Idiosyncratic Income Risk and Aggregate Fluctuations*

Davide Debortoli
UPF, CREi and Barcelona GSE

Jordi Galí
CREi, UPF and Barcelona GSE

July 2021

Abstract

We study the role of idiosyncratic income shocks for aggregate fluctuations within a simple heterogeneous household framework with no binding borrowing constraints. We derive analytically an Euler equation for (log) aggregate consumption, and show that the impact of idiosyncratic risk on aggregate consumption works through two channels: (i) changes in average consumption uncertainty and (ii) changes in the cross-sectional dispersion of consumption. We show that these two channels are related and tend to offset each other. Their net effect is captured by a sufficient statistic, the consumption-weighted average of changes in uncertainty. We apply this framework to two example economies —an endowment economy and a New Keynesian economy— and show that the net effect of heterogeneity is quantitatively small. By contrast, that effect becomes more significant when considering that borrowing constraints are binding for a sizable fraction of the population.

JEL Classification Numbers: E21, E32, E50

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

*We thank Sushant Acharya, Susanto Basu, Benjamin Moll and seminar participants at the University of Maryland, Boston College, London Business School, Bank of Canada, Ecole Polytechnique/CREST and CREI Faculty lunch for their comments and suggestions, 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, Advanced Grant 882332-HEMPEF (Galí), 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 CEX2019-000915-S, the Ramon y Cajal fellowship RyC-2016-20476 (Debortoli) and from the Generalitat de Catalunya, through CERCA and SGR Programme (2017-SGR-1393).

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 growing popularity of Heterogeneous Agent New Keynesian (HANK) models are a reflection of this emerging view. 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 and effects of monetary policy.[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 to seek to understand the mechanisms through which household heterogeneity may have an effect on aggregate fluctuations. Our model features idiosyncratic income

[1]For instance, Krussel and Smith (1998) study the role of income and wealth heterogeneity within a real business cycle model, and find that that the behavior of macroeconomic aggregates can be almost perfectly described using only the mean of the wealth distribution. See also Heathcote, Storesletten and Violante (2009), Guvenen (2011) and Krueger, Mitman and Perri (2016) for useful surveys of this earlier literature.

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

shocks as the only exogenous source of heterogeneity, in an environment where the only asset available is a riskless one-period bond and where borrowing constraints are not binding in equilibrium. Thus, and in order to isolate as much as possible the role of household heterogeneity, we abstract from two ingredients often featured in heterogenous agents (HA) models, namely: (i) multiple assets (including physical capital) with different degrees of liquidity and (ii) binding borrowing constraints.

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 resulting Euler equation in the HA economy includes two terms that are missing in its representative agent (RA) counterpart, which tend to have opposite effects on aggregate consumption dynamics. A first additional term captures variations in individual consumption *uncertainty*, averaged across households. Such a term is also present in the RA economy, but it is of second order relative to aggregate consumption variations and, hence, usually ignored. A second additional term captures changes in the *dispersion* of consumption across households. That term is trivially zero in the RA economy. We also show that these two channels are indeed related, and their combined effect can be captured by a cross-sectional consumption-weighted average of uncertainty changes, which can be viewed as a sufficient statistic for the impact of heterogeneity on aggregate consumption.

After deriving and discussing 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 idiosyncratic and aggregate endowment shocks. In that context, we study the mechanisms through which heterogeneity 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 role of heterogeneity in shaping the response of aggregate output to three aggregate shocks: a monetary policy shock, a preference/discount factor shock, and a technology shock. The simplicity of the models and the fact that the presence of idiosyncratic labor income shocks is the only departure from their RA counterparts allows us to isolate better their role in shaping aggregate fluctuations.

A central result of our analysis is to show how variations in *uncertainty* and *dispersion* may amplify or dampen the effects of aggregate shocks. In addition, in the context of the New Keynesian model, we discuss how these factors may exacerbate or mitigate the so-called forward-guidance puzzle (McKay, Nakamura and Steinsson (2016)).

We also show that variations in *uncertainty* and *dispersion* tend to have opposite effects on aggregate variables, with the net effect depending on the relative strength of these two channels. As an example, consider an aggregate shock that leads to a widespread and persistent increase in consumption *uncertainty*. That channel, by itself, would tend to reduce aggregate consumption, due to a precautionary savings motive. At the same time, to the extent that the largest reduction in consumption is concentrated among poorer (low-consumption) households, this would tend to offset the impact of precautionary savings on aggregate consumption (this is the *dispersion* channel).

From a quantitative viewpoint, we find that the net effect of these forces is very small in the two calibrated model economies that we analyze. By contrast, that impact becomes more significant when we introduce borrowing constraints that are binding for a non-negligible fraction of households at any point in time.

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 consumption uncertainty and dispersion, both from a qualitative and a quantitative perspective. Section 6 discusses the implications of introducing a binding borrowing constraint, and Section 7 concludes.

2 Related Literature

This paper belongs to a growing literature that studies the role of heterogeneity in aggregate economic fluctuations. In that literature, two main features are typically responsible for the differences in the behavior of aggregate variables relative to a representative agent economy: (i) uninsurable idiosyncratic income risk and (ii) the presence of binding borrowing constraints. However, understanding which is the exact role played by each of these factors remains a largely open question. This is what we seek to do in this paper.[3]

In this respect, our paper contributes to the literature developing *tractable* frameworks to isolate the channels through which heterogeneity operates. To that end, and following the original formulation of Campbell and Mankiw (1989), some studies in that literature (see e.g., Galí, López-Salido and Vallès (2007), Bilbiie (2008) 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.

Another branch of this literature (see e.g. Werning (2015), McKay et al. (2016), Bilbiie (2021), Ravn and Sterk (2021)) has considered economies with idiosyncratic income risk, but under assumptions that imply a degenerate wealth distribution in equilibrium.[4] That literature emphasizes the role played by the cyclicality of *income* inequality and liquidity for the propagation of aggregate shocks. We instead leave the equilibrium wealth distribution unrestricted, and emphasize the role played by the dynamics of *consumption* uncertainty for aggregate fluctuations, as well as the effectiveness of forward-guidance policies.

In related work, Acharya and Dogra (2020) consider an heterogenous household economy with CARA preferences, and no binding borrowing constraints, and where heterogeneity operates as a result of cyclical changes in the volatility of *income* shocks. We instead consider a framework with more standard CRRA preferences, associated with a non-trivial relationship between individual consumption, income and wealth. In that context, the dynamics of *consumption uncertainty* in response to aggregate shocks play a crucial role for the transmission of the latter, regardless of whether the volatility of income shocks is constant, or not.

We also show that the impact of heterogeneity can be summarized by a sufficient statistic, given by the cross-sectional consumption-weighted average of changes in uncertainty. In this respect, our paper is related to several qualitative and quan-

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

[4] For 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.

titative studies (see e.g. Auclert (2019), Auclert, Rognlie and Straub (2019) and Luetticke (2021) and the references therein) emphasizing the role of the cross-sectional distribution of variables other than consumption uncertainty (like the marginal propensity to consume, income, portfolio, etc.).

From a quantitative viewpoint, our finding that the net impact of heterogeneity on aggregate fluctuations is small is similar to that obtained by several authors in the literature, following Krusell and Smith (1998). In that respect, our contribution lies in the identification of two opposing forces.

3 Household Heterogeneity and the Euler Equation for Aggregate Consumption

Throughout we assume a continuum of households indexed by $j \in [0,1]$. Preferences are common to all households. They include a common discount factor $\beta \equiv \exp\{-\rho\}$ and a marginal utility of consumption given by $Z_t C_t(j)^{-\sigma}$, with $\sigma \geq 0$, and where $C_t(j)$ denotes household j's consumption in period t and $Z_t \equiv \exp\{z_t\}$ is a preference shock common to all households. We assume $\{z_t\}$ follows a stationary $AR(1)$ process $z_t = \rho_z z_{t-1} + \varepsilon_{z,t}$.

Our objective in this section is to derive an approximate Euler equation for aggregate consumption $C_t \equiv \int C_t(j)dj$. Our starting point is the individual consumption Euler equation

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

which is assumed to hold for all households $j \in [0,1]$ and for all t, where $R_t \equiv \exp\{r_t\}$ is the gross yield on a one-period riskless bond held between t and $t+1$.

We assume a common ergodic distribution for $C_t(j)$, and hence for $c_t(j) \equiv \log C_t(j)$.[5] A second order Taylor expansion of (1) around unconditional means $\Delta c_t(j) = z_t = 0$ yields:

$$c_t(j) \simeq \mathbb{E}_t\{c_{t+1}(j)\} - \frac{1}{\sigma}(r_t - \rho) + \frac{1}{\sigma}(1-\rho_z)z_t - \frac{\sigma}{2}v_t(j) \tag{2}$$

[5] See Ma et al. (2020) for a discussion of the conditions (satisfied in the models considered below) under which that ergodicity property holds.

where $v_t(j) \equiv \mathbb{E}_t\{\xi_{t+1}(j)^2\}$ with $\xi_t(j) \equiv c_t(j) - \mathbb{E}_{t-1}\{c_t(j)\}$. Note that $v_t(j)$ is a measure of uncertainty regarding one-period ahead individual consumption, with its (negative) effects on current consumption usually described as a "precautionary savings" motive.

Due to the presence of potentially "large" idiosyncratic income shocks in the background, we allow variations in $v_t(j)$ to be of the same order as fluctuations in z_t, r_t and other aggregate variables (henceforth denoted as $O(|y|)$).[6] This is in contrast with the representative household case, for which $v_t \equiv \mathbb{E}_t\{\xi_{t+1}^2\} \sim O(|y|^2)$, justifying its absence from common (approximate) representations of the consumption Euler equation in the literature.

Taking unconditional expectations:

$$r = \rho - \frac{\sigma^2}{2} v \tag{3}$$

where $r \equiv \mathbb{E}\{r_t\}$ and $v \equiv \mathbb{E}\{v_t(j)\}$. Thus we see that the interest rate in the stochastic steady state is decreasing in average consumption uncertainty.

Combining (2) and (3):

$$c_t(j) \simeq \mathbb{E}_t\{c_{t+1}(j)\} - \frac{1}{\sigma}\widehat{r}_t + \frac{1}{\sigma}(1-\rho_z)z_t - \frac{\sigma}{2}\widehat{v}_t(j) \tag{4}$$

where $\widehat{r}_t \equiv r_t - r$ and $\widehat{v}_t(j) \equiv v_t(j) - v$ are deviations from corresponding steady state values.

Integrating (4) over $j \in [0,1]$ and letting $\bar{c}_t \equiv \int c_t(j)dj$ denote *average log consumption* we obtain:

$$\bar{c}_t \simeq \mathbb{E}_t\{\bar{c}_{t+1}\} - \frac{1}{\sigma}\widehat{r}_t + \frac{1}{\sigma}(1-\rho_z)z_t - \frac{\sigma}{2}\widehat{v}_t \tag{5}$$

where $v_t \equiv \int v_t(j)dj$ is period t average consumption uncertainty (across households) and $\widehat{v}_t \equiv v_t - v$ is the corresponding deviation from its unconditional mean

[6]The reliance on a second-order approximation is without loss of generality, and only made for expositional purposes, as it facilitates the economic interpretation of the terms related to heterogeneity. Appendix A contains an equivalent formulation that does not rely on such approximation, and which is actually used for our quantitative exercises. Also, the assumption of "small" aggregate uncertainty justifies dropping additional terms involving the variance of innovations to z_{t+1} and their covariance with innovations to individual consumption, that would otherwise be part of the second order Taylor expansion of (1).

(over time).

At any point in time, *(log) aggregate consumption* $c_t \equiv \log \int C_t(j)dj$ will differ from *average log consumption* $\bar{c}_t$ due to Jensen's inequality. Let $x_t \equiv c_t - \bar{c}_t \geq 0$ denote the gap between the two variables. The size of that gap depends on the cross-sectional distribution of consumption, which is endogenous and time-varying. In a neighborhood of the symmetric allocation the following approximation holds $x_t \simeq \frac{1}{2}\int (c_t(j) - \bar{c}_t)^2 dj$, so we can think of x_t as capturing the cross-sectional dispersion of consumption at each point in time.[7] Henceforth we allow variations in x_t to be $O(|y|)$, i.e. the same order as aggregate variables and shocks, so that they cannot be ignored in our analysis. Now we can finally write the Euler equation for (log) aggregate consumption as:

$$c_t \simeq \underbrace{\mathbb{E}_t\left\{c_{t+1}\right\} - \frac{1}{\sigma}\widehat{r}_t + \frac{1}{\sigma}(1-\rho_z)z_t}_{\text{RA model}} \underbrace{-\frac{\sigma}{2}\widehat{v}_t - \mathbb{E}_t\left\{\Delta x_{t+1}\right\}}_{\text{HA component}}. \tag{6}$$

Under the assumption that $\lim_{T\to\infty}\mathbb{E}_t\left\{c_{t+T}\right\} = c$ and $\lim_{T\to\infty}\mathbb{E}_t\left\{x_{t+T}\right\} = x$ we can iterate (6) forward to obtain:

$$\widehat{c}_t = \underbrace{-\frac{1}{\sigma}\sum_{k=0}^{\infty}\mathbb{E}_t\left\{\widehat{r}_{t+k}\right\} + \frac{1}{\sigma}z_t}_{\text{RA model}} \underbrace{-\frac{\sigma}{2}\sum_{k=0}^{\infty}\mathbb{E}_t\left\{\widehat{v}_{t+k}\right\} + x_t}_{\text{HA component}} \tag{7}$$

Note that in the particular case of a representative household (i.e. the RA model) we have $x_t = 0$ (no cross-sectional dispersion) and $v_t \equiv \mathbb{E}_t\{\zeta_{t+1}^2\} \simeq 0$ ("small" aggregate uncertainty). In that case (6) reduces to the familiar Euler equation, with aggregate consumption fluctuations being determined by (current or expected) changes in the interest rate r_t and preference shocks z_t (the first two terms in (7))

In contrast, in the presence of idiosyncratic income shocks and incomplete markets, heterogeneity impacts aggregate consumption through two additional channels, summarized by the HA component in eq. (7): (i) current and anticipated

[7] The derivation is straightforward. Using the definition of aggregate consumption: $1 \equiv \int (C_t(j)/C_t)dj = \int \exp\{c_t(j) - c_t\}dj =\simeq \int \left(1 + (c_t(j) - c_t) + \frac{1}{2}(c_t(j) - c_t)^2\right) dj$. Letting $\bar{c}_t \equiv \int c_t(j)dj$ and rearranging terms gives $\bar{c}_t \simeq c_t - \frac{1}{2}\int (c_t(j) - c_t)^2 dj$. The previous approximation is meant to facilitate interpretation, but is not used in the solution of the models considered below.

fluctuations in average consumption uncertainty, captured by movements in $\widehat{v}_t$, and (ii) variations in consumption dispersion, captured by movements in x_t. In particular, we see that an increase in average uncertainty v_t, current or anticipated, tends to lower current consumption. This is the *precautionary savings channel* which, together with current and anticipated changes in the real interest and the preference shock, determines average log consumption $\bar{c}_t$. On the other hand, a higher consumption dispersion x_t is associated with a higher log aggregate consumption c_t, given average log consumption $\bar{c}_t$. We refer to this effect as the *dispersion channel*.

Our main goal is to understand how the presence of heterogeneity affects the response of aggregate consumption to an aggregate shock. With that purpose in mind let us define

$$\widehat{c}_t^H = -\frac{\sigma}{2}\sum_{k=0}^{\infty}\mathbb{E}_t\left\{\widehat{v}_{t+k}\right\} + x_t$$

which is the deviation between (log) aggregate consumption and its counterpart in the RA model, *conditional on the same path for the real rate*. Henceforth we refer to c_t^H as the *consumption heterogeneity factor*. The dynamic response of $\{c_{t+k}^H\}$ to a period t aggregate shock, ε_t, is thus given by:

$$[\partial c_{t+k}^H/\partial\varepsilon_t] = -\frac{\sigma}{2}\sum_{h=0}^{\infty}[\partial v_{t+k+h}/\partial\varepsilon_t] + [\partial x_{t+k}/\partial\varepsilon_t] \tag{8}$$

for $k = 0, 1, 2, ...$, i.e. heterogeneity affects the response of aggregate consumption to a an aggregate shock through the effect of the latter on average consumption uncertainty and consumption dispersion, for any given response of the real rate. Next we show that the two components of that response, uncertainty and dispersion, are closely related, with a likely dominant role for the response of average uncertainty. In addition we derive a sufficient statistic for the impact of heterogeneity on the response of aggregate consumption to an aggregate shock.

3.1 The Relationship between Dispersion and Uncertainty

To lay the ground for our derivation below, note that one can combine (4) and (5) to yield the following difference equation for $\widetilde{c}_t(j) \equiv c_t(j) - \bar{c}_t$, the deviation of individual log consumption from its average counterpart:

$$\widetilde{c}_t(j) = \mathbb{E}_t\{\widetilde{c}_{t+1}(j)\} - \frac{\sigma}{2}\widetilde{v}_t(j)$$

where $\widetilde{v}_t(j) \equiv v_t(j) - v_t$. Under our ergodicity assumption , $\lim_{T\to\infty} \mathbb{E}_t\{\widetilde{c}_{t+T}(j)\} = 0$, thus we can iterate forward the previous equation to yield:

$$\widetilde{c}_t(j) = -\frac{\sigma}{2}\sum_{k=0}^{\infty} \mathbb{E}_t\{\widetilde{v}_{t+k}(j)\} \tag{9}$$

i.e. deviations of individual (log) consumption from the mean differentials are inversely related to current and anticipated differentials in consumption uncertainty. Note that both real interest rates and preference shocks are absent from (9), since they affect all households identically.

Using the approximation $x_t \simeq \frac{1}{2}\int(c_t(j) - \overline{c}_t)^2 dj$ we can derive an expression for the impact of the dispersion channel on aggregate consumption:

$$\begin{aligned}
[\partial x_{t+k}/\partial\varepsilon_t] &\simeq \int \widetilde{c}_{t+k}(j)\ [\partial\widetilde{c}_{t+k}(j)/\partial\varepsilon_t]\ dj \\
&= -\frac{\sigma}{2}\int \widetilde{c}_{t+k}(j) \sum_{h=0}^{\infty}[\partial v_{t+k+h}(j)/\partial\varepsilon_t]\ dj \qquad (10)
\end{aligned}$$

where the second equality follows from (9) and the fact that $\int \widetilde{c}_{t+k}(j)dj = 0$. Thus, we see that the dispersion channel will amplify or dampen the impact of the precautionary savings channel depending on the correlation between the cumulative change in individual uncertainty with relative consumption. If that correlation is negative the dispersion channel will tend to offset the precautionary savings channel, capturing the fact that the latter will likely have a larger impact on low consumption households, thus limiting its aggregate effect.

Combining (8) and (10) we obtain

$$\begin{aligned}
[\partial c^H_{t+k}/\partial\varepsilon_t] &= -\frac{\sigma}{2}\sum_{h=0}^{\infty}[\partial v_{t+k+h}/\partial\varepsilon_t] - \frac{\sigma}{2}\int \widetilde{c}_{t+k}(j)\sum_{h=0}^{\infty}[\partial v_{t+k+h}(j)/\partial\varepsilon_t]\ dj \\
&= -\frac{\sigma}{2}\int(1+\widetilde{c}_{t+k}(j))\sum_{h=0}^{\infty}[\partial v_{t+k+h}(j)/\partial\varepsilon_t]\ dj \\
&\simeq -\frac{\sigma}{2}\int \frac{C_{t+k}(j)}{C_{t+k}}\sum_{h=0}^{\infty}[\partial v_{t+k+h}(j)/\partial\varepsilon_t]\ dj. \qquad (11)
\end{aligned}$$

The previous expression shows that the impact of heterogeneity on the response of aggregate consumption to an aggregate shock is inversely related to a consumption-weighted average of the cumulative increase in individual consumption uncertainty. Thus, we can think of tha latter as a *sufficient statistic* for the impact of heterogeneity on aggregate consumption fluctuations. Intuitively, a given *percent* change in the consumption of an individual household (as determined by $\sum_{h=0}^{\infty}[\partial v_{t+k+h}(j)/\partial \varepsilon_t]$) will have a larger (smaller) effect on aggregate consumption the higher (lower) is the relative consumption of that household. Note that if the sign of $\sum_{h=0}^{\infty}[\partial v_{t+k+h}(j)/\partial \varepsilon_t]$ is the same for all (or most) j, the sign of $[\partial c_{t+k}^{H}/\partial \varepsilon_t]$ will correspond to that of $\sum_{h=0}^{\infty}[\partial v_{t+k+h}/\partial \varepsilon_t]$, implying that the precautionary savings effect will likely dominate the dispersion effect.[8] This is the case in all the simulations of calibrated example economies reported and discussed below. Also, in those examples the net effect is generally small, so that the aggregate responses to shocks are similar to those obtained in the corresponding model with a representative agent.

Equation (11) clarifies that if individual consumption uncertainty remains unchanged in response to an aggregate shock, i.e. $\sum_{h=0}^{\infty}[\partial v_{t+k+h}/\partial \varepsilon_t] = 0 \ \forall j \in [0,1]$, then heterogeneity will have no impact on the behavior of aggregate consumption.

To fix ideas, it is useful to consider two examples where uncertainty is not affected by aggregate shocks. A first example is an economy where (log) individual consumption is a linear function of income (or cash on hand). In that case, the sensitivity of consumption to idiosyncratic shocks would be constant, and independent from aggregate shocks.[9] In this case, the response of aggregate consumption to aggregate shocks would be the same as in the corresponding RA economy.

A second example is an economy with zero-liquidity, which as mentioned in Section 2 is a limiting case often considered in the literature. In that case, since there is no borrowing or saving in equilibrium, individual consumption must coincide with individual income, in all periods. As a consequence, aggregate shocks have no effect on the households' ability to insure against (future) idiosyncratic

[8] Overturning that result would require that wealthy households experience a change in uncertainty sufficiently large and of a sign different from the change in average uncertainty. The presence of wealthy hand-to-mouth households, as in Kaplan and Violante (2014), constitutes a potential source of that phenomenon.

[9] This would be the case, for instance, under with linear-quadratic preferences, assuming a constant real interest rate. More generally, individual uncertainty is related to the coefficient of "prudence", see Kimball (1990).

shocks, so that the sufficient statistic in eq. (11) equals zero. The only remaining channel through which heterogeneity operates in that economy would be due to the presence of constrained households, and the associated risk that currently unconstrained households may become constrained at a future date —a channel that have purposefully ignored in our analysis.

3.2 A Brief Detour: Understanding Variations in Consumption Uncertainty

The discussion above has made clear the importance of variations in consumption uncertainty, current and anticipated, in shaping aggregate fluctuations in economies with household heterogeneity. In the present section we try to dig further in order to understand the sources of those variations.

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

$$c_t(j) = \mathcal{C}(s_t(j), S_t) \tag{12}$$

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 regarding the distribution of household-specific and economy-wide variables (wages, interest rates, etc) whose current and future values are relevant for today's consumption decision. The existence and properties of a consumption function like (12) can be established under standard assumptions.

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

$$\begin{aligned} \xi_t(j) &\equiv c_t(j) - \mathbb{E}_{t-1}\{c_t(j)\} \\ &= f^j_{t-1}(\nu_t(j), \varepsilon_t) \end{aligned} \tag{13}$$

where $f^j_{t-1}(\cdot)$ is a function satisfying $f^j_{t-1}(0,0) = 0$. In what follows, and in order to keep the algebra simple, we assume $\nu_t(j)$ and ε_t are scalars. See Appendix B for a generalization.

Under our assumptions, and using (13), we can approximate individual uncertainty $v_t(j)$ in period t as

$$\begin{aligned} v_t(j) &\equiv \mathbb{E}_t\{\xi_{t+1}(j)^2\} \\ &\simeq \psi_t(j)^2\sigma_\nu^2 + \varphi_t(j)^2\sigma_\varepsilon^2 \end{aligned}$$

where $\psi_t(j) \equiv \partial f_t^j(0,0)/\partial\nu_{t+1}(j)$, 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 $\sigma_\nu^2 \equiv \mathbb{E}\{\nu_t(j)^2\}$ for all $j \in [0,1]$ and $\sigma_\varepsilon^2 \equiv \mathbb{E}\{\varepsilon_t^2\}$ are, respectively, the variances of those shocks. Aggregating across households

$$v_t = \sigma_\nu^2 \int \psi_t(j)^2 dj + \sigma_\varepsilon^2 \int \varphi_t(j)^2 dj.$$

As shown in appendix B (for the general case with multiple shocks) the term $\sigma_\varepsilon^2 \int \varphi_t(j)^2 dj \sim O(|y|^2)$ if we are willing to impose bounds on the extent of heterogeneity in the elasticity $\varphi_t(j)$. Accordingly, we can write:

$$\begin{aligned} v_t &\simeq \sigma_\nu^2 \int \psi_t(j)^2 dj \\ &= \sigma_\nu^2 \left(\overline{\psi}_t^2 + var_j\{\psi_t(j)\}\right) \end{aligned} \tag{14}$$

where $\overline{\psi}_t \equiv \int \psi_t(j)dj$ and $var_j\{\psi_t(j)\} \equiv \int (\psi_t(j) - \overline{\psi}_t)^2 dj$ are the cross-sectional mean and variance of individual consumption elasticities *with respect to the idiosyncratic shock*.

Note that both statistics will generally vary in response to aggregate shocks, as a result of the latter's impact on state variables. Below, we study the response of v_t, and the factors behind it $\overline{\psi}_t$ and $var_j\{\psi_t(j)\}$ to different aggregate shocks in the context of two calibrated model economies.

Throughout our analysis, we maintain the assumption that the variance of idiosyncratic income shocks (σ_ν^2) is constant over time —i.e. the idiosyncratic *income* risk is acyclical. In principle, cyclical income risk could be another channel through which heterogeneity affects aggregate consumption. In related work, Acharya and Dogra (2020) consider a model with heterogeneous households under CARA preferences and no binding borrowing limit, and show that in that case the innovation in (the level of) consumption is a linear function of the idiosyncratic income shock, with a small coefficient of proportionality (approximatively equal to $r/(1+r)$), which is constant across households (i.e. the cross-sectional variance of the elasticity of consumption equals zero). In that setting, heterogeneity mainly operates through the cyclicality of income risk, which is however found to have a small quantitative impact on the behavior of aggregate variables. These findings can be rationalized through the lens of our analysis. Indeed, consider the particular case where the mean elasticity of consumption $\overline{\psi}_t = r/(1+r)$ is small —say of order $O(|y|)$ for plausible settings of r— and its cross-sectional variance $var_j\{\psi_t(j)\} = 0$. From eq. (14), average uncertainty is given by $v_t = \sigma_v^2 \overline{\psi}_t^2$, which means that changes in v_t are likely to be small (i.e. of order $O(|y|^2)$ or smaller), regardless of whether the variance of idiosyncratic income shocks σ_v is constant, or it fluctuates over time.[10]

4 The Role of Heterogeneity 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:

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

for $t = 0, 1, 2...$, where $B_t(j)$ represents holdings of one-period bonds, which yield a gross riskless real return R_t. The endowment, $Y_t \exp\{\eta_t(j)\}$, has two components: an aggregate component Y_t, which is common to all households, and follows an $AR(1)$ process (in logs) with autocorrelation $\rho_y \in [0,1)$; and an idiosyncratic component $\eta_t(j)$, which follows a stationary K-state Markov process with $\eta_t(j) \in [\eta_1, .., \eta_K]$, with $\eta_1 < \eta_2 < ... < \eta_K$, and where $\int_0^1 \exp\{\eta_t(j)\} dj = 1$. Note that by setting $\eta_t(j) = 0$ for all $j \in [0,1]$ and for 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.

We impose a borrowing constraint of the form

$$B_t(j) \geq \underline{B} \tag{15}$$

[10]See also Bayer et. al. (2019) and Ravn and Sterk (2020), among others, for examples of heterogenous household economies with cyclical idiosyncratic risk.

for all t. For most of the analysis below, we set $\underline{B} = -Y \exp\{\eta_1\}/r$, which constitutes the "natural debt limit," given a constant aggregate output and interest rate at their the steady state values (Y, r). The desire to avoid zero consumption (given that $\lim_{c \to 0} U_c = +\infty$) 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.[11]

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:

$$\widehat{r}_t = \underbrace{-\sigma(1-\rho_y)y_t}_{\text{RA model}} \underbrace{-\frac{\sigma^2}{2}\widehat{v}_t - \sigma \mathbb{E}_t\left\{\Delta x_{t+1}\right\}}_{\text{HA component}}. \tag{16}$$

The first term in (16) is the equilibrium real rate in the corresponding RA 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 heterogeneity on the interest rate is captured by the remaining two terms. On the one hand, an increase in average consumption uncertainty ($\hat{v}_t$) tends to increase the demand for precautionary savings, leading to a reduction in the equilibrium interest rate. On the other hand, an increase in $\mathbb{E}_t\left\{\Delta x_{t+1}\right\}$ occurs whenever the change in (log) average consumption is smaller than the change in aggregate (log) consumption —which is exogenous in this setting. This effect would tend to lower the equilibrium interest rate, relatively to a representative agent economy. To see why this is the case, consider the effects of a temporary increase (decrease) in aggregate endowment, so that households anticipate a decline (increase) in their future consumption. An increase in $\mathbb{E}_t\left\{\Delta x_{t+1}\right\}$ means that, on average, households expect a larger decline (smaller increase) in their (log) consumption, relative to the the decline (increase) in (log) aggregate consumption. The associated increase (or smaller decline) in the demand for savings leads to a lower equilibrium interest rate.

In summary, equation (16) implies that the role of heterogeneity for the response of the real interest rate to an aggregate endowment shock is a function of the

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

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_η	Autocorr. of idiosyncratic earnings	0.966
σ_η	Std. dev. of idiosyncratic earnings	0.5
Discretization		
n_η	Points in Markov Chain for η	11
n_a	Points in Markov Chain for Asssets	500

response of average uncertainty and dispersion to that shock. Next we turn to a quantitative assessment of the impact of an aggregate endowment shock on the real interest rate 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 an equilibrium steady-state real risk-free rate of 2 percent (in annual terms).

Regarding the idiosyncratic endowment shock, we assume that it follows an AR(1) process $\eta_t(j) = \rho_\eta \eta_{t-1}(j) + \nu_t(j)$, where $\nu_t(j) \sim N(0, \sigma_\eta \sqrt{1-\rho_\eta^2})$. Following Auclert et. al. (2020), we set $\rho_\eta = 0.966$ and $\sigma_\eta = 0.5$. This AR(1) process is then discretized into an 11 state Markov process according to the Rouwenhorst approach.[12] 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 the lower bound (the natural debt

[12] As a robustness check, Online Appendix A considers an alternative income process which combines a transitory and persistent component, and is a discrete-time (quarterly) version of the continous-time process in Kaplan, Moll and Violante (2018).

limit) and an upper bound set to 300 times quarterly income. For given values of the real interest rate and aggregate income, we solve for the households' policy functions using the endogenous gridpoints method described in 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 assets distribution implied by the households' choices satisfies the market clearing condition $\int B_t(j)dj = 0$.

For the transition dynamics, we adopt the sequence-state Jacobian approach described in Auclert et. al. (2020). This amounts to find 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).[13] 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) of aggregate shocks implies that idiosyncratic shocks are the only source of individual (and aggregate) uncertainty. As shown above, such uncertainty is a key determinant of the role of heterogeneity for aggregate consumption. In all our numerical exercises, in order to accurately capture the quantitative role of heterogeneity and its components, we do not rely on the second-order approximation to v_t and x_t described in section 1, but instead on the the exact representation contained in Appendix A.

4.2 Findings

We focus our discussion on the dynamic response of the real interest rate to an aggregate endowment shock. Figure 1 shows the responses of the real interest rate and (log) aggregate endowment to a 1 percent positive shock in the latter variable. The response of the real interest rate (left panel) is plotted for both our baseline model with heterogeneity (red line with circles) and for the corresponding

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

Figure 1: The Effects of an Aggregate Endowment Shock

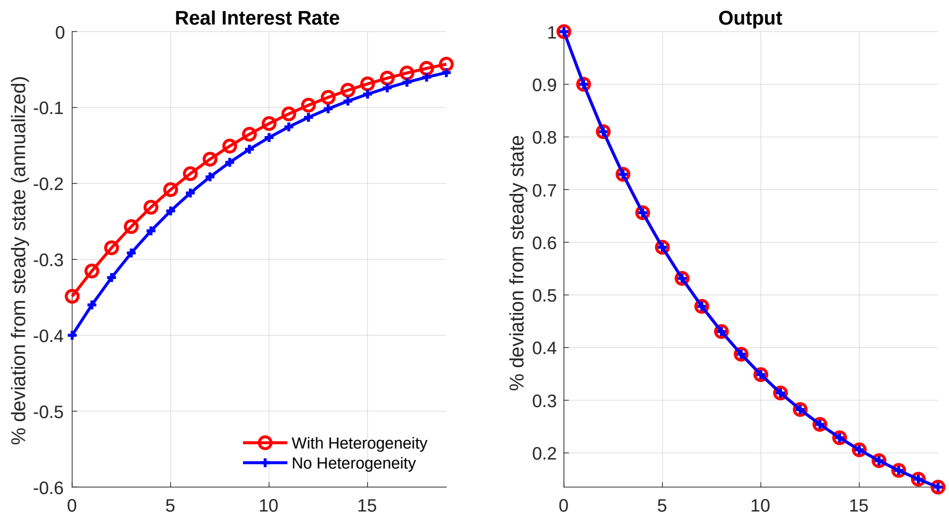

Notes: The figure shows the response of the annualized sreal interest rate (left panel) to a positive aggregate endowment shock (right panel) in in a model with heterogeneity (red line with circles) and without heterogeneity (blue line with crosses).

representative agent model (blue line with crosses). As it can be seen in the Figure, the two impulse responses are close to each other (at most 5 basis points in annual terms), thus suggesting that heterogeneity has a minor impact on the response of the real rate to an aggregate shock.

The role of heterogeneity is further investigated in Figure 2 through the lenses of our theoretical analysis —see eq. (22). The black line on the top panel displays the gap between the real rate responses shown in Figure 1. That evolution of that gap is the sum of two components, namely, $-\frac{\sigma^2}{2} dv_{t+k}$ (orange line) and $-\sigma(dx_{t+k+1} - dx_{t+k})$ (green line), respectively associated to the responses of average uncertainty and consumption dispersion.

Note that the overall effect of heterogeneity is positive, i.e. it dampens the decline in the interest rate, but quantitatively small. More interestingly, the figure shows 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 individual uncertainty v_t (bottom-left panel), which lowers the demand for savings and tends to increase the interest rate. On the other hand, the same shock causes consumption

Figure 2: Heterogeneity in an Endowment Economy

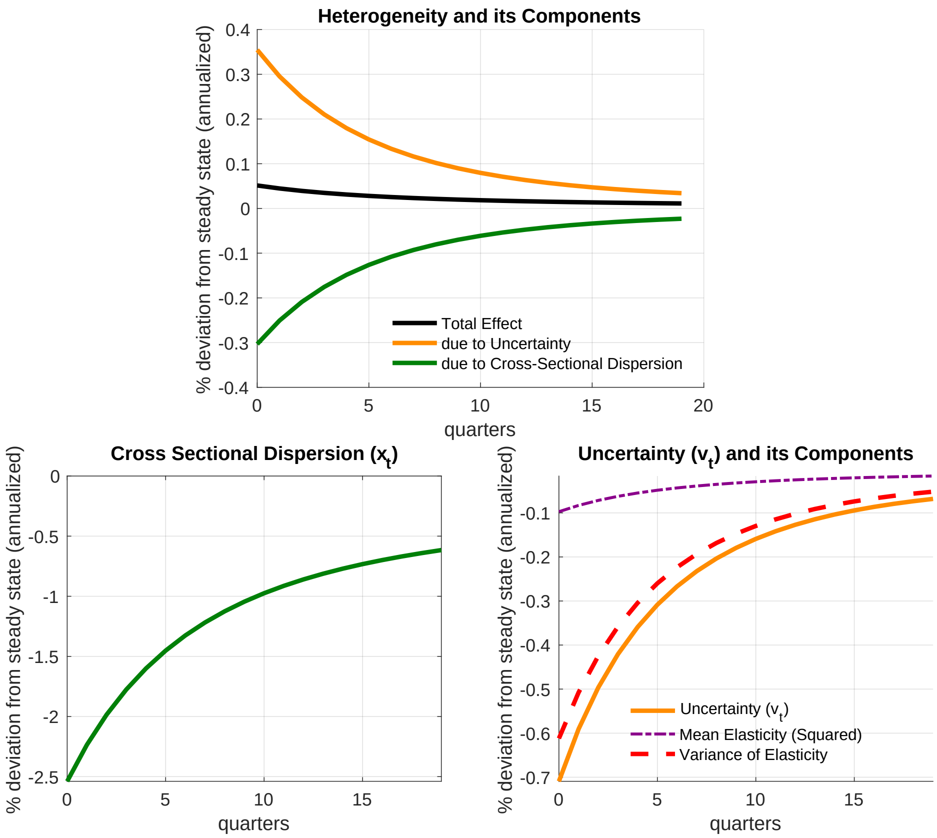

Notes: The figure shows the effect of heterogeneity on the response of the real interest rate to a positive aggregate endowment shock. The top panel shows the total effect (black line), together with the components due to uncertainty $-\frac{\sigma^2}{2}\widehat{v}_t$ (orange line) and cross-sectional dispersion $-\sigma\mathbb{E}_t\left\{\Delta x_{t+1}\right\}$ (green line). The bottom-left panel shows the response of dispersion x_t, and the bottom-right panel shows the response of uncertainty v_t (orange line), together with the responses of the variance of the elasticity (red dashed line) and of the (square of) mean elasticity (purple dash-dotted line). All figures are annualized.

dispersion x_t to decline on impact (bottom-right panel), after which it is expected to increase. The latter effect leads to a reduction in average individual consumption growth, for any given growth rate of aggregate consumption, which has a negative effect on the equilibrium interest rate. From a qualitative viewpoint, the effect of uncertainty dominates the effects of dispersion. From a quantitative viewpoint the impact of *each* of these two forces separately is similar to the overall change in the interest rate —which is about 35 basis points on impact. When considered jointly, however, their impact is an order of magnitude smaller.[14]

As mentioned in section 3.2, the behavior of average consumption uncertainty is related to the distribution of the elasticity of consumption with respect to the idiosyncratic shock. The bottom-right panel in Figure 2 shows that the decline in average consumption uncertainty discussed earlier is associated with a decline in both the cross-sectional mean $\overline{\psi}_t$ and in the cross-sectional variance $var_j\{\psi_t(j)\}$ of the consumption elasticity —with the decline in the variance being quantitatively more relevant.

The reason behind the decline in the cross-sectional mean and variance of the elasticity of consumption can be understood looking at Figure 3, which shows the steady state relationship between (log) consumption and the the elasticity of consumption $\psi_t(j)$.[15] As it can be seen, 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. This explains why an increase in aggregate income, which in and of itself causes an increase in consumption for most households, leads to a decline in $\overline{\psi}_t$, the mean elasticity of consumption. At the same time, the figure shows that the relationship between consumption and the elasticity of consumption is convex. Intuitively, the elasticity of consumption varies substantially as households get closer to their natural debt limit, but it is roughly constant (and small) for households with high income and wealth, who behave almost as permanent-income consumers. This explains why an increase in income also leads

[14]This 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.

[15]More precisely, the figure displays the range of elasticities $\varphi_t(j)$ 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.

Figure 3: Elasticity of Consumption in Steady State

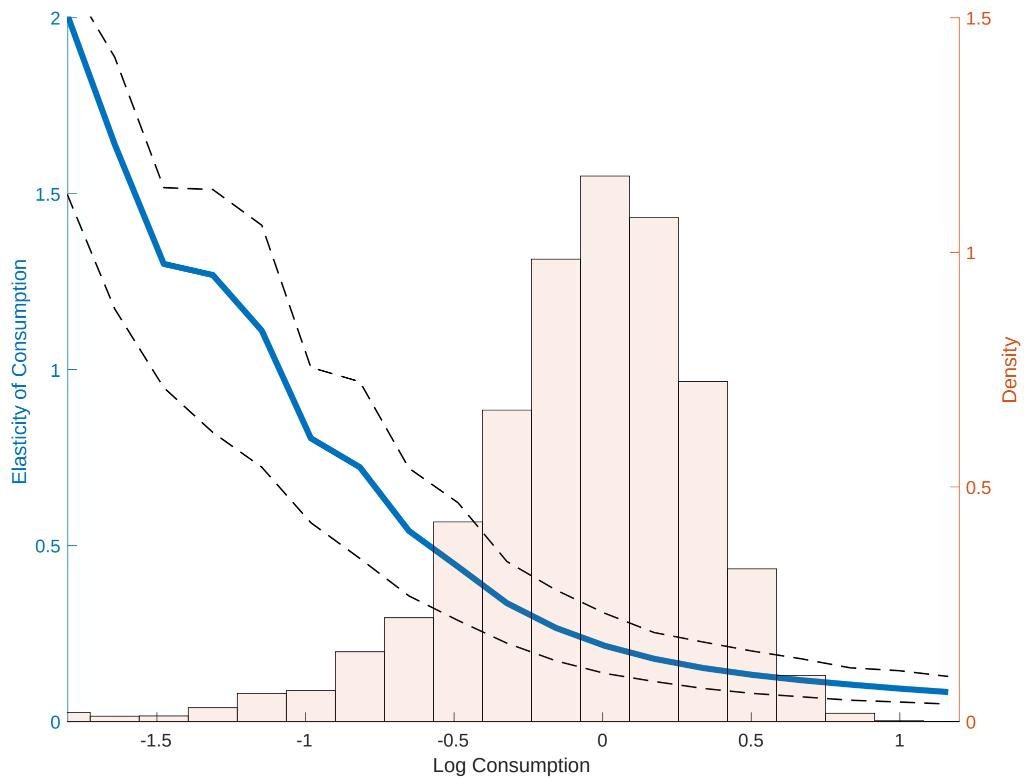

Notes: The figure shows the relationship between log consumption (horizontal axis), and the elasticity of consumption (left vertical axis) in steady state. For each value of consumption, the figure reports the average elasticity (solid blue line), the 5 95% interval of the distribution (black dashed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

to a reduction in $var_j\{\psi_t(j)\}$.

The relationship between the consumption elasticity $\psi_t(j)$ and consumption can also account for the initial reduction in consumption dispersion in response to the positive endowmment shock, since low-consumption households will on average increase their consumption more than proportionally than high-consumption households.

5 The Role of Heterogeneity in a New Keynesian Economy

We analyze an economy populated by a continuum of households, indexed by $j \in [0,1]$, with identical preferences given by $\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t(j), \mathcal{N}_t(j); Z_t)$. 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. $\mathcal{N}_t(j)$ denotes work hours. $Z_t \equiv \exp\{z_t\}$ is an exogenous preference shifter, common to all households. We assume $U(C, \mathcal{N}; Z) = \left(\frac{C^{1-\sigma}-1}{1-\sigma} - \frac{\mathcal{N}^{1+\varphi}}{1+\varphi} \right) Z$.

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:

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

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), $D_t(j)$ are real dividends, and $\eta_t(j)$ is an idiosyncratic productivity shifter which follows a stationary K-state Markov process identical to the one assumed in the previous section.[16] 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.

[16] The 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 effects on households, depending on their individual net wealth positions.

The supply side of the economy is kept as simple as possible, so that it is not affected by the presence of household heterogeneity. This allows us to focus on the impact of the latter on aggregate demand (which coincides with aggregate consumption in our simple model), in the spirit of Werning (2015). In particular, we assume a wage schedule

$$W_t = \mathcal{M}_w C_t^{\sigma} N_t^{\varphi} \tag{17}$$

where $C_t \equiv \int_0^1 C_t(j)dj$ denotes average consumption and where $\mathcal{M}_w > 1$ is a constant (gross) average wage markup. We implicitly assume $W_t \exp\{\eta_t(j)\} \geq C_t(j)^{\sigma} \mathcal{N}_t(j)^{\varphi}$ 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).

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

$$Y_t(i) = A_t N_t(i) \tag{18}$$

where $A_t \equiv \exp\{a_t\}$ 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:

$$\Pi_t (\Pi_t - 1) = \mathbb{E}_t \left\{ \Lambda_{t,t+1} \left(\frac{Y_{t+1}}{Y_t} \right) \Pi_{t+1} (\Pi_{t+1} - 1) \right\} + \frac{\epsilon}{\xi} \left(\frac{W_t (1-\tau)}{A_t} - \frac{1}{\mathcal{M}_p} \right) \tag{19}$$

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.[17] Aggregate profits are then given by $D_t = Y_t \Delta^p(\Pi_t) - W_t N_t$ where $\Delta^p(\Pi_t) \equiv 1 - (\xi/2)(\Pi_t - 1)^2$.

Combining eqs. (17)-(19), and taking a first-order approximation around the

[17] Formally, the subsidy is chosen such that $\mathcal{M}^p \mathcal{M}^w (1-\tau) = 1$.

zero-inflation steady state gives the well-know New Keynesian Phillips curve

$$\pi_t = \beta \mathbb{E}_t\{\pi_{t+1}\} + \kappa \tilde{y}_t \tag{20}$$

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 $y_t^n \equiv a_t(1+\varphi)/(\sigma+\varphi)$. Note that the latter is independent from monetary policy and, most importantly, it is also unaffected by heterogeneity factors.[18]

Regarding monetary policy, we assume the central bank controls directly the real interest rate, i.e. that $\widehat{r}_t$ follows an exogenous $AR(1)$ process $\widehat{r}_t = \rho_r \widehat{r}_{t-1} + \varepsilon_t^m$, where $\mathbb{E}_t\{\varepsilon_{t+1}^m\} = 0$. This specification allows us to isolate the (direct) effects of heterogeneity on aggregate demand, abstracting from the potential (indirect) effects due to the differential behavior of aggregate variables, which in turn may lead to a differential endogenous monetary policy response. [19]

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

$$C_t = Y_t \Delta^p(\Pi_t) \tag{21}$$

Market clearing in the bonds markets implies that $\int_0^1 B_t(j)dj = 0$ for all t. Finally, aggregate employment is given by $N_t = \int_0^1 N_t(i)di = Y_t/A_t$, which is assumed to be allocated uniformly across households, i.e. $\mathcal{N}_t(j) = N_t$ for all $j \in [0,1]$.

Up to a first-order approximation and in a neighborhood of the zero inflation steady state (21) can be written as

$$c_t = y_t$$

Combining the previous condition with Euler equation for aggregate consumption derived in section 2, we obtain a version of the dynamic IS equation:

$$y_t = \mathbb{E}_t\{y_{t+1}\} - \frac{1}{\sigma}\widehat{r}_t + \frac{1}{\sigma}(1-\rho_z)z_t - \frac{\sigma}{2}\widehat{v}_t - \mathbb{E}_t\{\Delta x_{t+1}\}$$

[18] Note also that the steady state output is given by $y = 0$.

[19] In Online Appendix B, we also consider a case where the central bank follows a Taylor-type rule for the real interest rate, and show that our main conclusions remain unaltered.

Iterating forward the previous condition and imposing $\lim_{T\to\infty} \mathbb{E}_t \{y_{t+T}\} = y$ and $\lim_{T\to\infty} \mathbb{E}_t \{x_{t+T}\} = x$ we obtain the following expression for (log) aggregate output

$$\widehat{y}_t = \underbrace{-\frac{1}{\sigma}\sum_{k=0}^{\infty} \mathbb{E}_t\{\widehat{r}_{t+k}\} + \frac{1}{\sigma}z_t}_{\text{RA model}} \quad \underbrace{-\frac{\sigma}{2}\sum_{k=0}^{\infty} \mathbb{E}_t\{\widehat{v}_{t+k}\} + \widehat{x}_t}_{\text{HA component}} \tag{22}$$

The first two terms in the previous expression correspond to equilibrium output in the RA version of the New Keynesian model. The second two terms reflect instead the impact of heterogeneity on equilibrium output, which is decreasing in current and anticipated average uncertainty —through its effects on precautionary savings— and increasing in the degree of current dispersion —since the higher is the dispersion, the higher must be the (log of) aggregate consumption associated with any given value of average (log) consumption.

Thus, the impact of heterogeneity on the response of aggregate output to a generic aggregate shock works through the effect of the latter on (i) current and anticipated average consumption uncertainty, which influences the response of average (log) consumption and on (ii) the cross-sectional consumption dispersion, which determines (log) aggregate consumption for any given average (log) consumption.

For illustrative purposes, let's consider the effects of an exogenous change in the real interest rate, and assume that interest rates and uncertainty are positively correlated, i.e. $\frac{\partial v_{t+h}(j)}{\partial r_{t+k}} > 0$ for all (or most) j. Such an assumption can be justified considering that a reduction (increase) in interest rates tends to improve (worsen) the individual ability to insure against idiosyncratic shocks —and in fact this is what happens in our quantitative examples.

Under that assumption, Eq. (22) implies that if the uncertainty channel (i) is the dominant factor, heterogeneity would amplify the effects of interest rate shocks. Otherwise, if the dispersion channel (ii) dominates the effects of interest rate changes would be dampened. As we argued in Section 3.1, we expect the uncertainty channel to be the dominant one in most applications.

Several studies in the literature have argued that the presence of household heterogeneity and idiosyncratic income risk may provide an explanation for the so-called forward-guidance puzzle, and namely that typical RA models tend to overestimate the effects of changes in future interest rates on current output. In this respect, Eq. (22) highlights that the puzzle could be exacerbated or mitigated, depending on whether the uncertainty or the dispersion channel dominate.

Indeed, it follows from equation (11) that the effect of heterogeneity for the impact of a "forward guidance" shock at a generic period $t+k$ is given by

$$\frac{\partial c_t^H}{\partial r_{t+k}} - \frac{\partial c_t^H}{\partial r_t} = -\frac{\sigma}{2} \int \frac{C_t(j)}{C_t} \sum_{h=0}^{k-1} \mathbb{E}_t \frac{\partial v_{t+h}(j)}{\partial r_{t+k}} dj.$$

The previous expression implies that, under the maintained assumption that $\frac{\partial v_{t+h}(j)}{\partial r_{t+k}} > 0$, the effects of forward guidance shocks are amplified, and thus the forward-guidance puzzle is exacerbated. Otherwise, but more unlikely, if the dispersion channel dominates, heterogeneity dampens the effects of forward-guidance shocks, and the puzzle is mitigated. A related point is made by Bilbiie (2021), who derives an analogous result in a two-agent economy with no idiosyncratic income shocks and with binding borrowing constraints for the subset of hand-to-mouth households.

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 $\eta_t(j)$. In addition, we set the (inverse) Frisch elasticity of substitution to unity ($\varphi = 1$).[20] 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 ξ 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 $\rho_z = \rho_r = 0.5$. We adopt the same numerical solution method described in Section 4.1.

[20]This calibration, together with the assumption of a steady state subsidy that corrects the distortions due to households' and firms' monopolistic power, implies that in steady state aggregate labor income $WN = 1$, and thus the steady state of the model is identical to the one considered for the endowment economy.

5.2 Findings

We now turn the attention to our baseline New Keynesian economy, and analyze how heterogeneity affects the response of aggregate variables to three exogenous shocks: monetary policy, preferences and technology. We discuss each in turn.

Figure 4 illustrates the response of aggregate variables to a 25 basis point expansionary monetary shock, which leads to a 100 basis point reduction in the (annualized) real interest rate. The figure compares the responses in a model with heterogeneity (red line with circles) with their representative agent counterparts (blue line with crosses). Regarding the response of output (top-left panel), the presence of heterogeneity tends to amplify the aggregate effects of a monetary shock. The effects are stronger on impact, and more persistent. However, from a quantitative viewpoint, the magnitude of this amplification seems relatively small —less that 0.05 percentage points at all horizons. The high persistence of the gap between the two output responses, combined with the extreme forward-lookingness of the model's Phillips curve, implies a slightly larger response of inflation in the model with heterogeneity (top-right panel).

Figure 5 analyzes the mechanisms behind the larger response of output in the presence of heterogeneity through the lens of the dynamic IS equation (22) derived above. The top panel shows the differential impulse response (black line) and its components, associated with changes in uncertainty (orange line) and dispersion (green line). An expansionary monetary shock leads to a persistent reduction in average uncertainty (bottom-right panel) which, other things equal, leads to an increase in current consumption, thus amplifying the original effects of the shock. At the same time, the monetary expansion causes a decline in current consumption dispersion (bottom-left panel), implying a smaller response of (log) aggregate consumption for any given change in average (log) consumption. Like in the endowment economy, and even though the total effect is small, each of the two opposing forces has a non-negligible size, of the same order of magnitude as aggregate output fluctuations. Thus, even though the net impact of heterogeneity is very small, it appears to affects the transmission mechanism in a significant way.[21]

The intuition behind these results can be summarized as follows: the increase in overall income and wealth accruing to most households leads to a reduction in

[21] Kaplan, Moll and Violante (2018) reach a similar conclusion, though focusing on the relative weight of direct vs. indirect effects, and in a model where liquidity constraints play a key role.

Figure 4: The Effects of a Monetary Policy Shock in a NK Economy

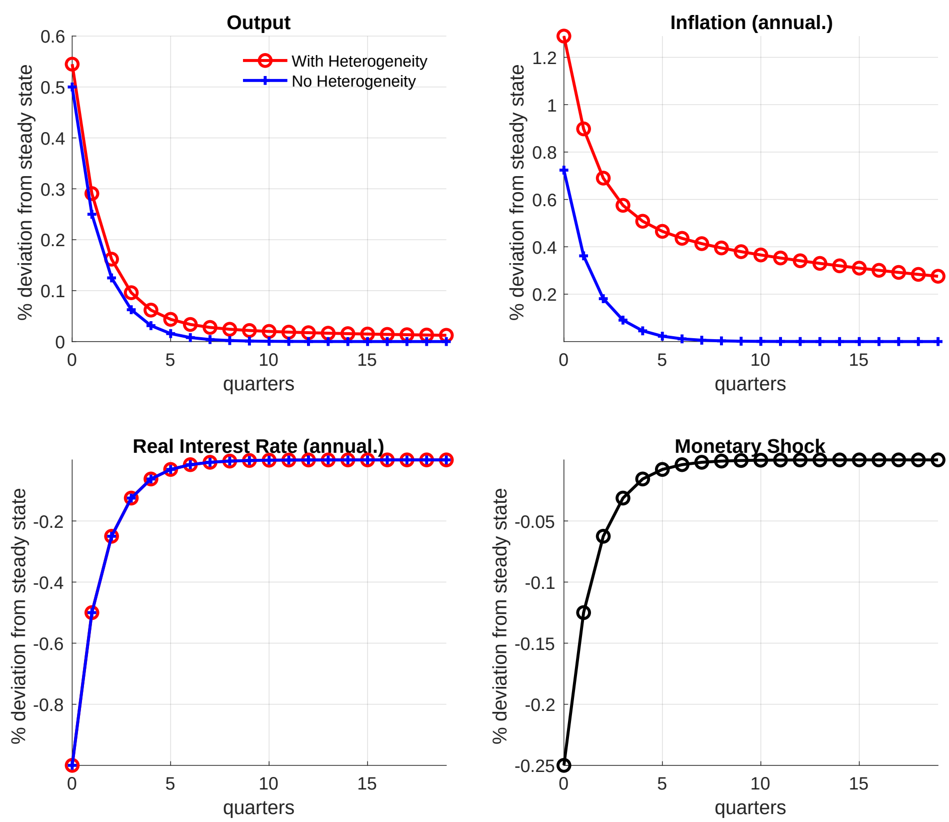

Notes: The figure shows the response of output, and of (annualized) inflation and real interest rate to a 25 basis point decrease in the real interest rate, in a model with heterogeneity (red line with circles) and without heterogeneity (blue line with crosses).

Figure 5: The Role of Heterogeneity in a NK Economy (Monetary Shock)

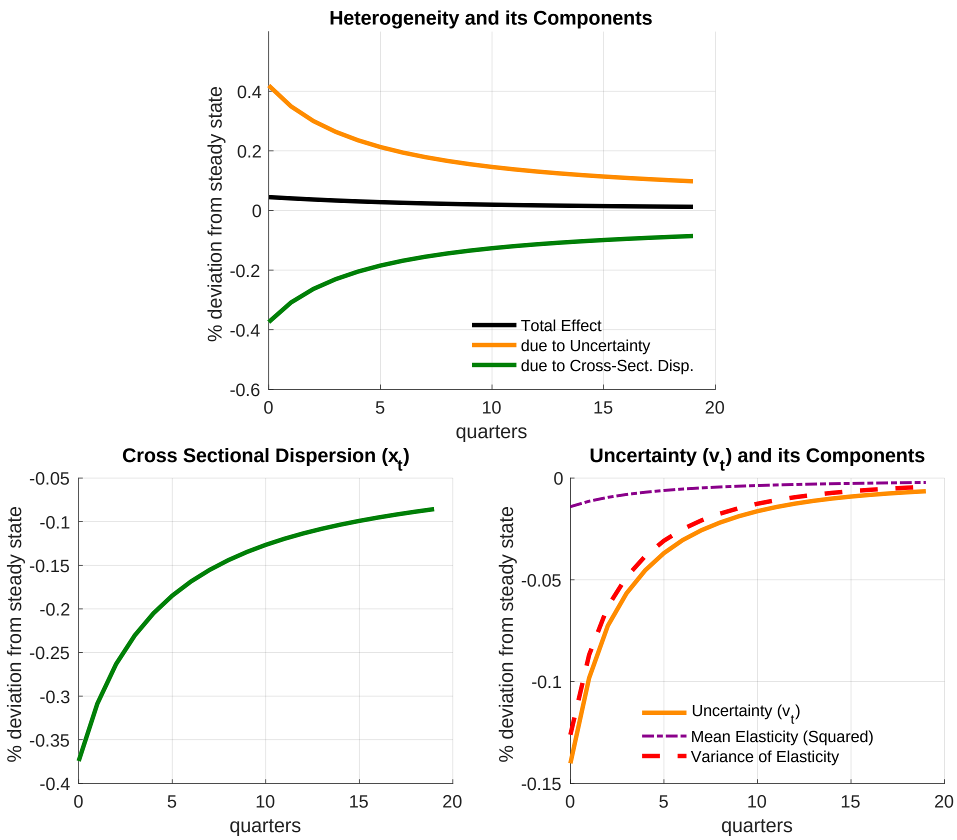

Notes: The figure shows the effect of heterogeneity on the response of aggregate output to a 25 basis point monetary shock. The top panel shows the total effect (black line) and the components due to uncertainty $-\frac{\sigma}{2}\sum_{k=0}^{\infty}\mathbb{E}_t\{\widehat{v}_{t+k}\}$ (orange line) and dispersion $\widehat{x}_t$ (green line). The bottom-left panel shows the response of dispersion $\widehat{x}_t$, and the bottom-right panel shows the response of uncertainty $\widehat{v}_t$ (orange line), together with the responses of the variance of the elasticity (red dashed line) and of the squared mean elasticity (purple dash-dotted line).

Figure 6: The Effects of a Preference Shock in a NK Economy

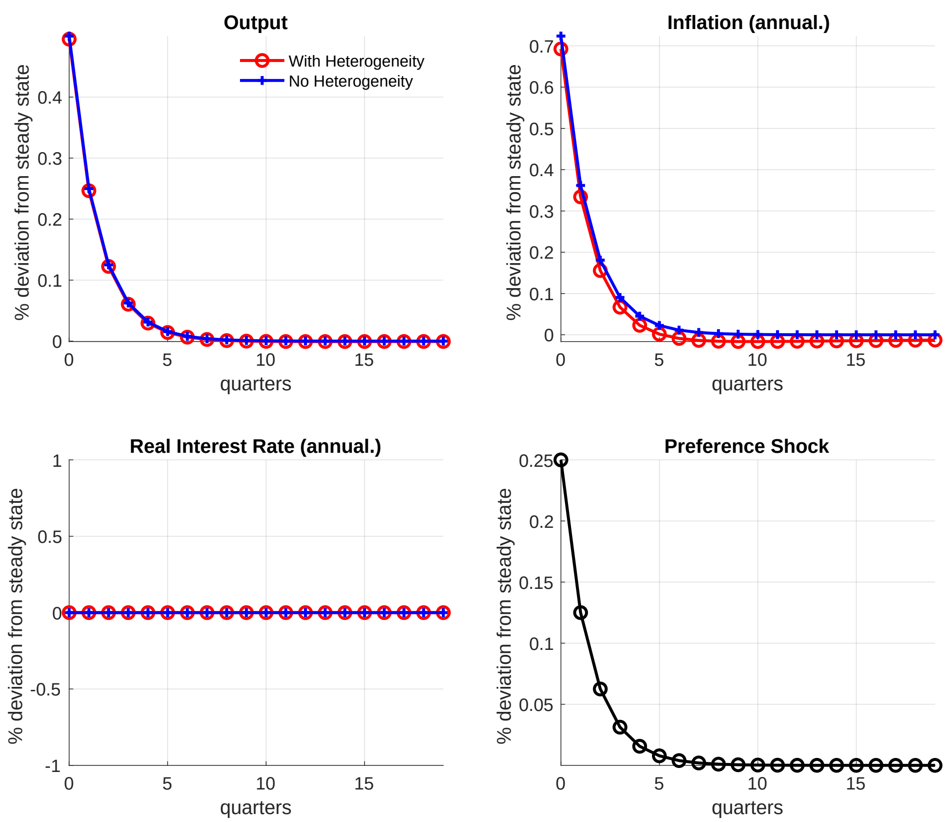

Notes: The figure shows the response of output, (annualized) inflation and (annualized) real interest rate to a 0.25 percent preference shock, in a model with heterogeneity (red line with circles) and without heterogeneity (blue line with crosses).

the mean and variance of the consumption elasticity, and hence an endogenous reduction in uncertainty, boosting the response of individual consumption (the uncertainty channel). That increase is proportionally larger for low consumption households, implying a response of aggregate consumption smaller (in percent terms) than that of average individual consumption, thus dampening the effects of the shock on aggregate demand and output (the dispersion channel).

Figure 6 displays the impulse responses to a positive preference shock. Note again the negligible difference difference in the response of output (and inflation) to that shock across the two environments (with and without heterogeneity). Interestingly, a look at the differential impulse response in Figure 7 points to a dampening effect of heterogeneity, resulting from an increase in average uncertainty, which raises precautionary savings and more than offsets the amplifying effects of the increase in dispersion. That pattern is the opposite from the one observed in response to an (equally expansionary) monetary policy shocks. The explanation for such a difference is that the preference shock raises the consumption elasticity for all levels of consumption, as illustrated in Figures 8. As a result, both the mean and variance of that elasticity increase, despite the increase in consumption due to the expansion.

Finally, Figures 9 and 10 show 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, with the impact of the decrease in both uncertainty and dispersion largely offsetting each other in a way similar to the case of a monetary policy shock.

6 The Role of Binding Borrowing Constraints

The quantitative analysis of the previous two sections has pointed to a very small impact of heterogeneity on the aggregate effects of different aggregate shocks. One may wonder to what extent the previous result is driven by our assumption of no binding borrowing constraints. In the present section we complement our quantitative analysis by studying what is the differential effect of introducing binding borrowing constraints. To that end, we determine the responses of aggregate variables to aggregate shocks, in versions of our economy where the borrowing limit $\underline{B}$ is tightened so that in steady state the constraint is binding for a fraction $\lambda = 20\%$

Figure 7: The Role of Heterogeneity in a NK Economy (Preference Shock)

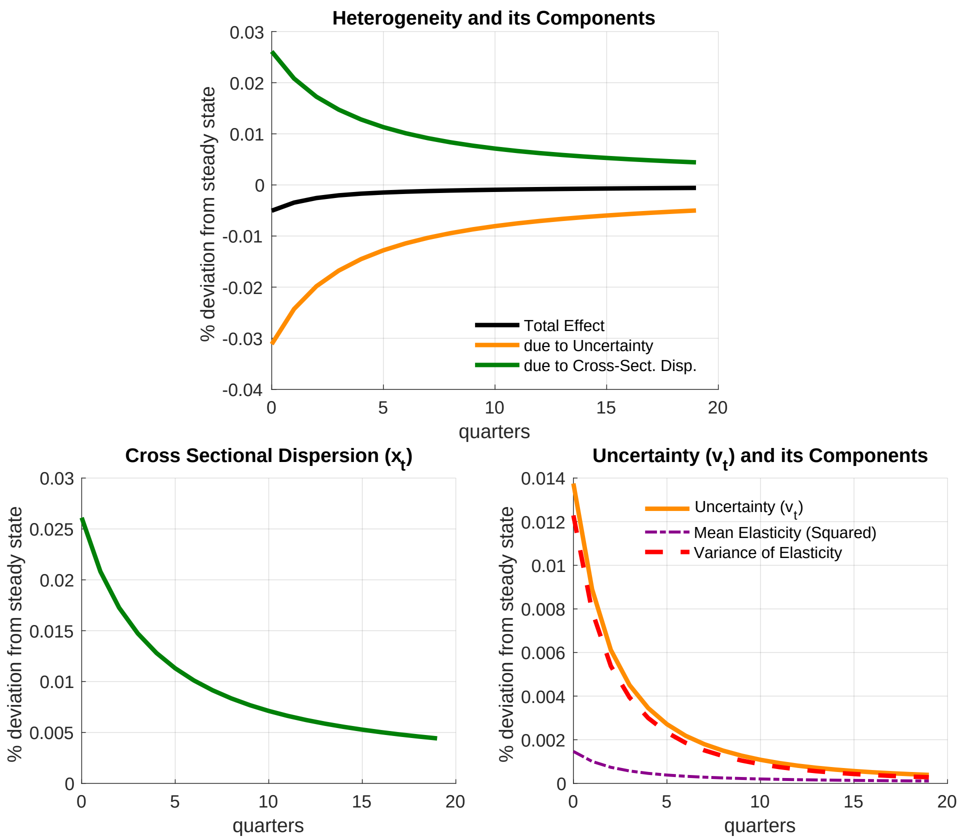

Notes: The figure shows the effect of heterogeneity on the response of aggregate output to a 0.25 percent preference shock. The top panel shows the total effect (black line) and the components due to uncertainty $-\frac{\sigma}{2}\sum_{k=0}^{\infty}\mathbb{E}_t\{\widehat{v}_{t+k}\}$ (orange line) and dispersion $\widehat{x}_t$ (green line). The bottom-left panel shows the response of dispersion $\widehat{x}_t$, and the bottom-right panel shows the response of uncertainty $\widehat{v}_t$ (orange line), together with the responses of the variance of the elasticity (red dashed line) and of the squared mean elasticity (purple dash-dotted line).

Figure 8: Preference Shocks and the Elasticity of Consumption

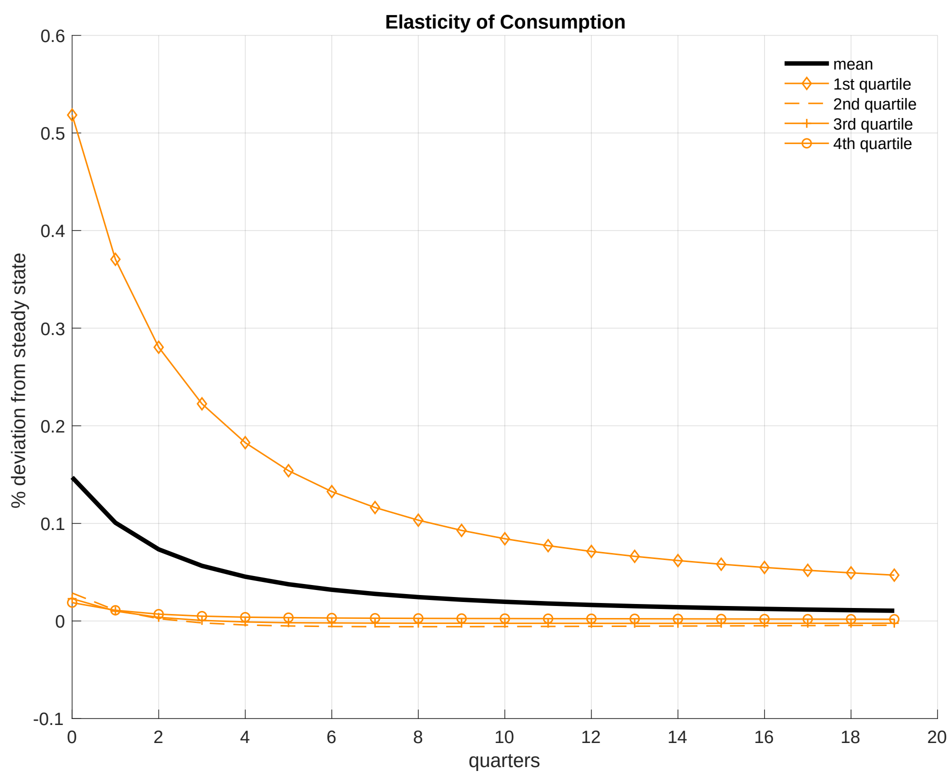

Notes: The figure shows the respone of the elasticity of consumption to idiosyncratic shocks by consumption quartile, as well as for the cross-sectional mean.

Figure 9: The Effects of a Technology Shock in a NK Economy

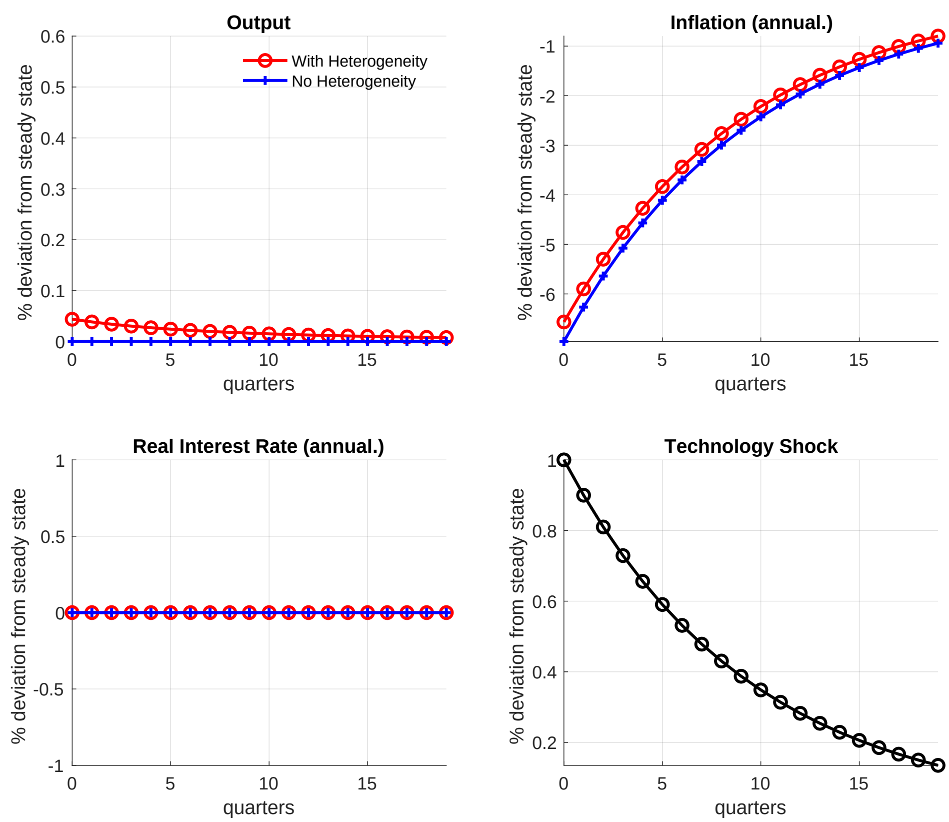

Notes: The figure shows the response of output, (annualized) inflation and (annualized) real interest rate to a 1 percent technology shock, in a model with heterogeneity (red line with circles) and without heterogeneity (blue line with crosses).

Figure 10: The Role of Heterogeneity in a NK Economy (Technology Shock)

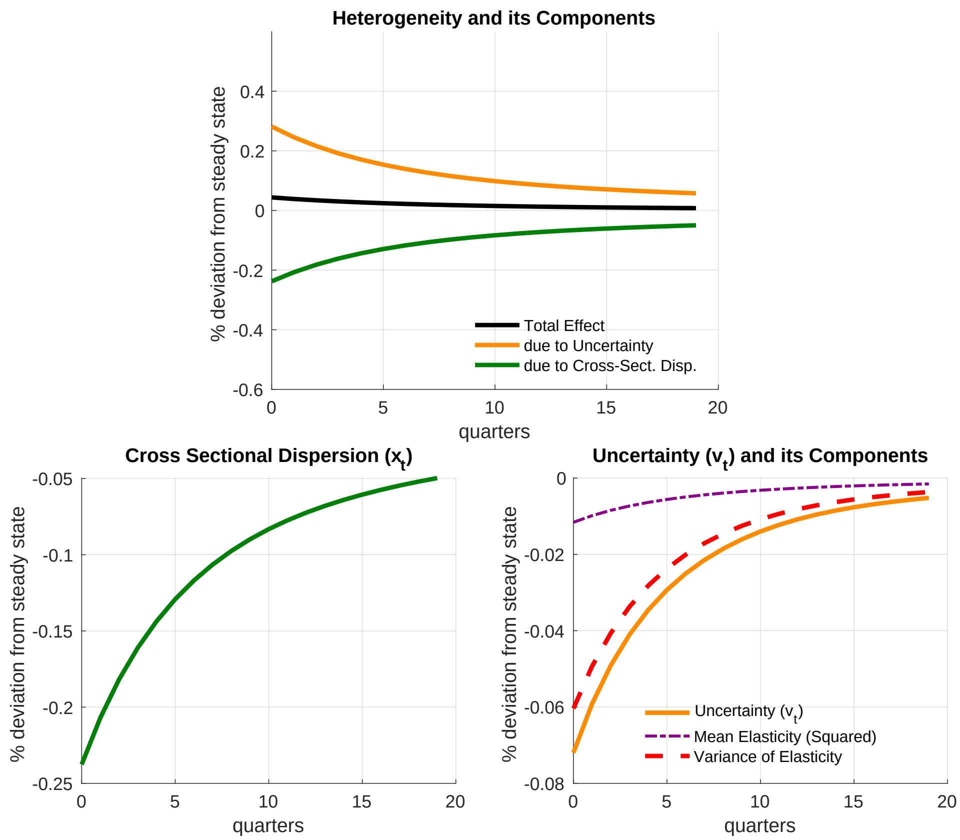

Notes: The figure shows the effect of heterogeneity on the response of aggregate output to a 1 percent technology shock. The top panel shows the total effect (black line) and the components due to uncertainty $-\frac{\sigma}{2}\sum_{k=0}^{\infty}\mathbb{E}_t\{\widehat{v}_{t+k}\}$ (orange line) and dispersion $\widehat{x}_t$ (green line). The bottom-left panel shows the response of dispersion $\widehat{x}_t$, and the bottom-right panel shows the response of uncertainty $\widehat{v}_t$ (orange line), together with the responses of the variance of the elasticity (red dashed line) and of the squared mean elasticity (purple dash-dotted line).

percent of households.

For brevity, we focus our analysis on the New Keynesian economy, and study the impact of heterogeneity in response to monetary, preference and technology shocks. Results are summarized in Figure 11 which compares the response of output to aggregate shocks in heterogeneous household economies with and without binding borrowing constraints, as well as in a RA economy. The key message from this exercise is that the presence of a borrowing limit that is occasionally binding for a fraction of households boosts significantly the role played by heterogeneity in shaping the response of output to aggregate shocks. In particular, the presence of a binding borrowing constraint leads to an amplification of the effects of monetary shocks (Panel (a)) and technology shocks (Panel (b)).[22] This follows from the presence of a substantial fraction of the population with a high elasticity of consumption out of income shocks —when the fraction of constrained households is increased from 0 to 20 percent, the average elasticity increases from 0.29 to 0.66, which correspond to an increase in the marginal propensity to consume from 0.013 to 0.25.

Instead, in the case of preference shocks (Panel (b)), the presence of a binding borrowing constraint has the opposite effect, and dampens the response of output —from 0.5 to 0.4 percentage points. Intuitively, this is because constrained households, whose consumption profile does not obey a consumption Euler equation, are not affected by a shock to intertemporal preferences. This implies a more muted response of aggregate consumption, relatively to an economy where all households are unconstrained, or to a RA economy.

7 Concluding Remarks

The objective of the present paper was to study the role of idiosyncratic income shocks for aggregate fluctuations within a simple heterogeneous household framework with no binding borrowing constraints. We derive analytically an Euler equation for (log) aggregate consumption, which helps us shed some light on the differential behavior of such an economy relative to its RA counterpart.

[22]The reason why the effects are stronger for technology shocks than for monetary shocks is related to our assumption of uniform distribution of firms' dividends, and the fact that profits are procyclical in response to technology shocks, and countercyclical in response to monetary shocks.

Figure 11: The Role of a Binding Constraint in a NK Economy

Panel (a): Response of Output to a Monetary Shock

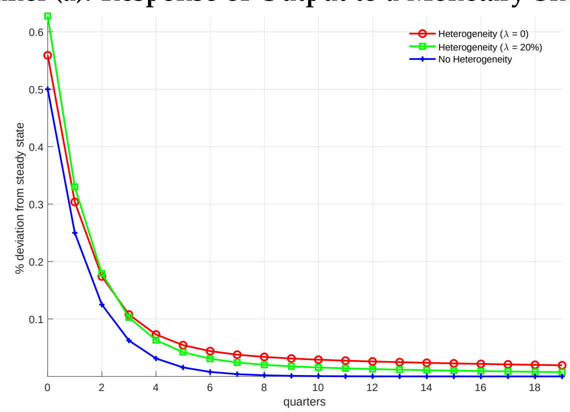

Panel (b): Response of Output to a Preference Shock

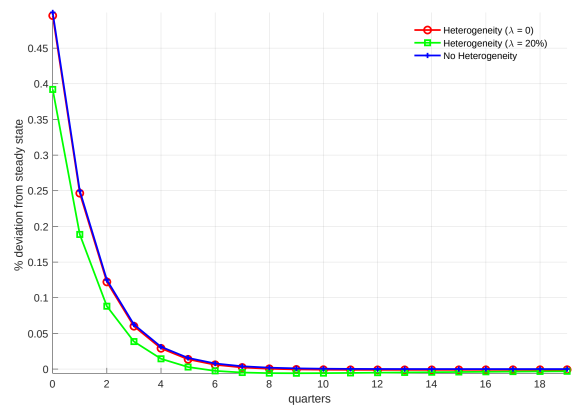

Panel (c): Response of Output to a Technology Shock

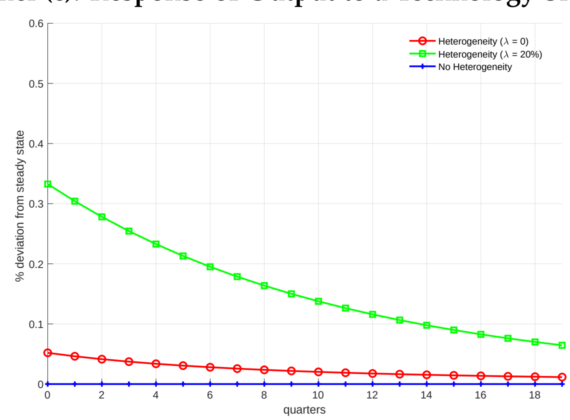

Notes: The figure plots the response of output in a heterogeneous household economy with binding constraints for a share $\lambda = 20\%$ of the population (green line with squares), with no binding constraints (red line with circles) and in an economy without heterogeneity (blue line w*4ith crosses), Panel (a) refers to a 25 basis point monetary shock, Panel (b) to a 25 percent preference shock and Panel (c) to a 1 percent technology shock.

We uncover two channels through which household heterogeneity has an impact on aggregate outcomes: (i) changes in average consumption uncertainty and (ii) changes in the cross-sectional dispersion of consumption. Both channels point to the key role of uncertainty dynamics in understanding the impact of heterogeneity. In particular, differences with RA outcomes are captured by a sufficient statistic, the consumption-weighted average of changes in uncertainty. We further show that individual consumption uncertainty is tightly linked to the elasticity of consumption with respect to idiosyncratic income shocks.

We conduct a quantitative analysis in the context of two full-fledged calibrated economies subject to aggregate shocks: an endowment economy and a New Keynesian economy. Our quantitative analysis shows small effects of heterogeneity, largely as a consequence of the mutually offsetting effects of the two channels mentioned above.

When binding borrowing constraints are introduced, affecting an empirically plausible fraction of households, the impact of heterogeneity is boosted significantly. The latter finding, combined with the near-irrelevance of other dimensions of heterogeneity, suggests that two-agent models of the sort used in the TANK literature may provide a reasonable approximation to the aggregate properties of richer HA models, an avenue we pursue in a companion paper (Debortoli and Galí (2018)).[23]

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 (2019): "The Intertemporal Keynesian Cross," mimeo.

[23]See e.g. Auclert, Rognlie and Straub (2019) for a discussion of possible limitations of conventional TANK models.

Adrien Auclert, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub (2020): “Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models,” *Econometrica*, Forthcoming.

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.

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. (2019): “Monetary Policy and Heterogeneity: An Analytical Framework,” mimeo.

Broer, Tobias, Niels-Jakob Harbo Hansen, Per Krusell, and Erik Oberg (2019): “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 Galï¿œ (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 hetereogeneity: 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 Agenst 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.

Appendices

A An Euler Equation for Aggregate Consumption: Alternative Representation

Consider an economy where the following Euler equation holds for all households $j \in [0,1]$

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

where $C_t(j)$ is individual consumption and $R_t \equiv \exp\{r_t\}$ is the gross yield on a one-period riskless bond and $Z_t \equiv \exp z_t$ is an aggregate preference shock following a stationary AR(1) process $z_t = \rho_z z_{t-1} + \varepsilon_{z,t}$.

Defining $c_t \equiv \log C_t(j)$, multiplying and dividing the RHS of (A.1) by $\exp\{\mathbb{E}_t [z_{t+1} - \sigma c_{t+1}(j)]\}$, and taking logs of the resulting expression yields

$$c_t(j) = \mathbb{E}_t\{c_{t+1}(j)\} - \frac{1}{\sigma}(r_t - \rho) + \frac{1}{\sigma}(1-\rho_z)z_t - \frac{1}{\sigma}v_t(j) \tag{A.2}$$

where $\rho \equiv -\log \beta$ and

$$v_t(j) \equiv \log\{\mathbb{E}_t \exp[(z_{t+1} - \mathbb{E}_t z_{t+1}) - \sigma(c_{t+1}(j) - \mathbb{E}_t c_{t+1}(j))]\}. \tag{A.3}$$

Notice that, in the case of perfect foresight about z_{t+1} and $c_{t+1}(j)$, eq. (A.3) implies that $v_t(j) = 0$. Thus, $v_t(j)$ can be interpreted as capturing the effects of individual uncertainty, as a consequence of aggregate or idiosyncratic shocks. This is consistent with what obtained in the main text (see eq. 2), where we relied on a second order approximation, under the assumption of "small" aggregate shocks, so that the term $v_t(j)$ would be (approximately) only a function of to the conditional variance of consumption "surprises" $\xi_{t+1} \equiv c_{t+1} - \mathbb{E}_t c_{t+1}$.

Intergrating eq. (A.1) across all households gives

$$\bar{c}_t = \mathbb{E}_t\{\bar{c}_{t+1}\} - \frac{1}{\sigma}(r_t - \rho) + \frac{1}{\sigma}(1-\rho_z)z_t - \frac{1}{\sigma}v_t \tag{A.4}$$

where $\bar{c}_t \equiv \int c_t(j)dj$ and $v_t \equiv \int v_t(j)dj$ denote, respectively, the cross-sectional average (log) consumption and uncertainty at a given point in time t.

Finally, defining the gap between (log) aggregate consumption and average (log) consumption $x_t \equiv c_t - \bar{c}_t$. As explained in the main text, such a gap is due to Jensen's inequality, and can then be interpreted as a proxy for the effects of dispersion of individual consumption aroung the cross-sectional mean.

Using the definition of x_t into (A.5) we obtain an Euler equation for aggregate consumption as

$$c_t = \mathbb{E}_t\{c_{t+1}\} - \frac{1}{\sigma}(r_t - \rho) + \frac{1}{\sigma}(1-\rho_z)z_t - \frac{1}{\sigma}v_t - \mathbb{E}_t \Delta x_{t+1}. \tag{A.5}$$

The latter expression, which is analogous to eq. (6) in the main text, highlights that the effects of household heterogeneity are captured by the terms associated with *uncertainty* (v_t) and *dispersion* (x_t).

B Variations in Uncertainty ($\widehat{v}_t$): The General Case

Let the consumption function for household j be given by:

$$c_t(j) = c^*(s_t(j), S_t)$$

where $s_t(j)$ is a vector of household-specific state and S_t is a vector of aggregate states. The latter contains all the information available at time t regarding the distribution of aggregate variables (wages, interest rates, etc) whose current and future values are relevant for today's consumption decision.

We can thus approximate the innovation in household j's consumption in period t as follows:

$$\begin{aligned} \xi_t(j) &\equiv c_t(j) - \mathbb{E}_{t-1}\{c_t(j)\} \\ &= f_{t-1}(j)[\eta_t(j), \varepsilon_t] \end{aligned}$$

where $\eta_t(j)$ and ε_t are vectors containing respectively the idiosyncratic and aggregate shocks (i.e. mutually orthogonal innovations in the individual and aggregate *exogenous* driving variables), and where $f_{t-1}(j)[\cdot]$ is a function satisfying $f_{t-1}(j)[0,0] = 0$.

Under our assumptions, individual uncertainty $v_t(j)$ in period t can thus be approximated as

$$\begin{aligned} v_t(j) &\equiv \mathbb{E}_t\{\xi_{t+1}(j)^2\} \\ &= \psi_t(j)'\Omega_\eta \psi_t(j) + \varphi_t(j)'\Omega_\varepsilon \varphi_t(j) + O(|y|^2) \end{aligned}$$

where $[\psi_{t-1}(j)', \varphi_{t-1}(j)'] \equiv \frac{\partial f_{t-1}(j)[0,0]}{\partial[\eta_t(j), \varepsilon_t]}$ is the Jacobian of $f_{t-1}(j)$ evaluated at 0, with $\Omega_\eta \equiv \mathbb{E}\{\eta_t(j)\eta_t(j)'\}$ for all $j \in [0,1]$ and $\Omega_\varepsilon \equiv \mathbb{E}\{\varepsilon_t \varepsilon_t'\}$ being the (diagonal) matrices containing the variance of the exogenous shocks on their diagonals. We can think of $\psi_{t-1}(j)$ and $\varphi_{t-1}(j)$ as the "local" elasticities of individual consumption with respect to those shocks.

Aggregating across households

$$v_t = \int \psi_t(j)'\Omega_\eta \psi_t(j) dj + \int \varphi_t(j)'\Omega_\varepsilon \varphi_t(j) dj + O(|y|^2)$$

Next we show that $\int \varphi_t(j)'\Omega_\varepsilon \varphi_t(j) dj \sim O(|y|^2)$ if we are willing to impose bounds on the extent of heterogeneity in $\varphi_t(j)$. Note that we can write

$$\int \varphi_t(j)'\Omega_\varepsilon \varphi_t(j) dj = \overline{\varphi}_t'\Omega_\varepsilon \overline{\varphi}_t + \int (\varphi_t(j) - \overline{\varphi}_t)'\Omega_\varepsilon(\varphi_t(j) - \overline{\varphi}_t) dj$$

where $\overline{\varphi}_t \equiv \int \varphi_t(j) dj$.

Furthermore, note that the innovation to *average* log consumption, $\overline{\xi}_t$, can be written as

$$\begin{aligned} \overline{\xi}_t &= \int \xi_t(j) dj \\ &\simeq \int f_{t-1}(j)[\eta_t(j), \varepsilon_t] dj \end{aligned}$$

Accordingly,

$$\begin{aligned} \mathbb{E}_t\{\overline{\xi}_{t+1}^2\} &\simeq \mathbb{E}_t\left\{\left(\int \psi_t(j)'\eta_t(j) dj + \overline{\varphi}_t'\varepsilon_t\right)^2\right\} \\ &\geq \overline{\varphi}_t'\Omega_\varepsilon \overline{\varphi}_t \end{aligned}$$

Our assumption $\mathbb{E}_t\{\overline{\xi}_{t+1}^2\} \sim O(|x_t|^2)$ thus implies $\overline{\varphi}_t'\Omega_\varepsilon \overline{\varphi}_t \sim O(|x_t|^2)$. Thus, if we assume $\int (\varphi_t(j) - \overline{\varphi}_t)'\Omega_\varepsilon(\varphi_t(j) - \overline{\varphi}_t) dj \sim O(|x_t|^2)$ it follows that $\int \varphi_t(j)'\Omega_\varepsilon \varphi_t(j) dj \sim$

$O(|x_t|^2)$ and

$$\begin{aligned} v_t &\simeq \int \psi_t(j)' \Omega_\eta \psi_t(j) dj \\ &= \overline{\psi}_t' \Omega_\eta \overline{\psi}_t + \int (\psi_t(j) - \overline{\psi}_t)' \Omega_\eta (\psi_t(j) - \overline{\psi}_t) dj \end{aligned}$$

which varies with the mean and dispersion of the individual elasticities of consumption *with respect to the idiosyncratic shock*. Both statistics will in turn change over time in response to changes in the distribution of wealth. In the simple endowment economy considered below,

$$v_t = [\overline{\psi}_t^2 + var_j\{\psi_t(j)\}]\sigma_\eta^2.$$

Online Appendices

A Robustness: 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 $\eta_t(i)$. In particular, we consider a discrete-time quarterly version of the continous-time process used in Kaplan, Moll and Violante (2018), which is the sum of two independent components $\eta_t(i) = \eta_{1,t}(i) + \eta_{2,t}(i)$. 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 continous-time components are discretized with 3 grid points for η_1 (transitory component) and 11 points for η_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 A.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 A.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 A.1: Distribution of (Log) Income Shocks in the Alternative Calibration

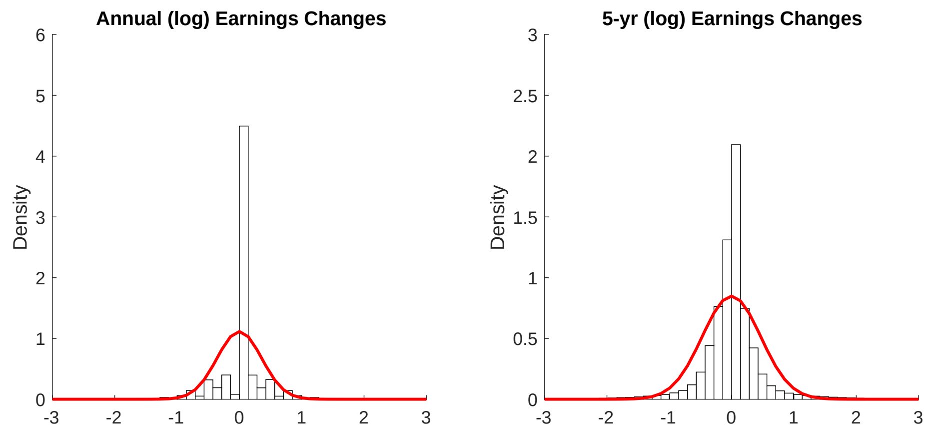

Notes: The figure shows the distribution of (log) earning changes at annual frequency (left panel) and at a 5yr frequency (right panel). In each panel, the histograms correspond to the distribution resulting from the (discretized) process with a transitory and apersistent component, while the solid line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

B 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 $\widehat{r}_t = \phi_\pi \pi_t + m_t$, 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).

Figures B.3-B.5 report the response of aggregate variables to, respectively, monetary shocks, preference shocks and technology shocks.

In response to all these shocks, and analogously to what shown in Figures 4, 6 and 9 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 A.2: The Effects of Monetary Shocks with Different Idiosyncratic Risk

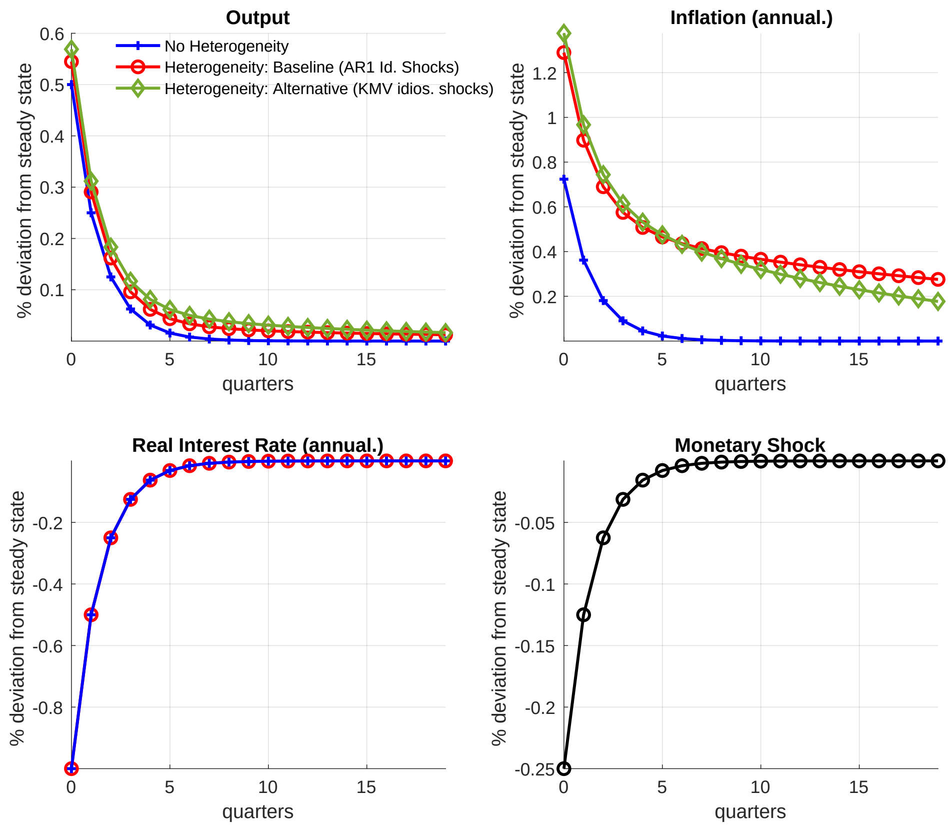

Notes: The figure shows the response of output, (annualized) inflation and (annualized) real interest rate to a 25 basis points expansionary monetary shock. The figure compares the responses in a model without heterogeneity (blue line with crosses), in the baseline heterogeneous household model with 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 B.3: The Effects of a Monetary Shock (Monetary Rule)

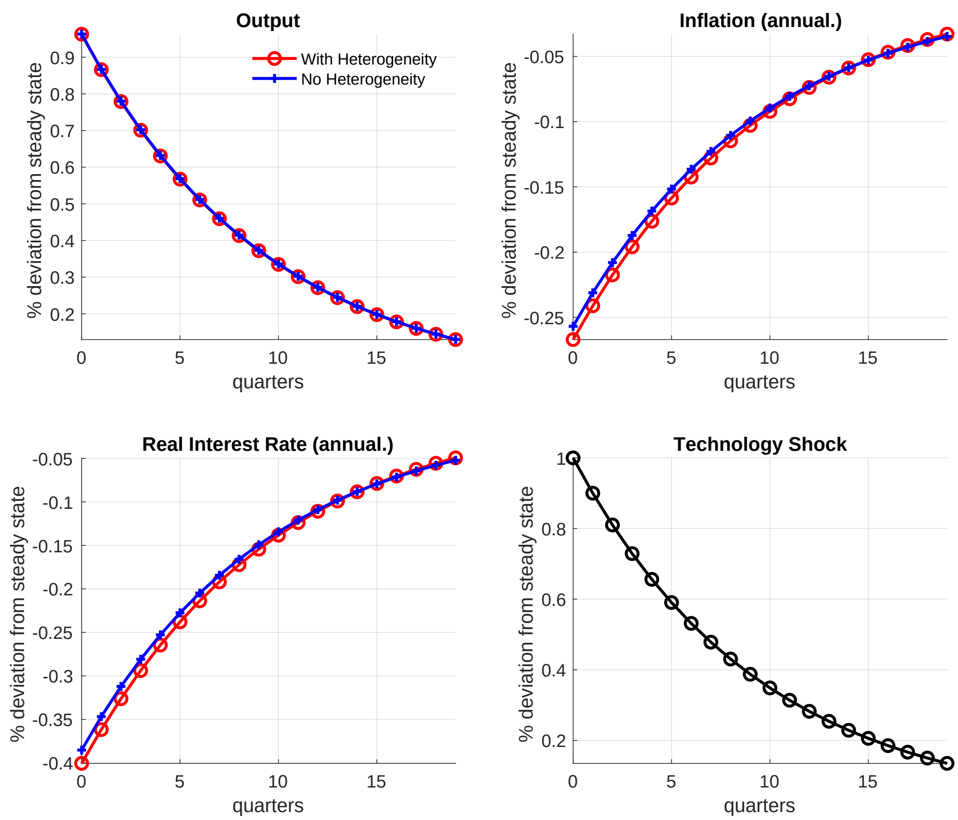

Notes: The figure shows the response of output, (annualized) inflation and (annualized) real interest rate to a 25 basis point monetary shock, in a model with heterogeneity (red line with circles) and without heterogeneity (blue line with crosses).

Figure B.4: The Effects of a Preference Shock (Monetary Rule)

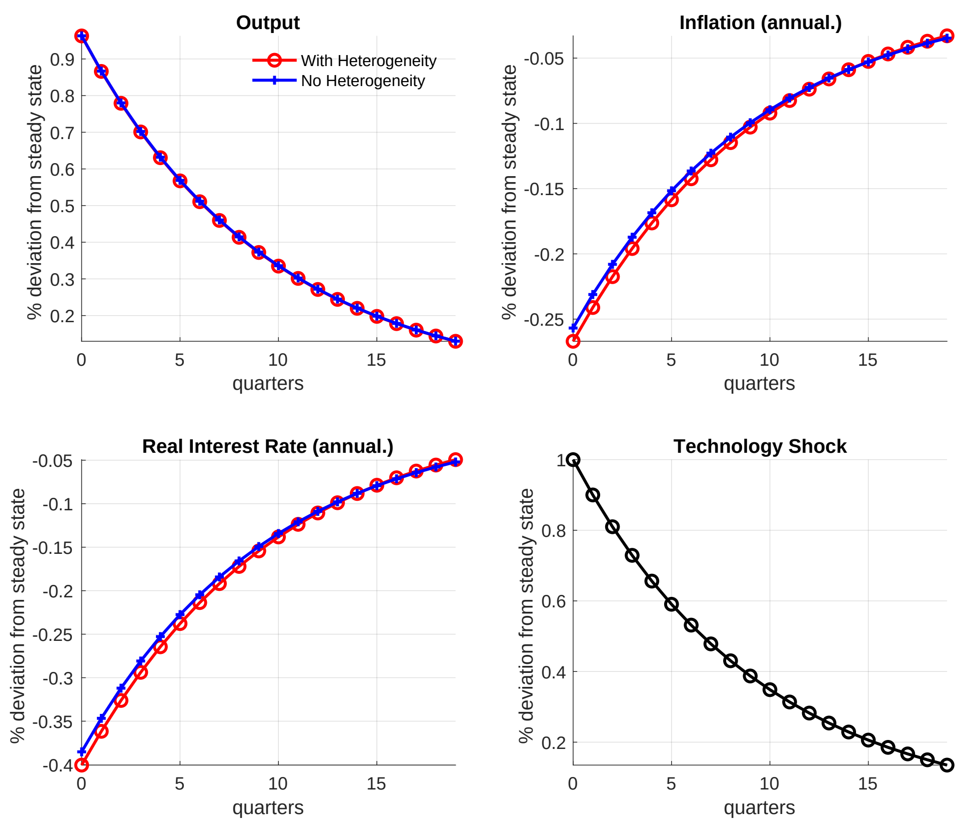

Notes: The figure shows the response of output, (annualized) inflation and (annualized) real interest rate to a 0.25 percent preference shock, in a model with heterogeneity (red line with circles) and without heterogeneity (blue line with crosses).

Figure B.5: The Effects of a Technology Shock (Monetary Rule)

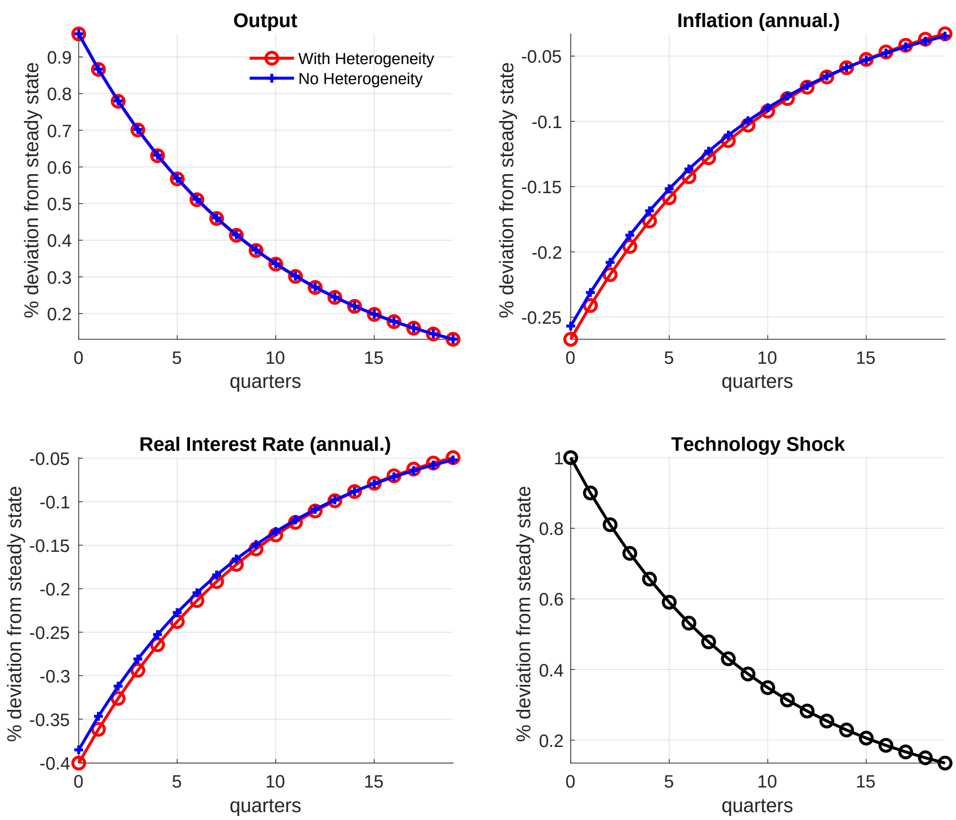

Notes: The figure shows the response of output, (annualized) inflation and (annualized) real interest rate to a 1 percent technology shock, in a model with heterogeneity (red line with circles) and without heterogeneity (blue line with crosses).