic risk and borrowing constraints, can be approximated by two-agent models that abstract from idiosyncratic risk.

Secondly, an implication of our findings is that idiosyncratic risk is likely to have a small impact on aggregate fluctuations in economies where fluctuations in consumption risk are concentrated among poorer (low consumption) households, as is the case in the quantitative example economies studied above –an endowment economy and a New Keynesian economy. Conversely, such idiosyncratic risk may be more relevant in economies where rich (i.e. high consumption) households experience large fluctuations in consumption risk, as it is likely to be the case in recent models in which a fraction of wealthy households behave in a hand-to-mouth fashion, possibly as a result of the low liquidity of their wealth (e.g. Kaplan et al. (2018)). Thus, and even though changes in consumption risk resulting from aggregate shocks may not (directly) impinge on the consumption of currently constrained households (wealthy or not), it may still be the case that those changes in consumption risk are relevant for households “close to the constraint,” which in the context of those models also include relatively wealthy (high consumption) households, with a consequent larger impact on aggregate consumption.

Thirdly, it should be clear that how aggregate shocks affect uncertainty for different types of households is ultimately an empirical question, and one which we plan to address in future work using micro data, in a similar spirit to the recent work of Berger et al. (2022).

References

- Acharya, Sushant and Keshav Dogra (2020): "Understanding HANK: Insights From a PRANK," *Econometrica*, 88 (3), 1113–1158.
- Auclert, Adrien (2019): "Monetary Policy and the Redistribution Channel," *American Economic Review*, 109 (6): 2333–67.
- Auclert, Adrien, Matthew Rognlie and Ludwig Straub (2018): "The Intertemporal Keynesian Cross," NBER working paper 25020.
- Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub (2021): "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 89 (5), 2375–2408.
- Bayer, Christian, Ralph Luetticke, 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.
- Berger, David, Luigi Bocola, and Alessandro Dovis (2022): "Imperfect Risk Sharing and the Business Cycle," working paper.
- Bilbiie, Florin O. (2008): "Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic," *Journal of Economic Theory*, 140 (1), 162–196.
- Bilbiie, Florin O. (2021): "Monetary Policy and Heterogeneity: An Analytical Framework," mimeo.
- Broer, Tobias, Niels-Jakob Harbo Hansen, Per Krusell, and Erik Oberg (2020): "The New Keynesian Transmission Mechanism: A Heterogeneous-Agent Perspective," *The Review of Economic Studies*, 87 (1), 77–101.
- Boppart, Timo, Per Krusell, and Kurt Mitman (2018): "Exploiting MIT Shocks in Heterogeneous-Agent Economies: The Impulse Response as a Numerical Derivative," *Journal of Economic Dynamics and Control*, 89, 68–92.
- Campbell, John Y. and N. Gregory Mankiw (1989): "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," *NBER Macroeconomics Annual*, 4, 185–216.
- Carroll, Christopher D. (2006): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 91 (3), 312–320.

Cao, Dan and Guangyu Nie (2017): "Amplification and Asymmetric Effects without Collateral Constraints," *American Economic Journal: Macroeconomics* 9 (3), 222-266.

Challe, Edouard, and Xavier Ragot (2011): "Fiscal Policy in a Tractable Liquidity-Constrained Economy," *The Economic Journal*, 121 (551), 273-317.

Challe, Edouard, Julien Matheron, Xavier Ragot, and Juan F. Rubio-Ramirez (2017): "Precautionary Saving and Aggregate Demand," *Quantitative Economics*, 8 (2), 435-478.

Debortoli, Davide and Jordi Galí (2018): "Monetary Policy with Heterogeneous Agents: Insights from TANK Models," mimeo.

Galí, Jordi, David López-Salido and Javier Vallés (2007): "Understanding the Effects of Government Spending on Consumption," *Journal of the European Economic Association*, 5 (1), 227-270.

Guvenen, Fatih (2011): "Macroeconomics with heterogeneity: a Practical Guide," *Economic Quarterly*, 255-326.

Kaplan, Greg and Giovanni L. Violante, (2014): "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82 (4), 1199-1239.

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

Kimball, Miles S. (1990): "Precautionary Saving in the Small and in the Large," *Econometrica*, 58 (1), 53-73.

Krueger, Dirk, Kurt Mitman and Fabrizio Perri (2016): "Macroeconomics and Household Heterogeneity," *Handbook of Macroeconomics*, vol.2, chapter 11, 843-921.

Krusell, Per and Anthony A. Smith, (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106 (5), 867-896.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2009): "Quantitative Macroeconomics with Heterogeneous Households," *Annual Review of Economics*, 1 (1), 319-354.

Luetticke, Ralph (2021): "Transmission of Monetary Policy with Heterogeneity in Household Portfolios," *American Economic Journal: Macroeconomics*, 13 (2), 1-25.

Marcet, Albert and Kenneth J. Singleton (1999): "Equilibrium Asset Prices and Savings of Heterogeneous Agents in the Presence of Incomplete Markets and Portfolio Constraints", *Journal of Economic Theory* 187, 1–38.

Ma, Qingyin, John Stachurski, and Alexis A. Toda (2020): "The Income Fluctuations Problem and the Evolution of Wealth," *Journal of Economic Theory* 187, 1–38.

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

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

Reiter, Michael (2009): "Solving Heterogeneous-Agent Models by Projection and Perturbation," *Journal of Economic Dynamics and Control*, 33 (3), 649–665.

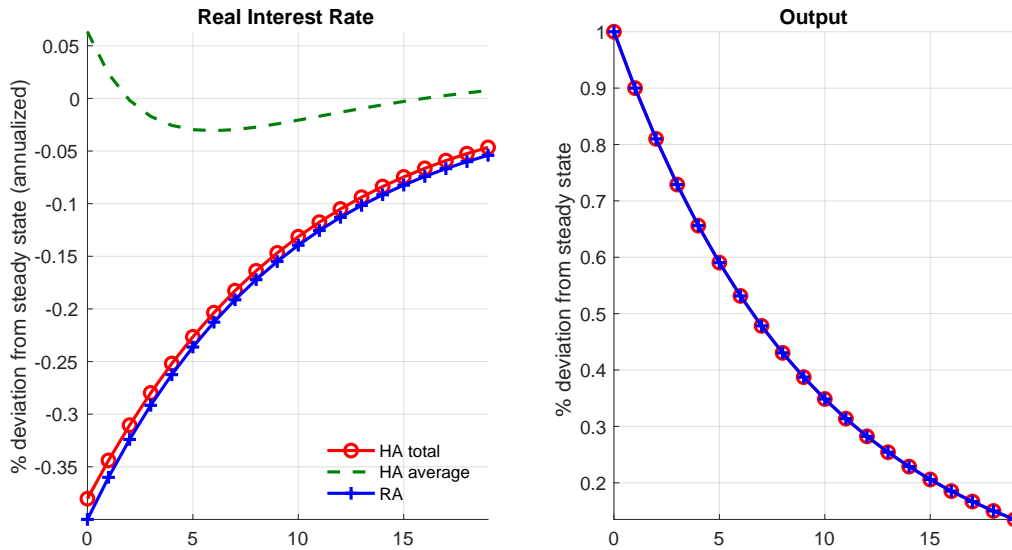
Werning, Iván (2015): "Incomplete Markets and Aggregate Demand," Working Paper 21448, National Bureau of Economic Research.

Tables and Figures

Table 1: Calibration of the Endowment Economy

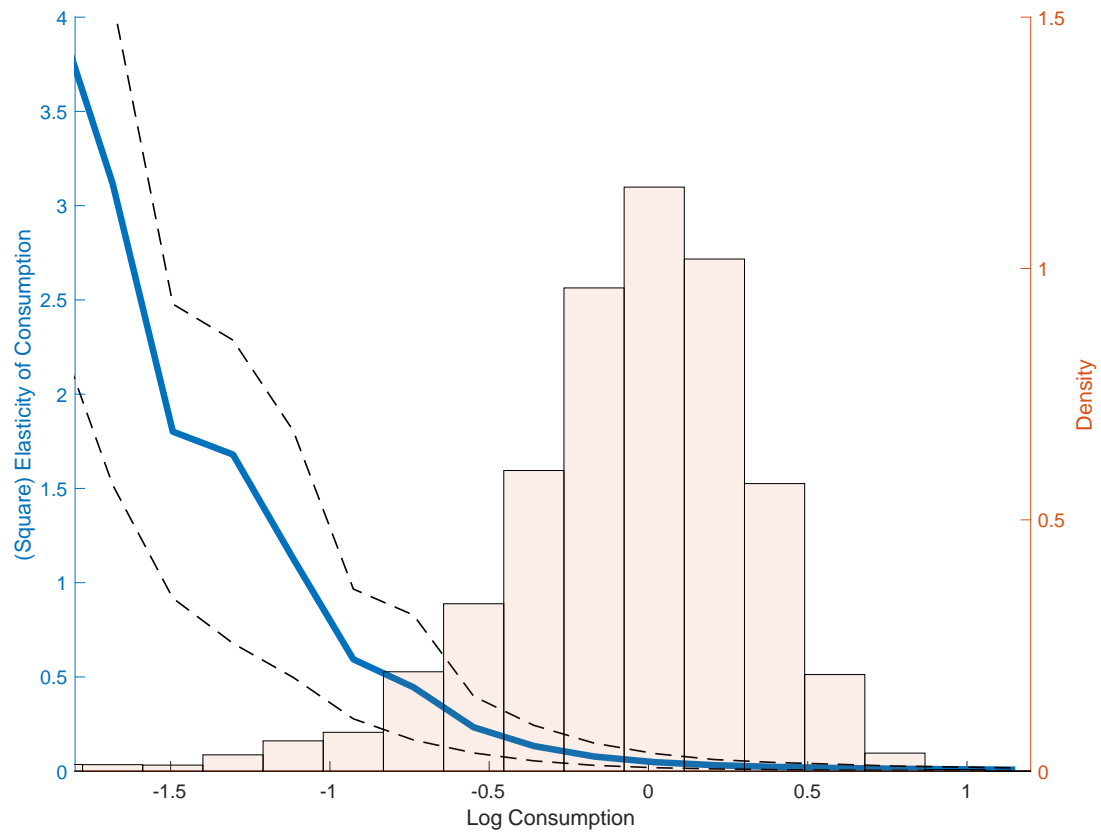
Parameter	Meaning	Value
Model parameters		
σ	Coefficient of Risk Aversion	1
\bar{r}	Steady State Interest Rate (annualized)	0.02
ρ_y	Autocorr. of agg. endowment shocks	0.9
ρ_ζ	Autocorr. of idiosyncratic earnings	0.966
σ_ζ	Std. dev. of idiosyncratic earnings	0.5
Discretization		
n_ζ	Points in Markov Chain for ζ	11
n_a	Points in Markov Chain for Assets	500

Figure 1: The Effects of an Aggregate Endowment Shock



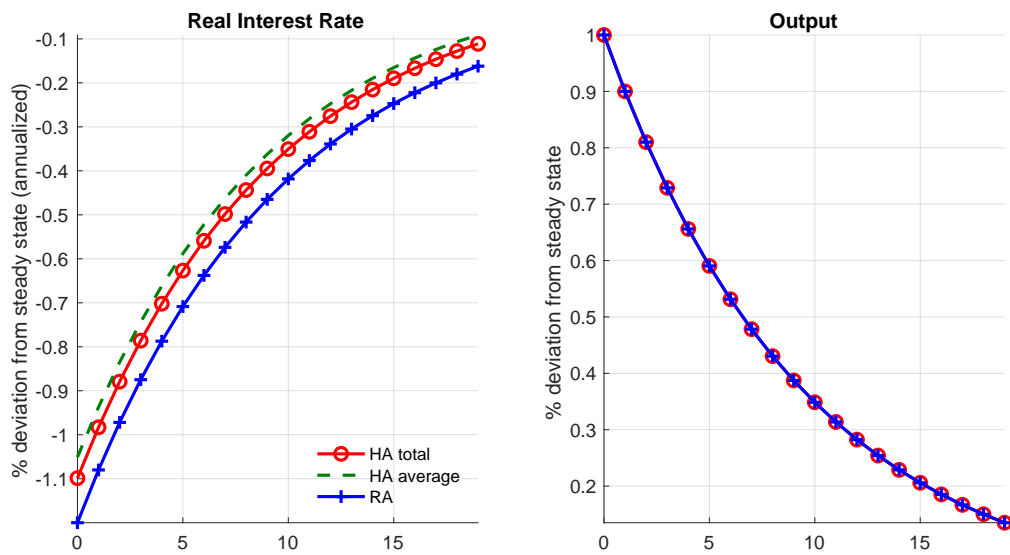
Notes: The figure shows the response of the annualized real interest rate (left panel) to a positive aggregate endowment shock (right panel) in a representative agent model (blue line with crosses), in the baseline model with heterogeneity (red line with circles), and in a model with heterogeneity but considering only the average consumption risk channel (dashed green line).

Figure 2: Elasticity of Consumption in Steady State



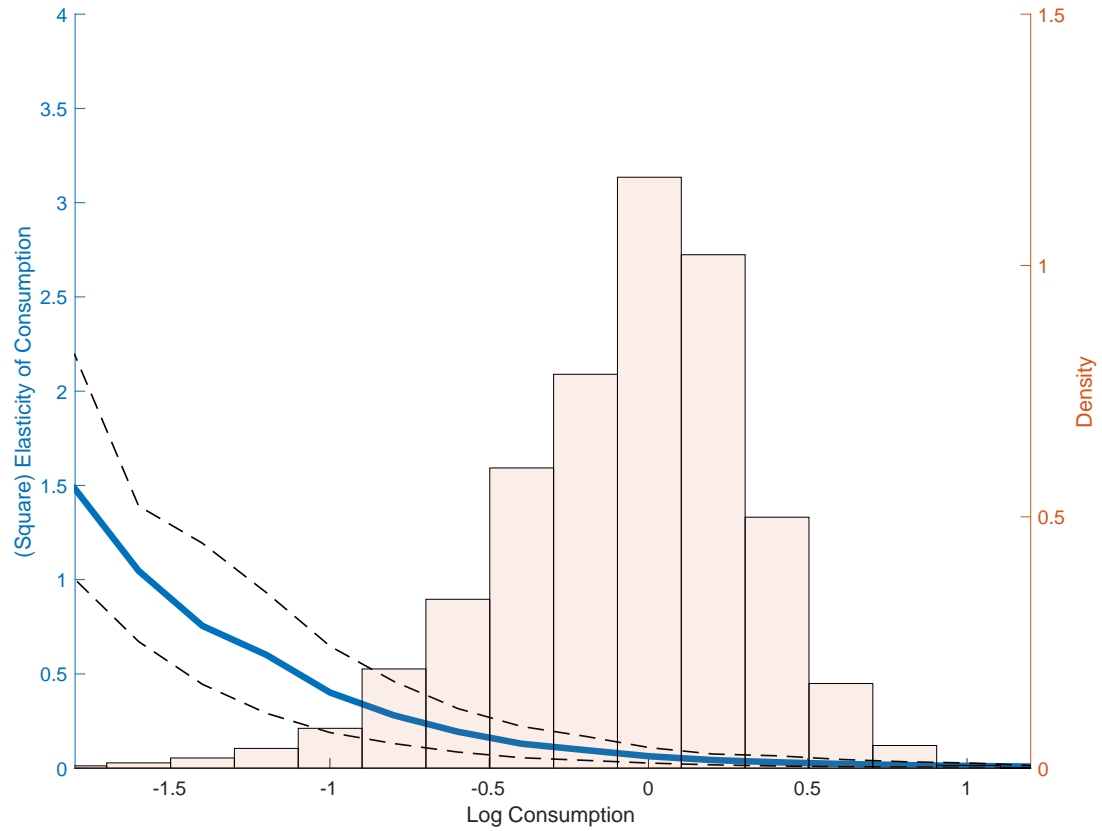
Notes: The figure shows the relationship between log consumption (horizontal axis), and the elasticity of consumption (left vertical axis) in steady state. For each value of consumption, the figure reports the average elasticity (solid blue line), the 5% - 95% interval of the distribution (black dashed lines), while the histogram indicate the steady state distribution (right vertical axis).

Figure 3: The Effects of an Aggregate Endowment Shock: $\sigma = 3$



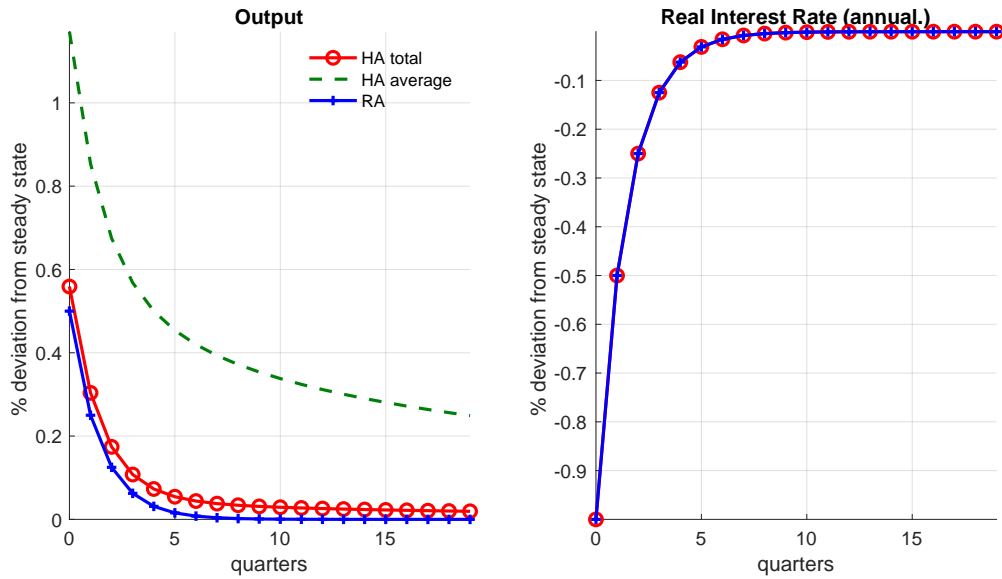
Notes: The figure shows the response of the annualized real interest rate (left panel) to a positive aggregate endowment shock (right panel) in a representative agent model (blue line with crosses), in the baseline model with heterogeneity (red line with circles), and in a model with heterogeneity but considering only the effect of average risk (dashed green line).

Figure 4: Elasticity of Consumption in Steady State: $\sigma = 3$



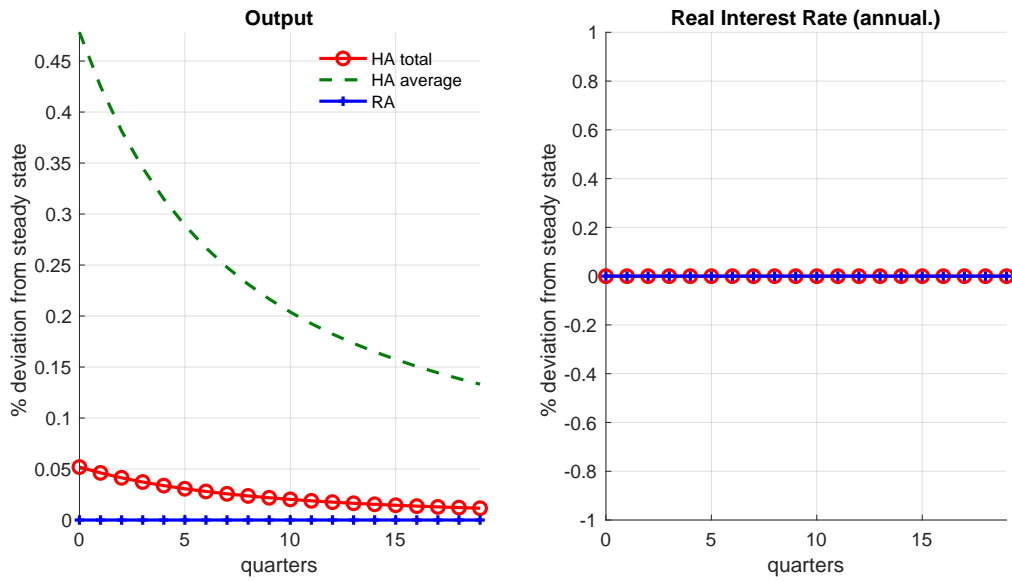
Notes: The figure shows the relationship between log consumption (horizontal axis), and the elasticity of consumption (left vertical axis) in steady state. For each value of consumption, the figure reports the average elasticity (solid blue line), the 5% - 95% interval of the distribution (black dashed lines), while the histogram indicate the steady state distribution (right vertical axis).

Figure 5: The Effects of a Monetary Policy Shock



Notes: The figure shows the response of output to a 1 percent decrease in the (annualized) real interest rate, in a representative agent model (blue line with crosses), in the baseline model with heterogeneity (red line with circles), and in a model with heterogeneity but considering only the effect of average uncertainty (dashed green line).

Figure 6: The Effects of a Technology Shock



Notes: The figure shows the responses of output and the real interest rate to a 1 percent positive technology shock, in a representative agent model (blue line with crosses), in the baseline model with heterogeneity (red line with circles), and in a model with heterogeneity but considering only the effect of average risk (dashed green line).

Appendices

A.1 Derivation of the approximate individual Euler equation

Our starting point is the individual Euler equation

$$C_t(j)^{-\sigma} = \beta R_t \mathbb{E}_t \{ C_{t+1}(j)^{-\sigma} \}$$

A second order approximation of $C_{t+1}(j)^{-\sigma}$ around $C_t(j)$ yields

$$C_t(j)^{-\sigma} \simeq \beta R_t \mathbb{E}_t \left\{ C_t(j)^{-\sigma} - \sigma C_t(j)^{-\sigma} \left(\frac{\Delta C_{t+1}(j)}{C_t(j)} \right) + \frac{\sigma(\sigma+1)}{2} C_t(j)^{-\sigma} \left(\frac{\Delta C_{t+1}(j)}{C_t(j)} \right)^2 \right\}.$$

Rearranging terms,

$$\mathbb{E}_t \left\{ \frac{\Delta C_{t+1}(j)}{C_t(j)} \right\} \simeq \frac{1}{\sigma} \left(1 - \frac{1}{\beta R_t} \right) + \frac{\sigma+1}{2} v_t(j)$$

where $v_t(j) \equiv \mathbb{E}_t \left\{ \left(\frac{\Delta C_{t+1}(j)}{C_t(j)} \right)^2 \right\}$, which corresponds to eq. (2) in the main text.

Letting $c_t(j) \equiv \log C_t(j)$ and using the Taylor expansion $\frac{\Delta C_{t+1}(j)}{C_t(j)} \simeq \Delta c_{t+1}(j) + \frac{1}{2} (\Delta c_{t+1}(j))^2$, we can rewrite the Euler equation in terms of (log) consumption:

$$\mathbb{E}_t \{ \Delta c_{t+1}(j) \} \simeq \frac{1}{\sigma} \left(1 - \frac{1}{\beta R_t} \right) + \frac{\sigma}{2} v_t(j) \quad (\text{A.1})$$

Evaluating the previous equation at the stochastic steady state (with $R_t = R$), and taking unconditional expectations we have

$$0 \simeq \frac{1}{\sigma} \left(1 - \frac{1}{\beta R} \right) + \frac{\sigma}{2} \mathbb{E} \{ v_t(j) \} \quad (\text{A.2})$$

Subtracting (A.2) from (A.1) and taking a first-order Taylor expansion of the resulting expression yields:

$$\mathbb{E}_t \{ \Delta c_{t+1}(j) \} \simeq \frac{1}{\sigma} \hat{r}_t + \frac{\sigma}{2} \hat{v}_t(j), \quad (\text{A.3})$$

where $\hat{r}_t \equiv \frac{1}{\beta R} \left(\frac{R_t - R}{R} \right)$ and $\hat{v}_t(j) \equiv v_t(j) - \mathbb{E} \{ v_t(j) \}$. Thus, it follows that $(\mathbb{E}_t \{ \Delta c_{t+1}(j) \})^2 \sim \mathcal{O}(|\varepsilon|^2)$ thus implying $\mathbb{E}_t \{ \Delta c_{t+1}(j)^2 \} \simeq \mathbb{E}_t \{ \zeta_{t+1}(j)^2 \}$, where $\zeta_t(j) \equiv c_t(j) - \mathbb{E}_{t-1} \{ \Delta c_t(j) \}$

is the innovation in household j 's (log) consumption. Accordingly, we have

$$v_t(j) \equiv \mathbb{E}_t \left\{ \left(\frac{\Delta C_{t+1}(j)}{C_t(j)} \right)^2 \right\} \simeq \mathbb{E}_t \{ \xi_{t+1}(j)^2 \}.$$

A.2 Derivation of an exact Euler equation for aggregate consumption

Starting from the individual Euler equation

$$C_t(j)^{-\sigma} = \beta R_t \mathbb{E}_t \{ C_{t+1}(j)^{-\sigma} \}$$

and multiplying and dividing the RHS by $[\mathbb{E}_t \{ C_{t+1}(j) \}]^{-\sigma}$ we have

$$C_t(j)^{-\sigma} = \beta R_t [\mathbb{E}_t \{ C_{t+1}(j) \}]^{-\sigma} \frac{\mathbb{E}_t \{ C_{t+1}(j)^{-\sigma} \}}{[\mathbb{E}_t \{ C_{t+1}(j) \}]^{-\sigma}}$$

or equivalently

$$C_t(j) V_t(j) = (\beta R_t)^{-\frac{1}{\sigma}} \mathbb{E}_t \{ C_{t+1}(j) \} \quad (\text{A.4})$$

where $V_t(j) \equiv \left[\frac{\mathbb{E}_t \{ C_{t+1}(j)^{-\sigma} \}}{[\mathbb{E}_t \{ C_{t+1}(j) \}]^{-\sigma}} \right]^{\frac{1}{\sigma}} \geq 1$ captures the effects of individual consumption risk on individual consumption choices, i.e. the "wedge" relative to the certainty-equivalence case.

Next, dividing and multiplying the LHS of (A.4) by aggregate consumption C_t , and integrating across households (and abstracting from aggregate uncertainty, as we do in our quantitative exercises) we get

$$C_t \int \frac{C_t(j)}{C_t} V_t(j) dj = (\beta R_t)^{-\frac{1}{\sigma}} \int \mathbb{E}_t \{ C_{t+1}(j) \} dj$$

which implies that

$$C_t = (\beta R_t)^{-\frac{1}{\sigma}} C_{t+1} V_t^{-1} \quad (\text{A.5})$$

where $V_t \equiv \int \frac{C_t(j)}{C_t} V_t(j) dj$.

Finally, in terms of log-deviations from steady state we have

$$\hat{c}_t = \hat{c}_{t+1} - \frac{1}{\sigma} \hat{r}_t - \hat{v}_t \quad (\text{A.6})$$

which is analogous to eq. (6) in the main text.

A.3 Derivation of the approximate individual Euler equation for a general utility function

The individual Euler equation under a general utility function $U(\cdot)$ is given by

$$U'(C_t(j)) = \beta R_t \mathbb{E}_t \{ U'(C_{t+1}(j)) \}.$$

Define $\sigma_t(j) \equiv -U''(C_t(j))C_t(j)/U'(C_t(j))$ (relative risk aversion) and $\varkappa_t(j) \equiv -U'''(C_t(j))C_t(j)/U''(C_t(j))$ (relative prudence). Approximating $U'(C_{t+1}(j))$ around $C_t(j)$ gives

$$U'(C_{t+1}(j)) \simeq U'(C_t(j)) + U''(C_t(j))\Delta C_{t+1}(j) + \frac{1}{2}U'''(C_t(j))(\Delta C_{t+1}(j))^2.$$

Substituting for $U'(C_{t+1}(j))$ in the Euler equation using the previous approximation we obtain

$$1 \simeq \beta R_t \mathbb{E}_t \left\{ 1 - \sigma_t(j) \frac{\Delta C_{t+1}(j)}{C_t(j)} + \frac{1}{2} \sigma_t \varkappa_t \left(\frac{\Delta C_{t+1}(j)}{C_t(j)} \right)^2 \right\}.$$

which gives the approximate Euler equation for aggregate consumption

$$\mathbb{E}_t \Delta C_{t+1}(j) \simeq - \frac{U'(C_t(j))}{U''(C_t(j))} \left(1 - \frac{1}{\beta R_t} \right) - \frac{1}{2} \frac{U'''(C_t(j))}{U''(C_t(j))} \mathbb{E}_t [\Delta C_{t+1}(j)]^2 \quad (\text{A.7})$$

Dividing by $C_t(j)$ and using our definitions of relative risk aversion and relative prudence gives

$$\mathbb{E}_t \left\{ \frac{\Delta C_{t+1}(j)}{C_t(j)} \right\} \simeq \frac{1}{\sigma_t(j)} \left(1 - \frac{1}{\beta R_t} \right) + \frac{\varkappa_t}{2} \mathbb{E}_t \left\{ \left(\frac{\Delta C_{t+1}(j)}{C_t(j)} \right)^2 \right\}.$$

Note that a CRRA utility implies $\sigma_t(j) = \sigma$ and $\varkappa_t(j) = \sigma + 1$, so the previous expression collapses to eq. (3) in the main text.

Alternatively, denoting with $\tilde{\sigma}_t(j)$ and $\tilde{\varkappa}_t(j)$ the coefficients of absolute risk aversion and prudence gives

$$\mathbb{E}_t \Delta C_{t+1}(j) \simeq \frac{1}{\tilde{\sigma}_t(j)} \left(1 - \frac{1}{\beta R_t} \right) - \frac{\tilde{\varkappa}_t(j)}{2} \mathbb{E}_t \left\{ \Delta C_{t+1}(j)^2 \right\}.$$

For example, the special case of CARA preferences implies that $\tilde{\sigma}_t(j) = \tilde{\varkappa}_t(j) \equiv \tilde{\sigma}$, so

the previous expression collapses to

$$\mathbb{E}_t \Delta C_{t+1}(j) \simeq \frac{1}{\bar{\sigma}} \left(1 - \frac{1}{\beta R_t} \right) - \frac{\bar{\sigma}}{2} \tilde{v}_t(j)$$

where $\tilde{v}_t(j) \equiv \mathbb{E}_t \left\{ \Delta C_{t+1}(j)^2 \right\} \simeq \mathbb{E}_t \left\{ [C_{t+1}(j) - \mathbb{E}_t C_{t+1}(j)]^2 \right\} \equiv \mathbb{E}_t \left\{ \xi_{t+1}(j)^2 \right\}$, which is analogous to eq. (3) in the main text. In this economy, under the assumption of i.i.d. idiosyncratic income shocks $y_t(j) \sim N(0, \sigma_{y,t})$, it can be shown that individual consumption is a linear function of cash-on-hand $x_t(j)$, i.e. $C_t(j) = C_t + \mu_t x_t(j)$, where μ_t denotes the marginal propensity to consume, and is constant across households (see Acharya and Dogra (2020)). It then follows that consumption risk $\tilde{v}_t(j) = \bar{v}_t = \mu_{t+1}^2 \sigma_{y,t}^2$, is common across households, and thus the distribution channel described in the main text is absent.

A.4 Derivation of the dynamic response of v_t

Recalling that $v_t \equiv \int \frac{C_t(j)}{C_t} v_t(j) dj$ we have

$$\begin{aligned} \frac{dv_{t+k}}{d\varepsilon_t} &= \int \frac{d[C_{t+k}(j)/C_{t+k}]}{d\varepsilon_t} v_t(j) dj + \int \frac{C_{t+k}(j)}{C_{t+k}} \frac{dv_{t+k}(j)}{d\varepsilon_t} dj \\ &= \int \frac{d \exp\{c_{t+k}(j) - c_{t+k}\}}{d\varepsilon_t} v_t(j) dj + \int \frac{C_{t+k}(j)}{C_{t+k}} \frac{dv_{t+k}(j)}{d\varepsilon_t} dj \\ &= \int \frac{d[c_{t+k}(j) - c_{t+k}]}{d\varepsilon_t} \frac{C_{t+k}(j)}{C_{t+k}} v_t(j) dj + \int \frac{C_{t+k}(j)}{C_{t+k}} \frac{dv_{t+k}(j)}{d\varepsilon_t} dj \end{aligned}$$

Next we derive an approximate expression for $\frac{d[c_{t+k}(j) - c_{t+k}]}{d\varepsilon_t}$. Combining the previous equation with (3) in the text and rearranging terms yields the difference equation

$$c_t(j) - c_t = \mathbb{E}_t \{(c_{t+1}(j) - c_{t+1})\} - \frac{\sigma}{2} v_t(j) + \frac{\sigma + 1}{2} v_t$$

which can be solved forward to obtain

$$c_t(j) - c_t = - \sum_{k=0}^{\infty} \left[\frac{\sigma}{2} \mathbb{E}_t \{v_{t+k}(j)\} + \frac{\sigma + 1}{2} \mathbb{E}_t \{v_{t+k}\} \right] + \mathbb{E} \{c_t(j) - c_t\} \quad (\text{A.8})$$

where we have used the fact that $\lim_{T \rightarrow \infty} \mathbb{E}_t \{c_{t+T}(j)\} = \mathbb{E} \{c_t(j)\}$ and $\lim_{T \rightarrow \infty} \mathbb{E}_t \{c_{t+T}\} = \mathbb{E} \{c_t\}$.

Using (A.8) as a reference, we can derive the dynamic response of (log) consumption

differential to an aggregate shock in period t :

$$\frac{d[c_{t+k}(j) - c_{t+k}]}{d\varepsilon_t} = - \sum_{h=k}^{\infty} \left[\frac{\sigma}{2} \frac{dv_{t+h}(j)}{d\varepsilon_t} + \frac{\sigma + 1}{2} \frac{dv_{t+h}}{d\varepsilon_t} \right] \sim \mathcal{O}(|\varepsilon|)$$

Accordingly, $\int \frac{d[c_{t+k}(j) - c_{t+k}]}{d\varepsilon_t} \frac{C_{t+k}(j)}{C_{t+k}} v_t(j) dj \sim \mathcal{O}(|\varepsilon|^2)$ and can thus be ignored in our approximation. Thus, it follows that

$$\frac{dv_{t+k}}{d\varepsilon_t} \simeq \int \frac{C_{t+k}(j)}{C_{t+k}} \frac{dv_{t+k}(j)}{d\varepsilon_t} dj$$

as found in the text.

B Robustness: Alternative Process for Idiosyncratic Income Shocks

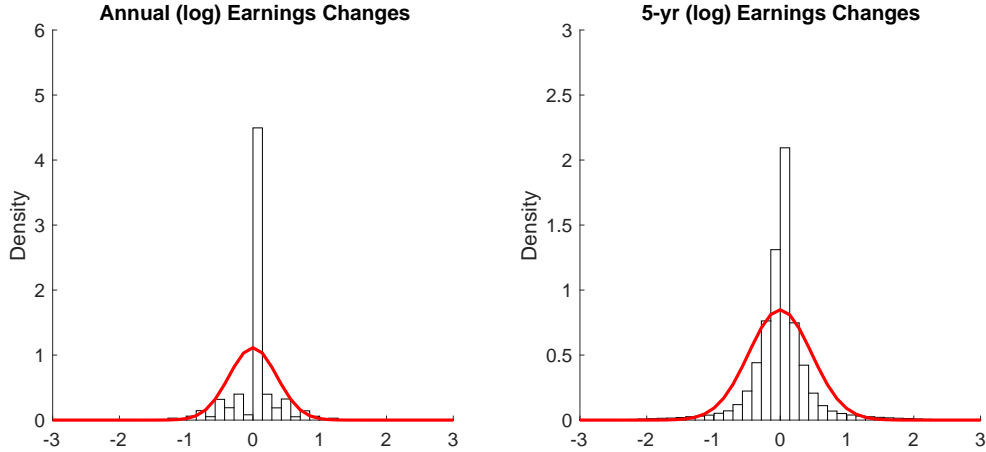
In this section, we study the role of heterogeneity in the New Keynesian economy described in Section 5, but considering an alternative process for the idiosyncratic income shocks $\zeta_t(i)$. In particular, we consider a discrete-time quarterly version of the continuous-time process used in Kaplan, Moll and Violante (2018), which is the sum of two independent components $\zeta_t(i) = \zeta_{1,t}(i) + \zeta_{2,t}(i)$. Both components evolve according to a “jump-drift” process, where jumps arrive at a Poisson rate $\lambda_1 = 0.080$ and $\lambda_2 = 0.007$ and where, conditionally on a jump, innovations are drawn from a normal distribution with mean zero and standard deviations $\sigma_1 = 1.74$ and $\sigma_2 = 1.53$. Between jumps, the processes drift toward zero at rates $\beta_1 = 0.0761$ and $\beta_2 = 0.009$, respectively. The two continuous-time components are discretized with 3 grid points for ζ_1 (transitory component) and 11 points for ζ_2 (persistent component) — see Section 4.2.2 and Appendix D in Kaplan, Moll, Violante (2018) for more details.

We calculate the corresponding Markov transition matrix at a quarterly frequency. The resulting discretized process gives rise to a leptokurtic distribution of income changes, as shown in Figure B.1. In particular, the values of the kurtosis are 14.8 for annual income changes, and 12.6 for 5-year changes, which are close to the empirical counterparts using data U.S. male earnings as in Guvenen et. al. (2015). We then recalibrate the discount factor to $\beta = 0.982$ so that the steady state real interest rate equal 2 percent per year, as in our baseline case.

Figure B.2 shows that the response of output to a monetary shock in this economy

(green line with diamonds) is remarkably close to the response obtained in our baseline calibration (red line with circles), and in turn similar its counterpart in a representative agent economy (blue line with crosses). A similar result is obtained in response to other shocks (results are omitted for brevity, and available from the authors upon request).

Figure B.1: Distribution of (Log) Income Shocks in the Alternative Calibration



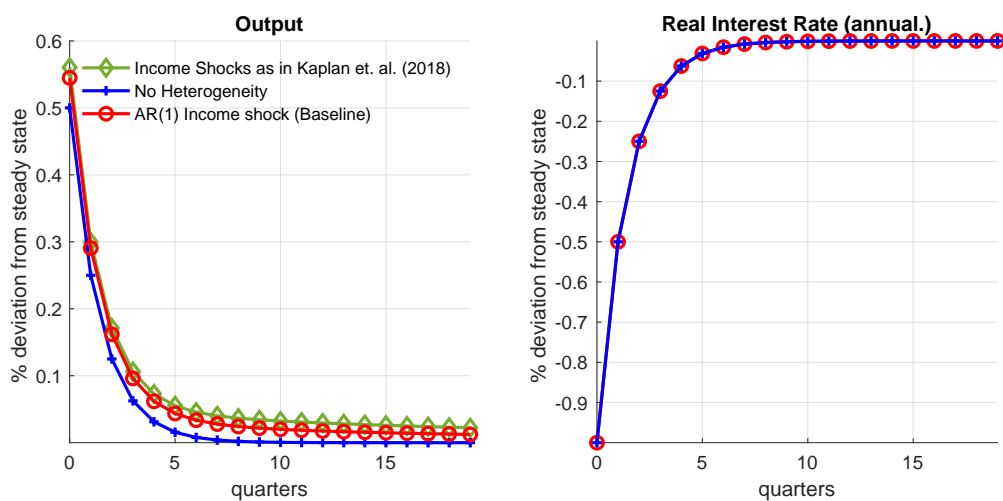
Notes: The figure shows the distribution of (log) earning changes at annual frequency (left panel) and at a 5yr frequency (right panel). In each panel, the histograms correspond to the distribution resulting from the (discretized) process with a transitory and a persistent component, while the solid line indicates the normal distribution with the same mean and variance.

C Robustness: Monetary Policy Rule

In this appendix, we study the role of heterogeneity in the New Keynesian economy described in Section 5, assuming that the central bank follows a Taylor-type rule for the real interest rate $\hat{r}_t = \phi_\pi \pi_t + m_t$, where m_t is a monetary shock, which is assumed to follow an AR(1) process, with auto-correlation coefficient $\rho_m = 0.5$. We set the coefficient $\phi_\pi = 0.5$, in line with the original estimates of Taylor (1999).

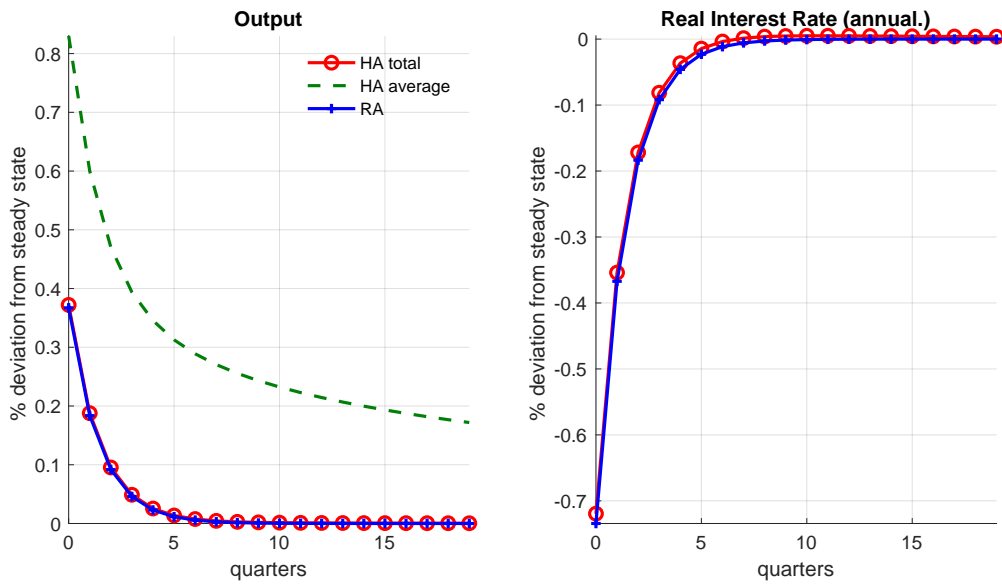
Figure C.3-C.4 report the response of aggregate variables to monetary and technology shocks, respectively. In response to all these shocks, and analogously to what shown in Figures 5 and 6 in the main text, the responses of aggregate variables in an heterogeneous agent economy (red lines with circles) are similar to those obtained in the corresponding model with a representative agent (blue line with crosses).

Figure B.2: The Effects of Monetary Shocks with Alternative Idiosyncratic Risk Process



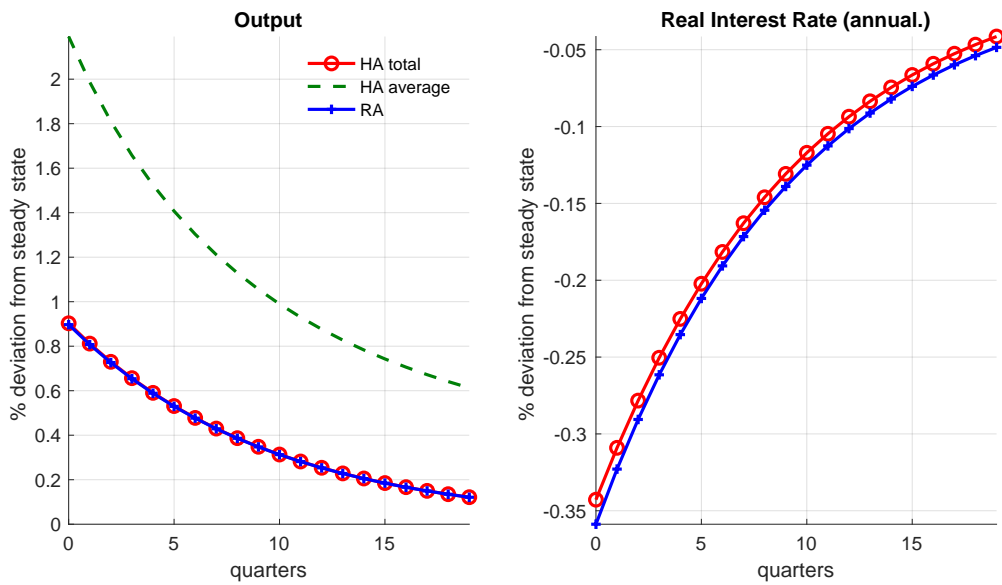
Notes: The figure shows the response of output and (annualized) real interest rate to a 25 basis points expansionary monetary shock. The figure compares the responses in a model without heterogeneity (blue line with crosses), in the baseline heterogeneous household model with AR(1) idiosyncratic income shocks (red line with circles), and in a model with idiosyncratic shocks with a transitory and a persistent component as in Kaplan, Moll and Violante (2018) (green line with diamonds).

Figure C.3: The Effects of a Monetary Shock (Monetary Rule)



Notes: The figure shows the response of output and (annualized) real interest rate to a 25 basis point monetary shock, in a representative agent model (blue line with crosses), in the baseline model with heterogeneity (red line with circles), and in a model with heterogeneity but considering only the effect of average consumption risk (dashed green line).

Figure C.4: The Effects of a Technology Shock (Monetary Rule)



Notes: The figure shows the response of output and (annualized) real interest rate to a 1 percent technology shock, in a representative agent model (blue line with crosses), in the baseline model with heterogeneity (red line with circles), and in a model with heterogeneity but considering only the effect of average consumption risk (dashed green line).