DISCUSSION PAPER SERIES

DP17080
(v. 3)

Monetary Policy with Heterogeneous Risk

Tommaso Monacelli and Roberto Colarieti

INTERNATIONAL MACROECONOMICS AND FINANCE

MONETARY ECONOMICS AND FLUCTUATIONS

CEPR

ISSN 0265-8003

Monetary Policy with Heterogeneous Risk

Tommaso Monacelli and Roberto Colarieti

Discussion Paper DP17080
First Published 02 March 2022
This Revision 02 March 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Monetary Policy with Heterogeneous Risk

Abstract

We lay out a heteroskedastic New Keynesian model, consistent with the evidence of pervasive cross-sectional variation in individual income risk in micro data. We obtain three main results. First, heterogeneous marginal propensities to consume (MPCs) result from the sensitivity of precautionary savings to realized earnings being heterogeneous across agents. Second, the response of aggregate output to demand shocks hinges crucially on how individual risk co-varies with the degree of individual income cyclicality across the income distribution. Third, the general equilibrium effects of monetary and fiscal policy can be suitably summarized by a set of observable cross-sectional sufficient statistics. Depending on the sign of those statistics, income heteroskedasticity may: (i) dampen or amplify the response of output to demand shocks; (ii) affect the local determinacy of the equilibrium; (iii) attenuate or exacerbate the forward guidance puzzle. We conjecture an incomplete information framework with Bayesian learning as a plausible microfoundation for individual income heteroskedasticity.

JEL Classification: E52, E32, E21

Keywords: Heteroskedasticity, precautionary savings, aggregate demand, monetary policy

Tommaso Monacelli - tommaso.monacelli@unibocconi.it
Bocconi University and CEPR

Roberto Colarieti - roberto.colarieti@unibocconi.it
Bocconi University, IGIER

Monetary Policy with Heterogeneous Risk*

Roberto Colarieti
Bocconi University

Tommaso Monacelli
Bocconi University, IGIER and CEPR

March 2022

Abstract

We lay out a *heteroskedastic* New Keynesian model, consistent with the evidence of pervasive cross-sectional variation in individual income risk in micro data. We obtain three main results. First, heterogeneous marginal propensities to consume (MPCs) result from the *sensitivity* of precautionary savings to realized earnings being heterogeneous across agents. Second, the response of aggregate output to demand shocks hinges crucially on how individual risk co-varies with the degree of individual income cyclicality across the income distribution. Third, the general equilibrium effects of monetary and fiscal policy can be suitably summarized by a set of observable cross-sectional sufficient statistics. Depending on the sign of those statistics, income heteroskedasticity may: (i) dampen or amplify the response of output to demand shocks; (ii) affect the local determinacy of the equilibrium; (iii) attenuate or exacerbate the forward guidance puzzle. We conjecture an *incomplete information* framework with Bayesian learning as a plausible microfoundation for individual income heteroskedasticity.

*JEL codes E52, E32, E21. Correspondence: roberto.colarieti@unibocconi.it, tommaso.monacelli@unibocconi.it. We thank Pierfrancesco Mei and Gianluca Violante for insightful comments. All errors are our own.

1 Introduction

A recent literature in macroeconomics has developed models that incorporate heterogeneity and uninsurable idiosyncratic risk into the standard New Keynesian apparatus with a representative agent (aka HANK models). There is, however, a large debate on *which* dimension of heterogeneity is essential to deepen our understanding of aggregate fluctuations as well as the channels of transmission of monetary and fiscal policy.

Traditionally the literature on heterogeneous agent models has been centered on the assumption that individual households face a *homoskedastic* income process: that is, the variance of the income process is homogeneous in the cross-section and constant over time. However, the evidence from micro data points to significant cross-sectional *heteroskedasticity* in the income process. Relying on a rich administrative dataset for the U.S., Guvenen et al. (2021) show that the second moment of individual income (log) growth varies substantially along the realized earnings distribution. In particular, and uniformly across the age profiles, individual income risk depends significantly on *past* realized individual income. This feature of the income distribution has been, however, largely ignored in the recent HANK literature.

We explore the implications of heteroskedasticity in individual income by means of a CARA (constant absolute risk aversion) utility model with a Gaussian heteroskedastic income process, building on the work by Caballero (1990) and, more recently, Acharya and Dogra (2020). Relative to the homoskedastic CARA setup of Acharya and Dogra (2020), our framework is able to account *both* for idiosyncratic risk and heterogeneity in marginal propensities to consume (MPCs) in an analytically tractable way, making the general equilibrium interaction between risk and the cross-sectional distribution of income extremely transparent.

Recent papers have shown that the *cyclicality* of income risk, i.e., how individual risk behaves during expansions and contractions, can be a source of either amplification (in the case of counter-cyclical risk) or dampening (in the case of pro-cyclical risk) of aggregate fluctuations.[1] For instance, agents react to negative aggregate demand shocks by decreasing their consumption relatively more when risk is counter-cyclical, since aggregate contractions are also associated with higher individual risk and precautionary savings. While this notion of risk-cyclicality captures the idea that idiosyncratic income risk may be *on average* higher during recessions,[2] it disregards the fact that risk varies *across the earnings distribution*. Within our setting, realized changes in earnings endogenously affect the risk faced by each agent, and thus their optimal consumption and savings decisions. Suppose, for instance, that the variance of the individual income process is an inverse (and non-linear) function of realized earnings. In this case, a negative realization of income also increases individual risk and, in turn, agent's precautionary savings become larger, at the margin. As a result, the *sensitivity* of precautionary savings to income realizations will be heterogeneous

[1] For instance, Acharya and Dogra (2020) and Bilbiie (2018) for income risk, Ravn and Sterk (2021) for unemployment risk.

[2] See, for instance, Storesletten et al. (2004).

across individuals.

We employ our heteroskedastic NK framework to study three questions that have been at the core of the recent HANK literature: (a) the transmission of monetary and fiscal shocks; (b) equilibrium determinacy; (c) the so-called forward guidance puzzle (henceforth FGP).

We reach three main results. First, and unlike the corresponding homoskedastic case, the CARA-Gaussian setup augmented with heteroskedasticity features a non-degenerate *distribution* of MPCs out of transitory income. This is due to realized earnings affecting income volatility at the margin, and therefore also precautionary savings and consumption choices, as long as the marginal change in volatility varies over the earnings distribution. In particular, and conditional on the variance being a convex function of income, our model is able to generate MPCs which are decreasing over the income distribution, in line with some empirical evidence.[3] To gain intuition for this result, suppose that income volatility is decreasing and convex in realized earnings. Then any additional unit of income determines a decline in precautionary savings that is, at the margin, stronger for low-income households, who in turn react by spending more of the income shock. As a result, they exhibit relatively higher MPCs.

Second, the presence of heteroskedasticity has general equilibrium implications that are absent both in a representative-agent (RA) model and in a corresponding version of our framework with homoskedasticity. The interaction between income risk and the income distribution determines indirect effects that shape the responses of key aggregate variables to demand shocks, the determinacy properties of the equilibrium, and the FGP.

Third, the aforementioned general equilibrium indirect effects of heteroskedasticity can be described by means of three cross-sectional sufficient statistics: (1) the expected marginal change in income volatility out of individual earnings (*individual risk* effect); (2) the average sensitivity of individual income to aggregate income (*income cyclicality* effect); (3) the covariance between the marginal change in income volatility and the sensitivity of individual to aggregate income (*risk-cyclicality covariance* effect). These sufficient statistics allow to reduce the dimensionality of the model when describing the dynamic response of output to demand shocks. In particular, statistics (1) and (2) jointly capture how, on average, risk varies as a function of individual income in response to a variation in aggregate output. Statistic (3), on the other hand, captures the extent to which individual risk varies differently across agents depending on how individual income is affected by aggregate income fluctuations.

Depending on the signs of those sufficient statistics, heteroskedasticity could either amplify or dampen the response of output to demand shocks, affect the local determinacy of the equilibrium, and can either attenuate or exacerbate the FGP. In particular, negative signs for (1) and (3) amplify the effect of demand shocks on output, make the Taylor principle insufficient for local determinacy,

[3] For instance, Jappelli and Pistaferri (2014) estimate (using survey data from the Bank of Italy Survey of Household Income and Wealth) that annual MPCs are higher by 11% in the lowest quintile of the income distribution relative to the highest quintile. Johnson et al. (2006) also show (with CEX data) that MPCs out of the 2001 U.S. tax rebates are higher for low-income households.

and exacerbate the FGP, while for positive values of (1) and (3) the opposite occurs (i.e., there is dampening of demand shocks, the Taylor principle is no longer a necessary condition for equilibrium determinacy, and there is attenuation of the FGP).[4] Conversely, both in the benchmark complete markets RA model and in the homoskedastic version of our CARA framework, statistics (1) and (3) are always equal to zero.

Notice that statistics (1) and (2) jointly describe the size of the effect of *average* changes in risk on output. For instance, a stronger average increase in income volatility generates a larger increase in precautionary savings, and therefore a larger contraction in consumption following a negative demand shock. Statistic (3), on the other hand, captures the *interaction* between income risk and income redistribution over the business cycle. Suppose that a negative demand shock has a first-round *direct* contractionary effect on output. To start with, lower aggregate income affects each agent depending on the heterogeneous sensitivity of individual income to aggregate income. In our heteroskedastic setup, however, any change in individual earnings affects the volatility of the income process, and therefore saving and consumption decisions. If the covariance term (3) is negative, those agents whose income falls relatively more out of a contraction in aggregate output are also the ones who experience a marginally stronger increase in idiosyncratic risk. At the margin, those agents exhibit a relatively stronger precautionary saving motive and will decrease their consumption relatively more. In other words, those agents whose income is more heavily affected by the contraction in aggregate output are also the ones who have a stronger incentive to decrease their consumption. This indirect *propagation through-risk* channel generates an amplification of the primitive shock on aggregate output.

As for equilibrium determinacy, a negative sign of the risk-cyclicality covariance (3), like in the previous example, makes the Taylor principle insufficient for local determinacy and requires monetary policy to react more strongly to a demand shock to stabilize the economy. On the contrary, a positive sign for the covariance (3) makes the equilibrium determinate also under an interest rate peg.[5]

Moreover, a negative covariance (3) exacerbates the FGP. This results from the interaction of two mechanisms. For one, with a negative covariance, agents expect that a monetary easing T periods ahead has an amplified positive effect on output T periods ahead (as previously described). In addition, agents expect higher output T periods ahead to cause output in $T-1$ to increase more than one-to-one. Overall, news about higher future output are compounded backward, resulting in a stronger increase in current output the longer the horizon T of the monetary easing. Conversely, a positive covariance (3) attenuates the FGP since these mechanisms operate in the opposite direction (i.e., future news are discounted).

Finally, we provide a microfoundation for the link between individual income volatility and lagged earnings, based on *incomplete information.* We assume that households cannot observe the

[4] In general, statistic (2) has to be positive, in the sense that average individual income has to be increasing in output.

[5] Unlike the well-known result by Sargent and Wallace (1975).

volatility of their income process and use available information to make inference on the unknown variance parameter. In particular, we impose heterogeneity in the information sets across agents, since each household observes both aggregate variables and the realization of her own income process, but *not* those of other households. Their current earnings realization is the private signal used to update the estimate of future volatility, via a Bayesian learning procedure. Therefore, heterogeneity in signals translates into heterogeneity in the estimates of the income variance. Our main result is to show that, in this framework, heteroskedasticity in the individual income process can be derived as an implication of informational frictions and heterogeneity in the households' information sets.

1.1 Relation to the literature

Our work relates to several strands of the literature. Concerning the propagation of monetary and fiscal policy, Kaplan et al. (2018) disentangle *direct* and *indirect* effects of monetary policy on output and show that, in a HANK model, policy transmission works mainly through the indirect channel. We highlight a new general equilibrium channel of transmission of monetary and fiscal policy, stemming from the interaction between income redistribution and changes in income volatility. Recent papers (such as Auclert, 2019; Auclert et al., 2018; Debortoli and Galí, 2022; Wolf, 2021) have highlighted the role of sufficient statistics as a way to reduce the dimensionality of HA models and to summarize the equilibrium effects of fiscal and monetary policy. We contribute further by showing that the heteroskedasticity channel of output responses to demand shocks can be suitably captured by a few observable sufficient statistics.

There is a large literature on the theory of income risk and market incompleteness. In particular, within a CARA utility framework, Caballero (1990) shows how income risk produces a drift in the consumption process leading to excess consumption growth. Wang (2002) enriches the CARA model by Caballero (1990) to study the role of conditionally heteroskedastic income, postulating a linear relation between conditional income volatility and lagged earnings. Compared to those papers, our model, while making use of CARA utility, imposes only a differentiability requirement to the conditional variance (as a function of lagged earnings), without restricting the dependence to be linear. Wang (2003) and Acharya and Dogra (2020) combine the consumption model of Caballero (1990) with a general equilibrium setting. We contribute further by showing the general equilibrium effects of heteroskedasticity.

The HANK setup of Bayer et al. (2019) assume a heteroskedastic process for labor productivity where the evolution of the variance over time depends on exogenous shocks. Unlike them, we focus on the cross-sectional dimension and we take heteroskedasticity as endogenously determined by the income distribution. Bilbiie (2018) builds a two-agent NK model augmented with idiosyncratic risk, modeled through a Markov transition structure between hand-to-mouth (HtM) and non-HtM states where the transition probabilities co-move with aggregate output. Since non-HtM agents are identical across them, they all face the same individual risk and aggregate fluctuations are linked to the cyclical property of the transition probabilities. In addition, Acharya and Dogra (2020) build a CARA utility model with a Gaussian income process where the volatility is identical across households, but depends over time on aggregate output, capturing the cyclicality of individual income risk. Unlike these papers, we assume that individual income is heteroskedastic in the cross-section and we show that cross-sectional heterogeneity in idiosyncratic risk interacts with the earnings distribution, thereby affecting aggregate fluctuations in response to demand shocks.

Debortoli and Galí (2022) study the effects of heterogeneous changes in individual consumption risk over the business cycle. Their setup has two main features. First, a precautionary motive due to incomplete insurance. Second, heterogeneity in MPCs. In their model, this last feature is due to heterogeneous assets positions, hence to a heterogeneous distance from a natural borrowing limit. Differently from them, in our model heterogeneity in MPCs is rooted in heteroskedasticity.

In addition, we contribute to the literature on *tractable* heterogeneous agent models (for instance, Acharya et al., 2020; Acharya and Dogra, 2020; Bilbiie, 2008, 2018, 2020; Bilbiie et al., 2013; Challe, 2020; Challe and Ragot, 2016; Ravn and Sterk, 2021) building a general equilibrium HA framework where we integrate heteroskedasticity and idiosyncratic risk with heterogeneity in MPCs and a non-degenerate wealth distribution in an analytically tractable and transparent way.

Recently, a large debate concerning the forward guidance puzzle (FGP), starting from Del Negro et al. (2015), focused on precautionary savings and market incompleteness as a way to solve the puzzle (for instance, see McKay et al., 2016, 2017). On the contrary, Werning (2015) argues that market incompleteness per se is not sufficient to solve the FGP and shows that risk countercyclicality may exacerbate the puzzle. Bilbiie (2018) summarizes these arguments, showing that pro-cyclical inequality and/or pro-cyclical risk can solve the FGP. We claim that the interaction between income redistribution (following positive news) and marginal changes in conditional income volatility can mitigate the FGP, although only if our covariance sufficient statistic is positive.

1.2 Evidence on heteroskedasticity

A large empirical evidence supports the hypothesis of cross-sectional heteroskedasticity in the individual income process.[6] The clearest empirical contribution comes from the recent literature on non-parametric estimation of income processes.[7] In particular, Guvenen et al. (2021) show, across the earnings distribution, a striking evidence of heterogeneity in the standard deviation of the growth in log-earnings. They use a large panel dataset drawing a representative 10% sample of the U.S. population from the Social Security Administration (SSA), and measure the presence of nonlinearities and non-normalities in individual earnings dynamics over the life cycle. They document the existence of a significant empirical relation between age/percentiles of the income distribution and higher moments of the income process (i.e., standard deviation, skewness, kurtosis).

Figure 1, taken from Figure C.1 in Guvenen et al. (2021), suggests the existence of an asym-

[6] See Meghir and Pistaferri (2004).

[7] For the literatue on non-parametric estimation of the income process, see Arellano (2014), Arellano et al. (2017), Arellano and Bonhomme (2017), and Guvenen et al. (2014).

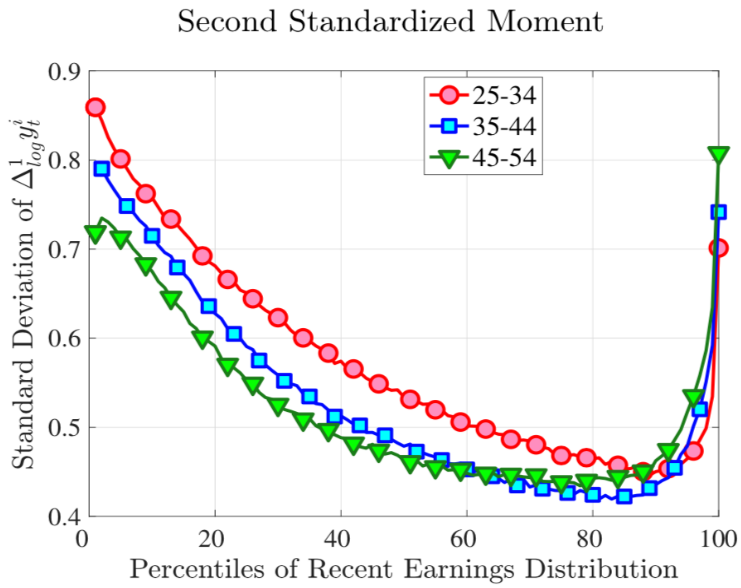

FIGURE 1: EVIDENCE ON HETEROSKEDASTICITY

Notes: The figure shows the dispersion of individual income growth across the earnings distribution. The standard deviation of one-year growth in log-earnings is plotted against the earnings distribution for different age groups (corresponding to different curves). The figure is taken from Guvenen et al. (2021).

metric U-shaped relation between the standard deviation of the one-year[8] log-earnings growth and the percentiles of lagged earnings distribution. That relation is decreasing and convex up to the 95^{th} percentile and sharply increasing for higher percentiles.

The rest of the paper is structured as follows. Section 2 presents the model. Section 3 describes the microfoundation of heteroskedasticity with incomplete information. Section 4 concludes.

2 A heteroskedastic New Keynesian model

Consider a continuum of measure 1 of agents/households in the interval $[0,1]$ indexed by $i \in [0,1]$. Let $y_{i,t}$ denote the income of agent i and assume that the earnings process can be characterized as the sum of two components:

$$y_{i,t} = y^d_{i,t} + \epsilon_{i,t}, \tag{1}$$

where $y^d_{i,t}$ and $\epsilon_{i,t}$ denote the deterministic and the stochastic components of income respectively. We make the following assumptions.

Assumption 1. *For every agent $i \in [0,1]$, the path of the deterministic component of income $\left(y^d_{i,t}\right)_{t=0}^{\infty}$ is known in every period $t \geq 0$.*

[8]Figure 3 in Guvenen et al. (2021) shows that the same relation holds also for five-year log-earnings growth. Here, we focus on the one-year evidence since it gives more weight to transitory components of the income shocks, in line with our modeling assumptions below.

Assumption 2. *For every agent* $i \in [0,1]$*, the process for the stochastic component is Gaussian, conditional on time* $t-1$ *information set* $\mathcal{I}_{t-1}$*, with its variance being a function of lagged deterministic income:*

$$\epsilon_{i,t} \mid \mathcal{I}_{t-1} \sim \mathcal{N}(0, \sigma_{i,t}^2), \quad \sigma_{i,t}^2 \coloneqq \sigma^2(y_{i,t-1}^d). \tag{2}$$

Assumption 3. *The random variables* $(\epsilon_{i,t})_{i\in[0,1]}$ *are conditionally stochastically independent:*

$$\epsilon_{i,t} \perp\!\!\!\perp \epsilon_{j,t} \mid \mathcal{I}_{t-1} \quad \forall i,j \in [0,1],\ i \neq j.$$

Assumption 3 is a standard requirement of cross-sectional independence for the stochastic components of the income process. Assumption 2 is the main requirement that we impose on the income process, consistently with the empirical evidence from Guvenen et al. (2021), shown in Figure 1, of cross-sectional heteroskedasticity in income. As in Guvenen et al. (2021), where the variance of the *unpredicatable* component varies over the income distribution, we assume that the unpredictable component $\epsilon_{i,t}$ of the income process is heteroskedastic.[9]

Note that the path of deterministic income being known by the agent (Assumption 1) and the conditional volatility of the stochastic component being a function solely of lagged deterministic earnings (Assumption 2) make the entire path of variances $\left(\sigma_{i,\tau}^2\right)_{\tau=t}^{\infty}$ known by the agent i at time t. In turn, these assumptions imply that the drift of the optimal consumption process is deterministic.

We argue that Assumption 2 allows for large gains in terms of analytical tractability without loss of generality. Alternatively, we may in fact consider an identical two-period version of our infinite-horizon CARA-Gaussian model, and assume that the variance of the income process is a function of the *entire* lagged income (not only of its deterministic component). Suppose that individual income is stochastic and normally distributed with mean y_t and variance $\sigma_{i,t}^2$: $y_{i,t} \mid \mathcal{I}_{t-1} \sim \mathcal{N}(y_t, \sigma_{i,t}^2)$. Then, one could model income volatility as: $\sigma_{i,t}^2 \coloneqq \sigma^2(y_{i,t-1})$. As we show in Section 3 this two-period version of the model admits a closed-form solution even under this relaxed parametric assumption on the overall income process. Moreover, we also prove (always in Section 3) that the implications of the infinite horizon and the two-period models are identical and that the main results of the paper can be derived (and coincide) in both frameworks.

2.1 Households

In every period t, each household can purchase a risk-free (real) bond $a_{i,t+1}$ (paying one unit of good in $t+1$) at the price $1/(1+r_t)$. Additionally, the household can consume a perishable consumption good $c_{i,t}$, using its total (net) income $y_{i,t}$ and its beginning-of-period wealth $a_{i,t}$.

Households maximize their intertemporal utility – discounted with a factor $\beta \in (0,1)$ – over an infinite-horizon, choosing consumption and savings (labor supply is inelastic) under the sequence of

[9]While Guvenen et al. (2021) look at log-earnings, without loss of generality, we focus on earnings in *levels*, since the CARA utility model needs to be specified in such a way for analytical purposes, as natural in the CARA utility literature. For example, see Caballero (1990), Wang (2002), Wang (2003), and Acharya and Dogra (2020).

budget constraints and the initial condition $a_{i,0}$ on bond holdings.[10] The problem can be formulated recursively with the Bellman operator:

$$V_t(a_{i,t}, y_{i,t}^d, \epsilon_{i,t}) := \max_{(c_{i,t}, a_{i,t+1})} \left\{ u(c_{i,t}) + \beta \cdot \mathbb{E}_t \left[V_{t+1}(a_{i,t+1}, y_{i,t+1}^d, \epsilon_{i,t+1}) \right] \right\}$$

$$\text{s.t.} \quad c_{i,t} + \frac{a_{i,t+1}}{1 + r_t} = y_{i,t} + a_{i,t},$$

$$a_{i,0} \quad \text{given,}$$

where $\mathbb{E}_t(\cdot) := \mathbb{E}(\cdot \mid \mathcal{I}_t)$ is the expectation operator conditional on time t information set and $V_t(\cdot)$ is the time t value function. In addition to the endogenous state $a_{i,t}$, the exogenous states $y_{i,t}^d$ and $\epsilon_{i,t}$ are relevant for the household's decision, since $y_{i,t}^d$ determines the volatility of $\epsilon_{i,t+1}$, while $\epsilon_{i,t}$ is the current realization of the income shock.

2.1.1 Solution to the consumer problem

Let $u(\cdot)$ be a CARA - exponential utility function: $u(c) = -\frac{1}{\gamma}\exp\{-\gamma c\}$ with γ being the coefficient of absolute risk aversion. The Euler equation derived from the utility maximization problem is:

$$\exp\{-\gamma c_{i,t}\} = \beta(1 + r_t) \cdot \mathbb{E}_t\big[\exp\{-\gamma c_{i,t+1}\}\big]. \tag{3}$$

The solution of the problem is characterized by equation (3), together with the sequence of budget constraints, the initial condition and the transversality condition on bonds:

$$\lim_{T \to \infty} \mathbb{E}_t\big[\beta^T u'(c_{i,T}) \cdot a_{i,T}\big] = 0.$$

A procedure to find a closed-from solution for consumption is proposed by Caballero (1990) and is based on a "guess and verify" of the stochastic process followed by optimal consumption.[11] We extend that approach to the more general case in which r_t and the variance of the income process are not constant over time, and $\beta \neq 1 + r_t$.

Proposition 1. *The optimal consumption plan follows the stochastic difference equation:*

$$c_{i,t} = c_{i,t+1} - \frac{1}{\gamma}\log\big[\beta(1 + r_t)\big] - \frac{\gamma \mu_{t+1}^2}{2}\sigma_{i,t+1}^2 - \mu_{t+1}\epsilon_{i,t+1}, \tag{4}$$

[10]The usual no-Ponzi-game condition applies as an additional constraint:

$$\lim_{T \to \infty} \prod_{t=0}^{T} \frac{1}{1 + r_t} a_{i,T+1} = 0.$$

[11]An alternative solution method is followed by Acharya and Dogra (2020). They guess a solution for consumption of the form $c_{i,t} = \mathcal{C}_t + \mu_t(y_{i,t} + a_{i,t})$ and find the value of the coefficients $\mathcal{C}_t, \mu_t$ recursively.

where μ_t, the MPC out of wealth, follows the recursion:

$$\mu_t = \frac{\mu_{t+1}(1+r_t)}{1+\mu_{t+1}(1+r_t)}. \tag{5}$$

Proof. See Appendix A.1.

Equation (4) is worth a few considerations. Notice that, relative to a standard permanent-income hypothesis (PIH) specification, the volatility term $(\gamma\mu_{t+1}^2/2)\cdot\sigma_{i,t+1}^2$ captures the effects of precautionary savings. Higher future income volatility depresses current consumption since agents have an incentive to save more as a form of partial insurance. In our setting, however, and unlike Acharya and Dogra (2020) who derive a similar equation in a homoskedastic CARA framework, the variance of future earnings that appears in (4) is *agent specific* since it is a function of current household income and it is, therefore, heterogeneous across agents.

2.2 Heterogeneity in MPCs

As in the corresponding homoskedastic CARA framework, MPCs in our model depend on current and future real interest rates. An increase in the real interest rate makes MPCs higher as a consequence of a stronger income effect (as can be seen from equation (5)). Consider a marginal increase in earnings: the household optimally saves part of it (or reduces its borrowing) obtaining higher future returns (since the interest rate has increased). These higher returns determine a positive income effect that optimally induces the agent to spend a larger fraction of the income shock (i.e., the MPC is higher).

In our heteroskedastic framework, however, there are two further implications for MPCs. First, there is a distinction between the MPC out of *wealth*, that coincides with μ_t (and with the MPC in the homoskedastic case) and the MPC out of a transitory *income* shock. Second, there exists a non-degenerate distribution of MPCs out of transitory income, since different positions over the earnings distribution affect future income volatility and the precautionary saving motive. We summarize these findings in the following proposition.

Proposition 2. *Assume that $\sigma^2:\mathbb{R}\to\mathbb{R}_+$ is differentiable with respect to lagged (deterministic) income. Then:*

(i) MPCs out of wealth are identical across agents and differ from MPCs out of transitory income if the marginal change in future income volatility out of current earnings is non-null, i.e., if $\partial\sigma_{i,t+1}^2/\partial y_{i,t}\neq 0$.

(ii) MPCs out of transitory income are heterogeneous across agents if the marginal change in future income volatility out of current earnings is non-constant over the earnings distribution, i.e., if $\partial\sigma_{i,t+1}^2/\partial y_{i,t}$ is non-constant.

In particular:

$$MPC_{i,t} := \frac{\partial c_{i,t}}{\partial y_{i,t}} = \mu_t - \frac{1}{1+\mu_{t+1}(1+r_t)} \frac{\gamma \mu_{t+1}^2}{2} \frac{\partial \sigma_{i,t+1}^2}{\partial y_{i,t}}. \tag{6}$$

Proof. See Appendix A.1.

Interestingly, our framework is able to combine three different characteristics of the heterogeneous agents literature in a tractable way: (i) precautionary savings (captured by the variance term appearing in equation (4)), (ii) heterogeneity in MPCs out of transitory income, and (iii) difference between MPCs out of wealth and income.[12] These results contrast with the standard homoskedastic framework, where MPCs out of income are identical across agents and coincide with the MPC out of wealth.

Notice that the empirical evidence from Figure 1 (from Guvenen et al., 2021) suggests that $\partial \sigma_{i,t+1}^2 / \partial y_{i,t}$ is both *negative* and *increasing* up to the 95$^{\text{th}}$ percentile. Consistent with this fact, and assuming that income volatility $\sigma_{i,t+1}^2$ is a decreasing and convex function of lagged income, it is also immediate to reproduce the finding (of some empirical literature on MPCs) whereby MPCs are decreasing over the earnings distribution.

Lemma 1. *If the variance of individual income $\sigma^2 : \mathbb{R} \to \mathbb{R}_+$ is twice differentiable and strictly convex in lagged income, then the marginal propensity to consume of agent i is decreasing in income:*

$$\frac{\partial MPC_{i,t}}{\partial y_{i,t}} < 0. \tag{7}$$

The result from Lemma 1 is easy to obtain from equation (6). In particular:

$$\frac{\partial MPC_{i,t}}{\partial y_{i,t}} = -\frac{1}{1+\mu_{t+1}(1+r_t)} \frac{\gamma \mu_{t+1}^2}{2} \cdot \underbrace{\frac{\partial^2 \sigma_{i,t+1}^2}{(\partial y_{i,t})^2}}_{>0} < 0.$$

Intuitively, if the variance of future income is *decreasing* and *convex* in current income, a marginal increase in current income has the effect, at the margin, of weakening the precautionary saving motive relatively more for poor than rich households. Thus agents at the bottom percentiles of the income distribution feature a higher MPC and, at the margin, increase consumption relatively more than agents at the top percentiles.

2.3 Aggregation and general equilibrium

We aggregate consumption integrating with respect to the Lebesgue measure. Let $c_t := \int_{[0,1]} c_{i,t} \, di$ and let us assume that the function $\sigma^2 : \mathbb{R} \to \mathbb{R}_+$ is measurable and integrable. Then, integrating

[12] MPCs out of housing wealth (see Guren et al., 2021) and out of financial wealth (see Chodorow-Reich et al., 2021) are estimated to be 2.67 cents and 3.2 cents per dollar, respectively, i.e., significantly smaller than usual estimates of MPCs out of income.

the recursion for agents' optimal consumption (equation (4)), we obtain the aggregate consumption function:

$$c_t = c_{t+1} - \frac{1}{\gamma}\log[\beta(1+r_t)] - \frac{\gamma\mu_{t+1}^2}{2}\int_{[0,1]}\sigma_{i,t+1}^2 di. \tag{8}$$

Note that the previous result can be obtained by a continuous version of the Law of Large Numbers (LLN)[13] since the integral of the income process innovation can be thought of as the limit of the means, converging to the population average:

$$\int_{[0,1]}\epsilon_{i,t+1}\ di = 0 = \mathbb{E}_t(\epsilon_{i,t+1}).$$

The economy has a supply-side that produces aggregate output y_t and a government that purchases part of this output, through government spending g_t.[14] Therefore, the goods market clearing requires $y_t = c_t + g_t$. Imposing this condition, we obtain a dynamic aggregate demand (AD) equation:

$$y_t = y_{t+1} - \frac{1}{\gamma}\log[\beta(1+r_t)] - \frac{\gamma\mu_{t+1}^2}{2}\int_{[0,1]}\sigma_{i,t+1}^2\ di + g_t - g_{t+1}. \tag{9}$$

In addition, the real and the nominal side of the economy are assumed to be generically linked via a Phillips-curve relation between current inflation π_t, future inflation, and output (implicitly, through the function $P(\cdot)$):

$$P(\pi_t, \pi_{t+1}, y_t) = 0. \tag{10}$$

Finally, the monetary authority follows a given (active) monetary policy rule. We consider two different specifications for monetary policy.

1. *Interest rate peg.* The monetary authority pegs the real interest rate to its steady-state value r:

$$r_t = r + v_t, \tag{11}$$

   where v_t is a stochastic process characterizing exogenous monetary shocks.

2. *Taylor rule.* The monetary authority sets the nominal interest rate i_t as function of current inflation, through the feedback Taylor coefficient $\phi_\pi > 0$:

$$1 + i_t = (1+r)\cdot(1+\pi_t)^{\phi_\pi}\exp(v_t). \tag{12}$$

2.3.1 Equilibrium

Given a process for $(g_t, v_t)_{t=0}^\infty$, an equilibrium is an allocation $(y_t)_{t=0}^\infty$, a vector of prices $(i_t, \pi_t, r_t)_{t=0}^\infty$ and a sequence of MPCs out of wealth $(\mu_t)_{t=0}^\infty$ such that:

- the dynamic AD – equation (9) – holds;

[13] See Uhlig (1988), Proposition 1.
[14] See Appendix A.2.

- the MPCs out of wealth recursion – equation (5) – holds;
- the PC – equation (10) – holds;
- monetary policy follows a given monetary policy rule, either the peg rule (equation (11)) or the Taylor rule (equation (12));
- the Fisher equation holds:

$$1 + r_t = \frac{1 + i_t}{1 + \pi_{t+1}}.$$

2.4 Sufficient statistics

We are now ready to establish one of our key results. We show that the equilibrium response of output to alternative demand shocks can be fully described through a small set of observable sufficient statistics, capturing the heterogeneous consumption choices of different agents. Before turning to a more general characterization under price rigidity, we assume below that monetary policy keeps a constant real interest rate (up to an exogenous shock v_t).[15]

Proposition 3. *The equilibrium response of aggregate output to either government spending or real interest rate shocks is captured by the following cross-sectional sufficient statistics:*

(i) Expected marginal change in the income volatility out of individual earnings:

$$\underbrace{\mathbb{E}\left(\frac{\partial \sigma^2_{i,t+1}}{\partial y_{i,t}}\right)}_{\textit{individual risk effect}}.$$

(ii) Expected (individual) earnings sensitivity to aggregate output:

$$\underbrace{\mathbb{E}\left(\frac{\partial y_{i,t}}{\partial y_t}\right)}_{\textit{income cyclicality effect}}.$$

(iii) Covariance between the marginal change in (individual) income volatility out of (individual) earnings and the (individual) earnings sensitivity to aggregate output:

$$\underbrace{Cov\left(\frac{\partial \sigma^2_{i,t+1}}{\partial y_{i,t}}, \frac{\partial y_{i,t}}{\partial y_t}\right)}_{\textit{risk-cyclicality covariance effect}}.$$

[15] Henceforth we also assume that it is possible to differentiate the variance function with respect to lagged income and to exchange the derivative and the integral.

Moreover, the output multiplier of a fully transitory government spending shock is:

$$\Theta_t := \frac{\partial y_t}{\partial g_t} = \left\{1 + \frac{\gamma \mu_{t+1}^2}{2}\left[\mathbb{E}\left(\frac{\partial \sigma_{i,t+1}^2}{\partial y_{i,t}}\right) \cdot \mathbb{E}\left(\frac{\partial y_{i,t}}{\partial y_t}\right) + Cov\left(\frac{\partial \sigma_{i,t+1}^2}{\partial y_{i,t}}, \frac{\partial y_{i,t}}{\partial y_t}\right)\right]\right\}^{-1}, \tag{13}$$

while the output multiplier of a fully transitory real interest rate shock is:

$$\frac{\partial y_t}{\partial r_t} = -\Theta_t \cdot \frac{1}{\gamma(1 + r_t)}. \tag{14}$$

Proof. See Appendix A.1.

Consider statistics (i)-(iii). All combined they capture how individual *risk* varies over the business cycle and how this effect is heterogeneous across the income distribution.

Statistic (i) captures the expected marginal effect of realized income on individual risk (*individual risk effect*). Statistic (ii) captures the degree of cyclicality of individual income with respect to aggregate income (*income cyclicality effect*). Finally, statistic (iii) describes how the effect on individual risk co-varies with the degree of income cyclicality (*risk-cyclicality covariance effect*). Hence, the size of both the government spending and real interest rate multipliers depends on the interaction between the cross-sectional variation in the precautionary saving motive and how the distribution of earnings is shaped by aggregate output.

Consider the (product of the) first two sufficient statistics in (13), i.e., (i) the individual risk effect and (ii) the income cyclicality effect. Suppose that a negative demand shock decreases output on impact and that, in turn, individual earnings decline. If, *on average*, individual risk is decreasing in individual income (so that the product of (i) and (ii) is negative), following the negative shock, agents face *more risk* on average and therefore feature a larger precautionary saving motive. Therefore, they decrease consumption relatively more than if average risk were unaffected.

However, the latter precautionary savings response is *not homogeneous* across agents. This can be gauged from statistic (iii), which captures the risk-cyclicality covariance effect. This statistic describes the interaction between how individual risk varies at the margin and how sensitive earnings are to aggregate output over the income distribution. Intuitively, consider a negative demand shock. As aggregate income decreases, individual earnings may react in the cross-section in a heterogeneous fashion. Then, changes in individual earnings affect future income volatility, and in turn precautionary savings and individual MPCs. In turn, the heterogeneous consumption response at the margin feeds back into aggregate demand. For instance, suppose that idiosyncratic risk is decreasing in individual income, i.e., $\partial \sigma_{i,t+1}^2 / \partial y_{i,t} < 0$. If those agents whose income is more sensitive to aggregate income are also those for whom risk increases (at the margin) relatively more (following a decrease in their individual earnings), then the covariance term is negative, leading to amplification. In other words, amplification of the negative demand shock requires that those agents whose income is more sensitive to aggregate income are also those with a stronger increase in the precautionary saving motive and higher MPCs.

2.4.1 Homoskedasticity: a special case

It is useful to compare our results so far to a standard setting with homoskedasticity (i.e., assuming that the volatility of the income process is common for all the households). Under (cross-sectional) homoskedasticity, let us denote the common volatility of the income process as σ_t^2. The dynamic AD equation then becomes:

$$y_t = y_{t+1} - \frac{1}{\gamma}\log[\beta(1+r_t)] - \frac{\gamma\mu_{t+1}^2}{2}\sigma_{t+1}^2 + g_t - g_{t+1}.$$

Since σ_{t+1}^2 is identical across the earnings distribution, changes in individual income (deriving from changes in output) neither affect risk, nor precautionary savings, nor MPCs at the margin. Therefore:

$$\mathbb{E}\left(\frac{\partial\sigma_{i,t+1}^2}{\partial y_{i,t}}\right) = Cov\left(\frac{\partial\sigma_{i,t+1}^2}{\partial y_{i,t}}, \frac{\partial y_{i,t}}{\partial y_t}\right) = 0,$$

and the output multiplier of an aggregate demand shock becomes, under homoskedasticity:

$$\Theta_t^{hom} := \frac{\partial y_t}{\partial g_t} = 1.$$

Hence, depending on the sign of the three sufficient statistics outlined above, the aggregate output multiplier of a temporary demand shock under heteroskedasticity, Θ_t, could either exceed or not the multiplier under homoskedasticity, $\Theta_t^{hom} = 1$. Notice also that in the homoskedastic case the output multiplier of an aggregate demand shock coincides with the one in a representative-agent economy.[16]

2.5 Results in the general model

Let us turn to a more general analysis of the first-order effects of fiscal (government spending) and monetary shocks. Let us suppose that monetary policy follows the Taylor rule specification of equation (12).

2.5.1 Steady state

We can approximate the model around a zero-inflation steady state. We set government spending in the steady state equal to zero and we normalize aggregate consumption and output to 1. Hence:

$$y = c = 1, \quad g = 0, \quad \pi = 0.$$

[16] That is equal to 1. See Woodford (2011).

Let $F : \mathbb{R} \to [0,1]$ be the cumulative density function describing the stationary distribution of (deterministic) income, y_i^d.[17] Under some additional assumptions, we can prove that:

$$\bar{\sigma}^2 := \int_{\mathbb{R}} \sigma^2(y)\ dF(y) = \int_{[0,1]} \sigma^2(y_i^d)\ di, \tag{15}$$

where $\bar{\sigma}^2$ is defined as the average income volatility in the steady state. Moreover, the steady-state value for the real interest rate is determined from the dynamic AD equation (9):

$$\left(\frac{1+r}{r}\right)^2 \log[\beta(1+r)] = -\frac{\gamma^2}{2}\bar{\sigma}^2. \tag{16}$$

Assuming that $\bar{\sigma}^2 > 0$,[18] the solution to this equation exists and is unique.

Given the steady-state value for the real interest rate, the steady-state MPC out of wealth is determined from the MPC recursion, described in equation (5):

$$\mu = \frac{r}{1+r}. \tag{17}$$

Finally, let us define Γ as a function of the sufficient statistics (i)-(iii) in the steady state:

$$\Gamma := \int_{[0,1]} \left(\frac{\partial \sigma_i^2}{\partial y_i}\frac{\partial y_i}{\partial y}\right) di = \mathbb{E}\left(\frac{\partial \sigma_i^2}{\partial y_i}\right) \cdot \mathbb{E}\left(\frac{\partial y_i}{\partial y}\right) + Cov\left(\frac{\partial \sigma_i^2}{\partial y_i}, \frac{\partial y_i}{\partial y}\right). \tag{18}$$

Notice that $\Gamma = 0$ holds in the two limit cases of a representative-agent and in the homoskedastic version of our model. Details and proofs related to the steady state are in Appendix A.3.

2.5.2 Log-linear representation

We take a log-linear approximation around the previously defined steady state. A *hat* notation denotes variables in deviation from their steady-state values: $\hat{x}_t := \frac{x_t - x}{x}$ and $\hat{g}_t := g_t$. Additionally, suppose that the Phillips Curve admits a standard forward-looking linear approximation (with k being its slope).[19] The linearized model can be described by the following four equations.[20]

$$\hat{y}_t = \Theta \cdot \left[\hat{y}_{t+1} - \frac{1}{\gamma} \cdot (\hat{i}_t - \hat{\pi}_{t+1}) - \Lambda \cdot \hat{\mu}_{t+1} + (\hat{g}_t - \hat{g}_{t+1})\right], \tag{19a}$$

$$\hat{\mu}_t = \tilde{\beta} \cdot [\hat{\mu}_{t+1} + (\hat{i}_t - \hat{\pi}_{t+1})], \tag{19b}$$

$$\hat{\pi}_t = \tilde{\beta} \cdot \hat{\pi}_{t+1} + k \cdot \hat{y}_t, \tag{19c}$$

$$\hat{i}_t = \phi_\pi \cdot \hat{\pi}_t + v_t, \tag{19d}$$

[17] We can always assume that such a function exists in the steady state.

[18] Equivalently, assuming that a set of agents with non-zero measure has strictly positive variances in the steady state.

[19] We linearize the PC obtained under convex price adjustment costs as in Rotemberg (1982). See Appendix A.2.4.

[20] Note that the log-linear Fisher equation $\hat{r}_t = \hat{i}_t - \hat{\pi}_{t+1}$ has already been replaced into the other equilibrium equations.

where:

$$
\begin{aligned}
&\Theta := \left(1 + \frac{\gamma\mu^2}{2}\Gamma\right)^{-1} > 0, \\
&\Lambda := \gamma \cdot \mu^2 \cdot \bar{\sigma}^2 > 0, \\
&\tilde{\beta} := \frac{1}{1+r}.
\end{aligned}
\tag{20}
$$

Equation (19a) is the linearized version of the dynamic AD equation, equation (19b) is the linearized MPC recursion, equation (19c) is the linearized PC, and equation (19d) is the linearized Taylor rule. To conclude the description of the equilibrium, we assume that government spending and monetary shocks follow an AR(1) process:

$$
\begin{aligned}
\hat{g}_t &= \rho_g \cdot \hat{g}_{t-1} + \epsilon_t^g, \\
v_t &= \rho_v \cdot v_{t-1} + \epsilon_t^v,
\end{aligned}
$$

with $|\rho_g| < 1$, $|\rho_v| < 1$ and ϵ_t^g, ϵ_t^v being white noise serially uncorrelated processes, orthogonal between them, i.e., $\epsilon_t^g \perp \epsilon_t^v$.

Notice that parameter Λ in equation (19a) captures the (direct) marginal effect on output of changes in the MPC out of wealth. As in the CARA-homoskedastic model of Acharya and Dogra (2020), monetary policy is able to affect MPCs through interest rates. In addition, MPCs affect the pass-through of idiosyncratic income risk to consumption risk. Hence, a higher real interest rate increases the MPC and reduces output since income risk translates into a greater consumption risk (a larger fraction of income is spent and consumption inherits "more" of the volatility of income).

Heteroskedasticity and discounting/compounding

A key feature that distinguishes our framework is the presence, in the AD equation (19a), of the parameter Θ, given by the inverse of a linear transformation of Γ, in turn a function of the sufficient statistics outlined in Proposition 3. Θ also represents the steady-state output multiplier of government spending of Proposition 3. Notice that Θ augments the standard representative-agent AD equation with a *discount/compound* factor for future expectations (about $\hat{y}_{t+1}, \hat{g}_{t+1}$) and current shocks ($\hat{r}_t, \hat{g}_t$). Parameter Θ is greater (less) than $\Theta_t^{hom} = 1$ if the average marginal change in risk and/or the covariance statistics are negative (positive).[21] We explore below the implications of this compounding/discounting factor for two main issues: equilibrium determinacy and the size of output multipliers.

[21] Recall that the income cyclicality effect captured by statistic (ii) is always positive.

2.5.3 Equilibrium determinacy

The presence of the parameter Θ in (19a) affects the conditions for equilibrium determinacy, since it imposes restrictions on the parameter subspace where the model features a unique non-explosive local equilibrium solution. In general we can show that with larger values of Θ the conditions for equilibrium determinacy become more stringent. The reason is that expectations about future output get compounded more when Θ is high and equilibrium solutions tend to become explosive. Below we treat the issue of determinacy separately for the two specifications of monetary policy: an interest rate peg and a Taylor rule.

Determinacy with an interest rate peg

Let us consider the simple case of a monetary authority that follows an interest rate peg, i.e., $\hat{r}_t = 0$, and let us assume that prices are fully fixed, i.e., $\hat{\pi}_t = 0$ for every $t \geq 0$. Then, the equilibrium is described uniquely by the AD equation (19a):

$$\hat{y}_t = \Theta \cdot \left[\hat{y}_{t+1} + (\hat{g}_t - \hat{g}_{t+1}) \right]. \tag{21}$$

Equilibrium determinacy requires AD-discounting: $\Theta < 1$. In fact, when expectations about future output are discounted, the effect of future fiscal shocks or future deviations of output from its steady-state level have a smaller impact on current output the further they occur in the future. Consequently, when iterating forward to solve for current output:

$$\hat{y}_t = \lim_{T \to \infty} \Theta^T \cdot \hat{y}_{t+T} + \sum_{k=1}^{\infty} \Theta^k \cdot (\hat{g}_{t+k-1} - \hat{g}_{t+k}).$$

If $\Theta < 1$, then the limit term is null and the equilibrium path of output $(\hat{y}_t)_{t \geq 0}$ is bounded.[22] On the contrary, if $\Theta > 1$, then the limit term diverges and the solution for current output is explosive.

Determinacy under a Taylor rule

Turning to the Taylor rule specification for monetary policy, determinacy requires stronger restrictions to the parameter space when Θ is larger. Again, the greater Θ, the larger the compounding effect of future expectations (or the less they are discounted) back to current output and the stronger the reaction of monetary policy must be to stabilize the economy. In particular, for equilibrium determinacy, the inflation feedback coefficient ϕ_π has to rise with higher values of Θ, meaning that the monetary authority needs to overreact to inflation with a larger increase in the nominal interest rate.

[22] We also need to postulate that the sequence of changes in government spending is bounded. Let $\Delta \hat{g}_{t+1} := \hat{g}_{t+1} - \hat{g}_t$, then, we assume that:

$$\sup_{k \geq 1} |\Delta \hat{g}_{t+k}| < \infty.$$

We reduce the dimensionality of the model from four to three equations by substituting (19d) in the other three equations of the linear system (19) and rewrite it in matrix form as below:

$$\begin{pmatrix} \hat{y}_{t+1} \\ \hat{\mu}_{t+1} \\ \hat{\pi}_{t+1} \end{pmatrix} = \underbrace{\begin{bmatrix} \Theta^{-1} + \frac{k}{\tilde{\beta}}(\gamma^{-1} - \Lambda) & \frac{\Lambda}{\tilde{\beta}} & (\gamma^{-1} - \Lambda)\left(\phi_\pi - \frac{1}{\tilde{\beta}}\right) \\ -\frac{k}{\tilde{\beta}} & \frac{1}{\tilde{\beta}} & \frac{1}{\tilde{\beta}} - \phi_\pi \\ -\frac{k}{\tilde{\beta}} & 0 & \frac{1}{\tilde{\beta}} \end{bmatrix}}_{\Omega :=} \begin{pmatrix} \hat{y}_t \\ \hat{\mu}_t \\ \hat{\pi}_t \end{pmatrix} + \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \Delta\hat{g}_{t+1} + \begin{bmatrix} \gamma^{-1} - \Lambda \\ -1 \\ 0 \end{bmatrix} v_t, \tag{22}$$

where $\Delta\hat{g}_{t+1} = \hat{g}_{t+1} - \hat{g}_t$.

With this specification there are three forward-looking variables, $\hat{y}_t, \hat{\mu}_t, \hat{\pi}_t$, hence determinacy requires the matrix Ω to have three eigenvalues outside the unit circle. See Appendix A.4.

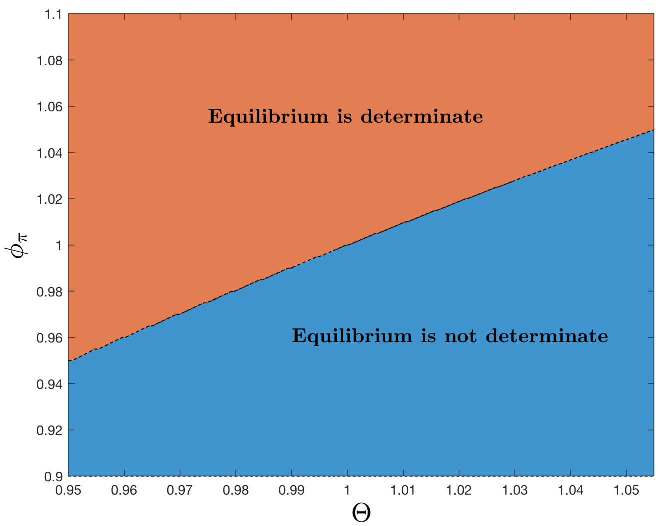

FIGURE 2: DETERMINACY REGION: Θ - ϕ_π.

Figure 2 plots the Taylor-rule coefficient ϕ_π as a function of Θ. The area depicted in red corresponds to the region of the parameter space where the equilibrium is determinate. Notice that in the case $\Theta = 1$, corresponding to either the representative-agent or the homoskedasticity cases, the Taylor principle, i.e., $\phi_\pi > 1$, is both sufficient and necessary for determinacy. More generally, let $\bar{\phi}_\pi$ be the threshold for ϕ_π such that for every $\phi_\pi > \bar{\phi}_\pi$ the model has a determinate equilibrium. Then, $\bar{\phi}_\pi$ is increasing in Θ.

Summarizing, the presence in the AD equation (19a) of either compounding or discounting (through the parameter Θ) modifies the extent to which the Taylor principle ($\phi_\pi > 1$) is sufficient for local determinacy. In particular *compounding* – $\Theta > 1$ – requires monetary policy to react more strongly to an increase in output ($\bar{\phi}_\pi$ has to be greater than 1), while *discounting* – $\Theta < 1$ – allows for a weaker reaction of monetary policy (ϕ_π can be less than 1).

2.5.4 Analytical solution for multipliers

The model is tractable enough that we can easily derive analytical solutions. We solve for the impulse response functions (IRFs) to fiscal and monetary shocks through the method of undetermined coefficients. Let us suppose that the solution for output, inflation and MPCs follows the guess:

$$\hat{x}_t = \psi_x^u \cdot u_t \quad \text{where} \quad u \in \{g, v\}, \quad x \in \{y, \mu, \pi\}.$$

Substituting the guess in the system of equations (19) (substituting also for the Taylor rule), rearranging, and matching coefficients, we obtain the solution for the reduced form elasticities to both shocks.[23] The general expressions for the multipliers out of persistent shocks can be found in Appendix A.5.

All reduced-form multipliers become analytically transparent when fiscal and monetary shocks are white noise, i.e., $\rho_g = \rho_v = 0$. In this case, we have, for the *fiscal shock*:

$$\begin{aligned} \psi_y^g &= \frac{\gamma\Theta}{\gamma + \Theta \cdot \phi_\pi \cdot k} > 0, \\ \psi_\mu^g &= \phi_\pi \cdot \tilde{\beta} k \cdot \psi_y^g > 0, \\ \psi_\pi^g &= k \cdot \psi_y^g > 0, \end{aligned} \tag{23}$$

and for the *monetary shock*:

$$\begin{aligned} \psi_y^v &= -\frac{\Theta}{\gamma + \Theta \cdot \phi_\pi \cdot k} < 0, \\ \psi_\mu^v &= \tilde{\beta} + \phi_\pi \cdot \tilde{\beta} k \cdot \psi_y^v > 0, \\ \psi_\pi^v &= k \cdot \psi_y^v < 0. \end{aligned} \tag{24}$$

A positive shock to government spending increases output and inflation. Monetary policy responds by rising the nominal, and therefore the real interest rate, leading to a rise in MPCs. Higher MPCs have an offsetting contractionary effect on output since they amplify the pass-through of income to consumption risk. As for the monetary policy multiplier, a contractionary increase in the nominal interest rate reduces output and inflation, and causes the MPC to increase. Most importantly, the coefficients for output and inflation of fiscal and monetary shocks $(\psi_y^g, \psi_\pi^g, \psi_y^v, \psi_\pi^v)$ are *increasing* in Θ in absolute value, consistent with the fact that higher values of Θ determine a general equilibrium amplification of demand shocks.

Note, again, that by imposing $\Theta = 1$ one obtains the IRFs for the model with homoskedasticity. It is therefore clear that $\Theta > 1$ generates amplification, while $\Theta < 1$ generates dampening of both

[23] We assume that $\phi_\pi > \rho_v, \rho_g$ and that $\gamma \cdot \Lambda < 1$. Both assumptions hold for our and, more generally, for empirically reasonable calibrations.

shocks:

$$\text{fiscal shock:}\quad |\psi_y^g(\Theta > 1)| > \underbrace{\frac{\gamma}{\gamma + \cdot\phi_\pi \cdot k}}_{\text{homosked.}} > |\psi_y^g(\Theta < 1)|,$$

$$\text{monetary shock:}\quad |\psi_y^v(\Theta > 1)| > \underbrace{\frac{1}{\gamma + \cdot\phi_\pi \cdot k}}_{\text{homosked.}} > |\psi_y^v(\Theta < 1)|.$$

2.5.5 Calibration

We can visualize the above analytical results by plotting the IRFs to a 1% annual increase in government spending and in the nominal interest rate, for three different values of Γ (hence of Θ). We calibrate the model to an annual frequency. Following Acharya and Dogra (2020), the coefficient of absolute risk aversion, $\gamma = -\frac{u''(c)}{u'(c)}$, is set equal to 3 and the slope of the Phillips Curve, k, is set equal to 0.1. According to the volatility estimates by Storesletten et al. (2004) and following Acharya and Dogra (2020), we assign a value of 0.5 to the steady-state average income volatility $\bar{\sigma}^2$. We set the discount factor β equal to 0.96 and a common autoregressive coefficient for the fiscal and monetary shock $\rho_g = \rho_v = \rho = 0.8$. Additionally, we assign a value of 1.5 to the Taylor rule coefficient ϕ_π (so that the equilibrium is locally determinate). See Table 1.

We can then derive the remaining parameters as follows. The real interest rate is computed from equation (16) for given values of $\bar{\sigma}^2, \beta, \gamma$. Then, the steady-state MPC out of wealth μ is immediately determined from (17), while Λ and $\tilde{\beta}$ are obtained from equation (20). Finally, the value for Θ depends on the underlying assumptions for Γ and is found through equation (20).

The dotted black line in Figures 3 and 4 denotes the homoskedastic benchmark, while the dashed red line represents the IRFs for $\Theta > 1$ (amplification scenario) and the solid blue line for $\Theta < 1$ (dampening scenario).

Parameter		**Value**
Coefficient of absolute risk aversion	γ	3
Discount factor	β	0.96
Slope of PC	k	0.1
Average income volatility	$\bar{\sigma}^2$	0.5
Taylor rule coefficient	ϕ_π	1.5
Persistence of shocks	ρ	0.8

Table 1: Calibrated Parameters

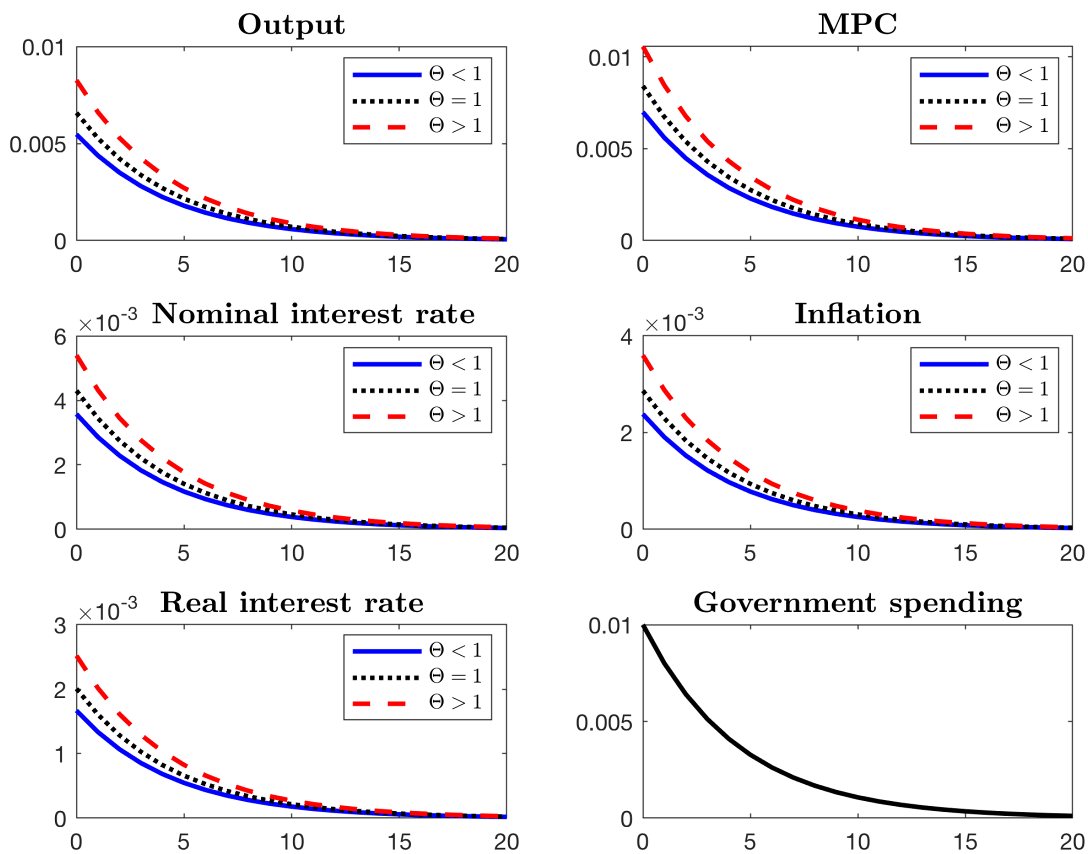

Figure 3: Impulse Responses to a Positive Government Spending Shock.

Notes: All variables are in % deviation from their steady-state value.

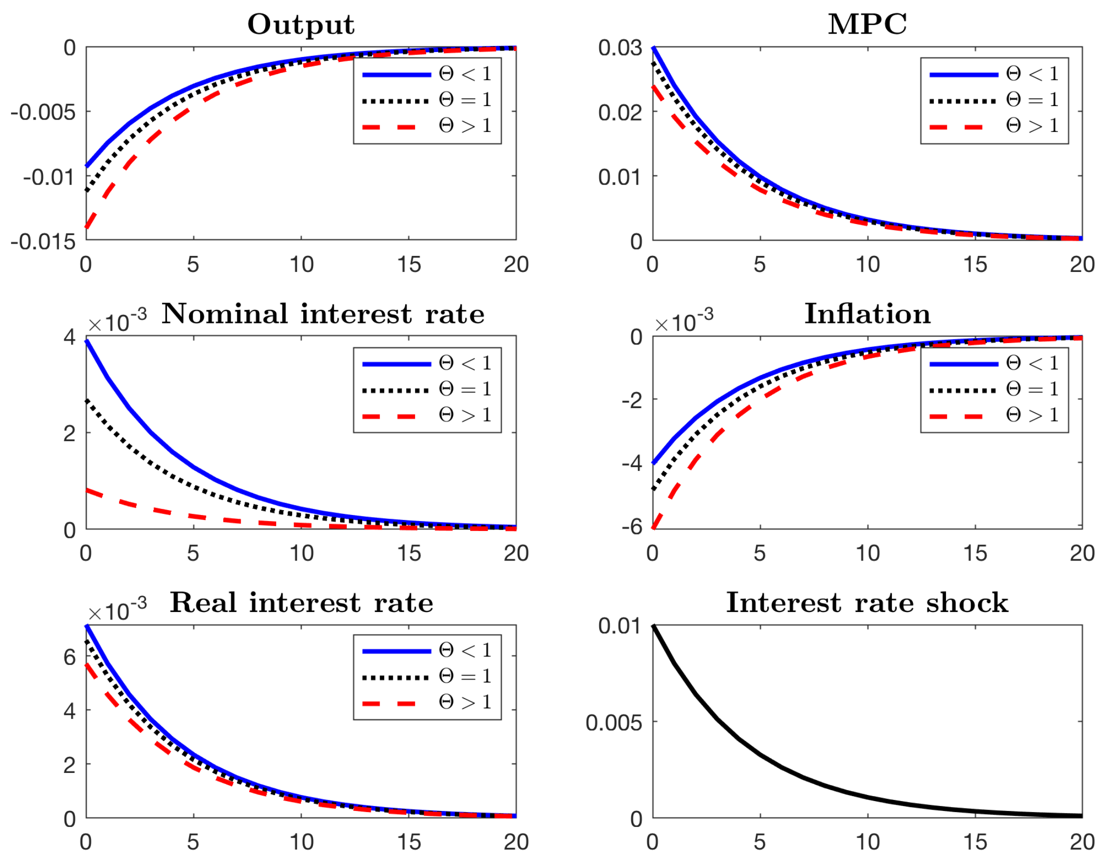

Figure 4: Impulse Responses to a Positive Interest Rate Shock.

Notes: All variables are in % deviation from their steady-state value.

2.6 Forward guidance puzzle

The expression *forward guidance puzzle* (FGP) (first used by Del Negro et al., 2015) denotes the puzzling prediction of standard New Keynesian models whereby changes in the nominal interest rate are more effective (on current output) the further they occur in the future. In this vein, news about monetary easing years ahead are at least as effective to boost economic activity as a cut in the current short-term interest rate. This puzzle occurs since agents perfectly anticipate the change in monetary policy and react by perfectly smoothing consumption over time. In addition, the positive feedback effect of monetary easing on output and inflation further reduces the real interest rate, amplifying the effects on current output of the announcement.

In our model, the FGP puzzle can be either worsened or solved, depending on the values of the sufficient statistics, that in turn determine the value of the parameters Γ and Θ (in the log-linear version, system of equations (19)).[24] Let us assume, for simplicity, that prices are fully fixed, i.e., $\hat{\pi}_t = 0$ for every $t \geq 0$ and consider a (credible) policy announcement of a nominal interest rate cut H periods ahead: $\hat{r}_{t+H} = \hat{i}_{t+H} = -\epsilon < 0$. We can rewrite the AD and the MPC recursions (equations (19a) and (19b)), iterating forward for H periods.

$$\hat{y}_t = \Theta^{H+1} \cdot \hat{y}_{t+H+1} - \frac{\Theta^{H+1}}{\gamma} \cdot \hat{i}_{t+H} - \Lambda \cdot \sum_{k=1}^{H+1} \Theta^k \hat{\mu}_{t+k}, \tag{25a}$$

$$\hat{\mu}_{t+k} = \tilde{\beta}^{H+1-k} \cdot \hat{\mu}_{t+H+1} + \tilde{\beta}^{H+1-k} \hat{i}_{t+H} \quad \text{for} \quad k = 0, ..., H. \tag{25b}$$

Additionally, using the same iterations, it is easy to show that $\hat{y}_{t+H+1} = \hat{\mu}_{t+H+1} = 0$. Substituting (25b) in (25a), we get:

$$\hat{y}_t = -\left[\frac{\Theta^{H+1}}{\gamma} + \Lambda \cdot \tilde{\beta}^{H+1} \sum_{k=1}^{H+1} \left(\frac{\Theta}{\tilde{\beta}}\right)^k\right] \cdot \hat{i}_{t+H} = \underbrace{\left[\frac{\Theta^{H+1}}{\gamma} + \Lambda \cdot \Theta^{H+1} \sum_{s=0}^{H} \left(\frac{\tilde{\beta}}{\Theta}\right)^s\right]}_{\Pi_H :=} \cdot \epsilon. \tag{26}$$

Let

$$\Phi_H := \left[\frac{\Theta^{H+1}}{\gamma} + \Lambda \cdot \Theta^{H+1} \sum_{s=0}^{H} \left(\frac{\tilde{\beta}}{\Theta}\right)^s\right] \tag{27}$$

be the *forward guidance (FG) multiplier*, expressed as a function of the time horizon H in which the interest rate cut will take place. Then, we can state that the FGP is "solved" when Φ_H is decreasing in H, in the sense that the further in the future the monetary easing is expected to occur, the less effective it will be on current output.

The formula for the FG multiplier can be decomposed in two parts, respectively describing two different channels through which future shocks affect current output.

- *Direct compounding/dampening* (direct C/D henceforth) channel, captured by the first term

[24] Recall that $\Gamma > 0 \iff \Theta < 1$.

in expression (27):

$$\frac{\Theta^{H+1}}{\gamma}.$$

This component describes the effects of future shocks on current output, *for a given path of MPCs.*

- *Change in MPCs* channel, captured by the second term in expression (27):

$$\Lambda \cdot \Theta^{H+1} \sum_{s=0}^{H} \left(\frac{\tilde{\beta}}{\Theta} \right)^{s}.$$

  This component describes the effects of future shocks on current output *through variations in the path of MPCs* (recall that in our model MPCs depend on current and future interest rates).

To study under which subset of the parameter space the FGP is either exacerbated or attenuated, we can take the derivative of Φ_H with respect to H and study its sign:[25]

$$\frac{\partial \Phi_H}{\partial H} = \Bigg[\underbrace{\log(\Theta) \cdot \frac{\Theta^{H+1}}{\gamma}}_{\text{derivative of 1}^{\text{st}} \text{ term in (27)}} + \underbrace{\frac{\Lambda\Theta}{\Theta - \tilde{\beta}} \left[\log(\Theta) \cdot \Theta^{H+1} - \log(\tilde{\beta}) \cdot \tilde{\beta}^{H+1} \right]}_{\text{derivative of 2}^{\text{nd}} \text{ term in (27)}} \Bigg].$$

In Figure 5 we show the contribution of each of the two channels (dashed and dotted lines for the direct C/D and the change in MPCs channels, respectively) and their cumulative effect on the FG multiplier (solid line) for different horizons H of the interest rate shock. We distinguish three scenarios for three different values of Θ (equivalently, of Γ, i.e., of our sufficient statistics). In the top panel we present the *discounting* case ($\Theta < 1$), in the middle one the homoskedastic benchmark ($\Theta = 1$) and in the bottom one the *compounding* scenario ($\Theta > 1$).

Notice that if $\Theta \geq 1$, then the second term in Φ_H is always increasing in H, while the first term is constant and increasing in H for $\Theta = 1$ and $\Theta > 1$, respectively. Therefore $\Theta \geq 1$ is a sufficient condition to exacerbate the FGP, as it can be seen from the bottom and middle panels in Figure 5.

Intuitively, the second term in Φ_H captures the impact of news on current output through changes in the path of MPCs. In fact, as described by Acharya and Dogra (2020), news about future interest rate cuts reduce MPCs. Hence, since they weaken the pass-through from income to consumption risk, they have the effect of stimulating current demand. In addition, in every period, if $\Theta > 1$ the effect of a decline in MPCs is further amplified (through the demand amplification

[25] The summation in Φ_H can be rewritten as:

$$\Lambda \cdot \Theta^{H+1} \sum_{s=0}^{H} \left(\frac{\tilde{\beta}}{\Theta} \right)^{s} = \Lambda \cdot \Theta^{H+1} \frac{1 - \frac{\tilde{\beta}^{H+1}}{\Theta^{H+1}}}{1 - \frac{\tilde{\beta}}{\Theta}} = \frac{\Lambda\Theta}{\Theta - \tilde{\beta}} \left[\Theta^{H+1} - \tilde{\beta}^{H+1} \right].$$

In addition, notice that when the derivative of Φ_H with respect to H is positive, the FGP is exacerbated.

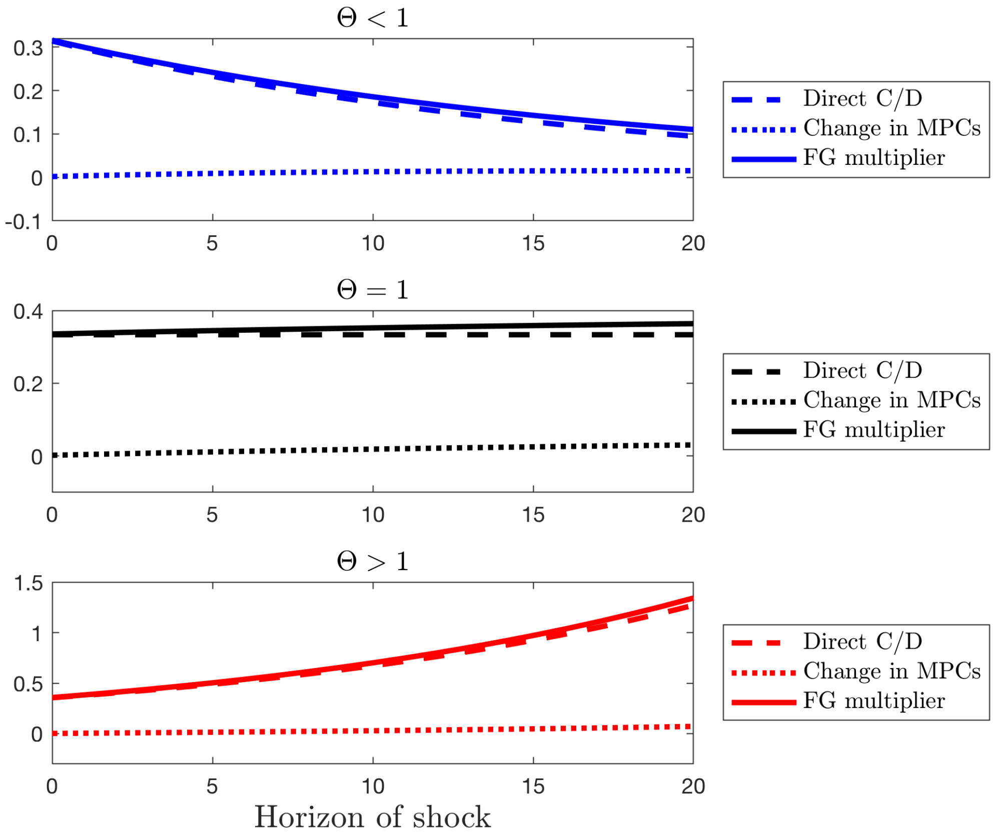

Figure 5: FG Multipliers: disentangling different channels.

Notes: The figure plots the "direct compounding/discounting (C/D)" channel, the "change in MPCs" channel, and the FG multiplier Φ_H as functions of the horizon H of the interest rate shock (all for three alternative values of Θ).

channel coming from heteroskedasticity described in Section 2.5.4) [26] and, then, compounded back to current output.

In addition to the previous MPC channel, suppose that $\Theta > 1$ (because, e.g., the correlation between the marginal change in risk and income sensitivities is negative). In this case, even keeping the path of MPCs fixed (i.e., isolating the effects of the "direct C/D channel"), positive news about an interest rate cut H periods ahead cumulate over time for two reasons. First, agents expect that the cut in $t+H$ will have an amplified effect on output in $t+H$ (through the usual demand amplification channel). Second, they also expect that output in $t+H-1$ will increase more than one-to-one with y_{t+H} since one-period ahead positive news are also amplified (through the usual demand amplification channel). Moving backward period-by-period, the sequential amplification effect compounds back to current output. Therefore, the further the horizon H, the more compounding takes place, making current output more responsive to the forward guidance than to a

[26] This last amplification effect is absent in Acharya and Dogra (2020), while in our AD equation (19a) Λ is premultiplied by Θ.

current interest rate cut.

Conversely, when considering the case of $\Theta < 1$ the "direct C/D channel" is active in the opposite direction (see the dashed line of the top panel in Figure 5), in the sense that dampening takes place and future news are discounted over time. Hence, if the path of MPCs were kept constant, positive values for our covariance statistic[27] (equivalently, for $\Theta < 1$) would be able to solve the FGP.

Turning to the "change in MPC" channel when $\Theta < 1$: the second term in Φ_H has a non-linear shape in H, increasing for smaller values of H and decreasing thereafter.[28] Intuitively, for longer horizons, the dampening effect on output of changes in MPCs (since Θ premultiplies Λ in our AD equation) prevails on the expansionary consequences of declines in MPCs, while this prediction is reversed when the cut occurs over shorter horizons.

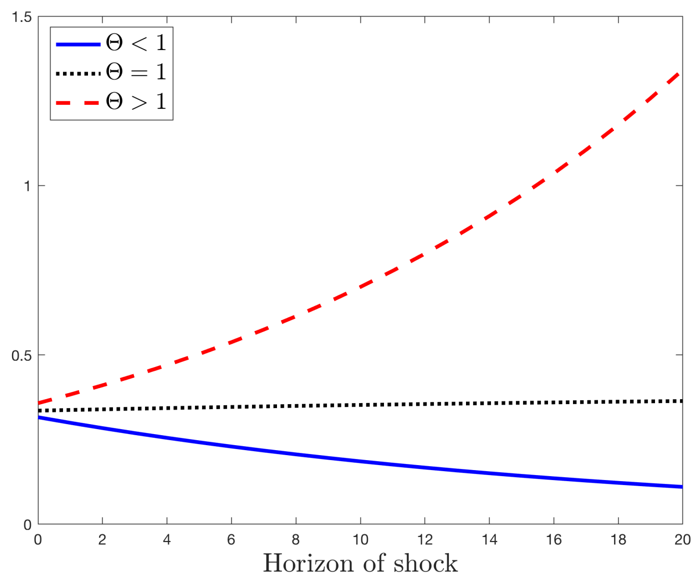

Figure 6: FG Multipliers

Notes: The figure plots the FG multiplier, Φ_H, as a function of the horizon H of the interest rate shock, for three alternative values of Θ.

Figure 6 summarizes the above-mentioned effects showing the FG multiplier for the cases of discounting (blue line) and compounding (red line) relative to the homoskedastic benchmark (black line). While $\Theta \geq 1$ is a sufficient condition to exacerbate the puzzle, $\Theta < 1$ may indeed solve the FGP. Notice, finally, that our model is not able to *simultaneously* attenuate the FGP and to generate

[27] Also, positive values for the cross-sectional average change in individual income risk.

[28] Studying the sign of the second term in the derivative of the FPG multiplier out of the time horizon, with $\Theta < 1$ (and $\Theta \neq \tilde{\beta}$) we can find the threshold $\hat{H}$ at which its sign is reversed:

$$\left(\frac{\Theta}{\tilde{\beta}}\right)^{\hat{H}+1} = \frac{\log(\tilde{\beta})}{\log(\Theta)}.$$

For $H < \hat{H}$, the second term is increasing in H, while it decreases for $H > \hat{H}$.

output amplification of demand shocks.[29] In fact, attenuation of the FGP requires $\Theta < 1$, while demand amplification needs $\Theta > 1$. This incompatibility occurs since the underlying mechanism that regulates output responses to both future news and demand shocks is similar and depends on the interaction between marginal changes in income risk and changes in the earnings distribution over the business cycle.

3 Incomplete information

So far we have assumed, in line with empirical evidence, that the volatility of the individual income process is a function of lagged income. In this section we show that a possible microfoundation for that functional relationship lies in *incomplete information*.

We assume that agents cannot observe the variance of the true income process and, in a Bayesian sense, must make inference on this moment using the realizations of their income as private signals. As a result, changes in realized earnings induce households to revise the estimates of their income volatility, modifying their *perception* of individual risk. In turn, this will lead agents to revise their optimal consumption and saving decisions.

The key difference relative to the previous setting featuring complete information lies in the nature of the heterogeneity in idiosyncratic risk. While in the complete information framework heteroskedasticity is an assumed structural feature of the income process, in the incomplete information setting the fundamental income process is assumed to be homoskedastic (in the cross-section). However, it is the *combination* of incomplete information *and* heterogeneity in signals (for income realizations differ across agents) that determines heterogeneity in the estimated volatility.

We first elucidate the role of heterogeneity in income signals in delivering heterogeneity in individual risk within a Bayesian learning procedure applied to a general *state-space* formulation. Then, we turn to analyse the general equilibrium effects of heterogeneity in the estimates of income volatility within an analytically tractable two-period version of the CARA-Gaussian model of the previous section.

3.1 General state-space formulation

Below we lay out a general state-space formulation for the evolution of the volatility of the income process.

3.1.1 Model and information structure

The economy is populated by a continuum of households in the interval $[0,1]$, indexed by $i \in [0,1]$. We assume a *state-space* model where the volatility σ_t^2 of individual income (homoskedastic in the cross-section) is the state variable, behaving according to a Markov process. Individual income is

[29] As common in the FGP literature with rational expectations, see Bilbiie (2018).

the *observation* variable, assumed to be i.i.d across agents and independent over time, conditional on the current variance state. Formally:

Assumption 4. *For every period t the fundamental individual income process for agent i is described by the following state-space formulation:*[30]

$$\sigma_t^2 \sim G_t(d\sigma_t^2 \mid \sigma_{t-1}^2), \tag{28a}$$

$$y_{i,t} \sim F_t(dy_{i,t} \mid \sigma_t^2), \tag{28b}$$

where G_t is the probability transition kernel[31] *for the variance (i.e., the state variable) from σ_{t-1}^2 to $d\sigma_t^2$, and F_t is the probability measure for individual income, conditional on the realization of time t variance σ_t^2.*

Notice that both G_t and F_t are common across agents, which in turn implies that the individual income process is identical across agents. In addition, the above state-space formulation requires the observation variable to be conditionally independent of past observations.

We postulate that each agent knows the state-space model in (28) and can observe the history of the earnings realizations and aggregate variables. However, the agent does not observe neither the realizations of the variance σ_t^2 (which needs to be estimated) nor other agents' earnings. In other words, we assume that the individual income process is a *private signal* that each household exploits to forecast her future income volatility, a key moment in the determination of her optimal consumption and saving decision. While income volatility evolves over time according to a Markov process, agents cannot observe it and use the history of current and past individual income realizations to obtain forecasts for the future variance of the income process.

3.1.2 Learning

We assume that rational agents *learn* over time about the variance of the income process in a Bayesian way, using the history of their earnings to derive state-predictive posterior distributions of future income volatility. Let $y_{i,0:t} := \left(y_{i,k}\right)_{k=0}^{t}$ denote the *history* of individual income realizations at time t (i.e., the vector of past and current realizations from 0 to t). Then, the output of agents' learning in every period t can be described by two conditional distributions.

- The *filtering distribution*: $\mathbb{P}_t(d\sigma_t^2 \mid y_{i,0:t})$, i.e., the distribution of the current (time t) income volatility conditional on the earnings history at time t.

- The *state-predictive distribution*: $\mathbb{P}_t(d\sigma_{t+1}^2 \mid y_{i,0:t})$, i.e., the distribution of future (time $t+1$) income volatility conditional on the earnings history at time t.

[30] We work here with probability measures, hence the notation $d\sigma_t^2, dy_{i,t}$.

[31] More formally, by definition of transition kernel, $G_t : (\mathbb{R}_+, \mathcal{B}(\mathbb{R}_+)) \to (\mathbb{R}_+, \mathcal{B}(\mathbb{R}_+))$ is a function such that $G_t(d\sigma_t^2 \mid \sigma_{t-1}^2)$ is a probability measure on $(\mathbb{R}_+, \mathcal{B}(\mathbb{R}_+))$ (where $\mathcal{B}(\mathbb{R}_+)$ is the Borel sigma-algebra on the set of positive real numbers) for every σ_{t-1}^2 and the map $\sigma_{t-1}^2 \to G_t(A \mid \sigma_{t-1}^2)$ is measurable for every set $A \in \mathcal{B}(\mathbb{R}_+)$.

To obtain the above mentioned distributions, the Bayesian learning procedure can be described by the following algorithm. Let $i \in [0,1]$ be a generic agent in the economy.

Initial stage, $t = 0$. Each agent i has a common prior $\mathbb{P}_0(d\sigma_0^2)$ over σ_0^2. After observing her first-period income $y_{i,0}$, agent i computes the filtering distribution through Bayesian updating:

$$\mathbb{P}_0(d\sigma_0^2 \mid y_{i,0}) \propto F_0(dy_{i,0} \mid \sigma_0^2) \cdot \mathbb{P}_0(d\sigma_0^2), \tag{29}$$

obtaining the state-predictive distribution for σ_1^2 as:

$$\mathbb{P}_0(d\sigma_1^2 | \, y_{i,0}) = \int_{\sigma_0^2 \in \mathbb{R}_+} G_1(d\sigma_1^2 \mid \sigma_0^2) \cdot \mathbb{P}_0(d\sigma_0^2 \mid y_{i,0}). \tag{30}$$

Recursive stages, $t \geq 1$. For the following periods, the learning procedure can be described recursively. At time $t \geq 1$, the prior is given by the state-predictive distribution computed by the agent in $t-1$, i.e., $\mathbb{P}_{t-1}(d\sigma_t^2 \mid y_{i,0:t-1})$.

When income is drawn at time t from the conditional distribution $F_t(dy_{i,t} \mid \sigma_t^2)$, agent i updates her prior, obtains the filtering distribution and, in turn, the new state-predictive distribution. In particular, the filtering distribution is given (up to a normalization constant) by:[32]

$$\mathbb{P}_t(d\sigma_t^2 \mid y_{i,0:t}) \propto F_t(dy_{i,t} \mid \sigma_t^2) \cdot \mathbb{P}_{t-1}(d\sigma_t^2 \mid y_{i,0:t-1}), \tag{31}$$

while the new state-predictive distribution is given by:

$$\mathbb{P}_t(d\sigma_{t+1}^2 | \, y_{i,0:t}) = \int_{\sigma_t^2 \in \mathbb{R}_+} G_{t+1}(d\sigma_{t+1}^2 \mid \sigma_t^2) \cdot \mathbb{P}_t(d\sigma_t^2 \mid y_{i,0:t}). \tag{32}$$

The above equation illustrates our key point. Namely, that the state-predictive distribution depends on the individual earnings history. This implies that agents that have experienced heterogeneous income realizations over time also derive heterogeneous distributions of future income volatility. This establishes a relationship between agent i's *estimated distribution* of the income variance and agent i's position in the earnings distribution. The existence of this link has, in turn, implications for each agent's optimal consumption and savings decisions: households with different earnings histories have different perceptions of their individual income risk and, therefore, make heterogeneous consumption choices.

[32] This expression is derived from Bayes' Theorem and from the assumption of conditional independence of income, i.e., $\mathbb{P}_t(dy_{i,t} \mid \sigma_t^2, y_{i,0:t-1}) = F_t(dy_{i,t} \mid \sigma_t^2)$:

$$\mathbb{P}_t(d\sigma_t^2 \mid y_{i,0:t}) = \frac{\mathbb{P}_t(dy_{i,t} \mid \sigma_t^2, y_{i,0:t-1})}{\mathbb{P}_t(dy_{i,t} \mid y_{i,0:t-1})} \cdot \mathbb{P}_{t-1}(d\sigma_t^2 \mid y_{i,0:t-1}) \propto$$
$$\propto F_t(dy_{i,t} \mid \sigma_t^2) \cdot \mathbb{P}_{t-1}(d\sigma_t^2 \mid y_{i,0:t-1}).$$

3.1.3 Consumer problem under imperfect information

As in Section 2.1 the agent can purchase, in every period, a risk-free real bond $a_{i,t+1}$ (paying one unit of good in $t+1$) at the price $1/(1+r_t)$ and consume a perishable consumption good $c_{i,t}$, using its total (net) income $y_{i,t}$ and its beginning-of-period wealth $a_{i,t}$. The recursive formulation of the consumer's problem requires to keep track of the filtering and the state-predictive distributions:

$$V_t(a_{i,t}, y_{i,t}) := \max_{(c_{i,t}, a_{i,t+1})} \left\{ u(c_{i,t}) + \beta \cdot \mathbb{E}_{i,t}\left[V_{t+1}(a_{i,t+1}, y_{i,t+1})\right] \right\}$$
$$\text{s.t.} \quad c_{i,t} + \frac{a_{i,t+1}}{1+r_t} = y_{i,t} + a_{i,t}, \quad a_{i,0} \text{ given.}$$

Notice that the expected value on the right-hand side of the Bellman equation is computed with respect to the distribution of future earnings conditional on the individual earnings history:

$$\mathbb{E}_{i,t}\left[V_{t+1}(a_{i,t+1}, y_{i,t+1})\right] = \int_{y_{i,t+1}\in\mathbb{R}} V_{t+1}(a_{i,t+1}, y_{i,t+1}) \cdot \mathbb{P}_t(dy_{i,t+1} \mid y_{i,0:t}),$$

where:

$$\mathbb{P}_t(dy_{i,t+1} \mid y_{i,0:t}) = \int_{\sigma^2_{t+1}\in\mathbb{R}_+} F_{t+1}(dy_{i,t+1} \mid \sigma^2_{t+1}) \cdot \mathbb{P}_t(d\sigma^2_{t+1} \mid y_{i,0:t}),$$

and $\mathbb{P}_t(d\sigma^2_{t+1} \mid y_{i,0:t})$ is given by equation (32). Importantly, the consumer problem is recursive: agent i does not need to know her earnings history in every period t, but can simply carry over her state-predictive distribution as a state variable. In particular, taking $\hat{\mathbb{P}}_t(d\sigma^2_t) := \mathbb{P}_{t-1}(d\sigma^2_t \mid y_{i,0:t-1})$ as a state variable in t, equations (31) and (32) describe the law of motion to obtain $\hat{\mathbb{P}}_{t+1}(d\sigma^2_{t+1}) = \mathbb{P}_t(d\sigma^2_{t+1} \mid y_{i,0:t})$ through the new income observation $y_{i,t}$.

3.2 An analytically tractable setup

The general consumption problem outlined in the previous section does not allow a closed-form solution. To illustrate the implications of heterogeneity in private income signals on agents' optimal consumption decisions we lay out a two-period version of the model. The economy runs for two periods: $\tau = t, t+1$. Compared to the general state-space model described above, we introduce a set of simplifying assumptions. First, the true variance of the income process is constant over time: $\sigma^2_{t+1} = \sigma^2_t = \sigma^2$. Second, the fundamental income process (that is i.i.d. across time and agents) follows a Gaussian distribution (conditional on the variance σ^2):

$$y_{i,t} \mid \sigma^2 \sim \mathcal{N}(y_t, \sigma^2),$$

where y_t is the period t mean. As in the state-space formulation described above, agents know both the process followed by their individual income and its mean y_t, yet not its second moment σ^2. Therefore, they use their private income signal to estimate the variance of income.

Each agent starts with a (common) prior distribution[33] $\sigma^2 \sim p(\sigma^2)$ at time t and, after observing her current income realization $y_{i,t}$, she updates her prior through Bayesian learning to obtain the posterior distribution: $\sigma^2 \mid y_{i,t} \sim p(\sigma^2 \mid y_{i,t})$. More compactly:

$$y_{i,t} \mid \sigma^2 \sim \mathcal{N}(y_t, \sigma^2) \quad i.i.d., \tag{33a}$$

$$\sigma^2 \sim p(\sigma^2), \tag{33b}$$

where (33a) indicates that, conditional on a given volatility, income follows a Gaussian distribution, and (33b) represents the prior held by the agent on the income variance.

The posterior distribution is proportional to the prior times the likelihood function:

$$p(\sigma^2 \mid y_{i,t}) \propto f_t(y_{i,t} \mid \sigma^2) \cdot p(\sigma^2), \tag{34}$$

where f_t is the Gaussian density of $y_{i,t}$ (as well as the likelihood function for any given $y_{i,t}$). Notice that equation (34) is a simplified version of equation (31).

Each agent i estimates the income volatility as the expected value of the posterior distribution. More formally, we define the estimator of the variance as:

$$\hat{\sigma}_i^2 := \mathbb{E}_{i,t}(\sigma^2 \mid y_{i,t}) = \int_{\mathbb{R}_+} \sigma^2 \cdot p(\sigma^2 \mid y_{i,t}) \cdot d\sigma^2, \tag{35}$$

where $\mathbb{E}_{i,t}(\cdot)$ denotes the expectation operator conditional on agent i's information set at time t.

Notice that the estimates of income volatility are heterogeneous across agents with different earnings draws. In other words, heteroskedasticty is now an outcome of the learning procedure.

Finally, we assume that each agent i uses the forecast $\hat{\sigma}_i^2$ *as if* it were the *correct* variance of the income process. This allows us to preserve the normality of the income process and to find closed-form solutions to the consumption problem.[34] It is therefore possible to define the relevant income process for consumption and saving decisions: $y_{i,t+1} \mid \mathcal{I}_{i,t} \sim \mathcal{N}(y_{t+1}, \hat{\sigma}_i^2)$ where $\mathcal{I}_{i,t}$ is agent i's information set at time t. In other words, after concluding their learning procedure, agents take their estimated future volatility as the true one and, conditional on their current income realization $y_{i,t}$), they know that their future income follows a Gaussian density with variance given by the estimated one.

3.2.1 Two-period consumer problem

The two-period version of the CARA-Gaussian model under incomplete information can be obtained as a special case of the infinite horizon model of Section 2.1 by imposing two conditions. First,

[33] We assume that the probability distributions admit densities. We use lower-case letters to denote densities, while capital letters to denote probability measures.

[34] Otherwise we would have to work with a non-Gaussian distribution for unconditional income, and we would not be able to obtain closed-form solutions.

we set the MPC out of wealth in the second period $(t+1)$ equal to 1:[35] $\mu_{t+1} = 1$. Second, we substitute $\sigma^2_{i,t+1}$ with the estimate of the income volatility $\hat{\sigma}^2_i$ obtained through Bayesian learning. See Appendix B.1 for a full derivation of the two-period model.

Imposing these conditions on equation (4) (after taking conditional expectations) and using the household's budget constraints, one can derive optimal current consumption as a function of the real interest rate, permanent income, and the estimate of income volatility:

$$c_{i,t} = \underbrace{-\frac{1}{(2+r_t)\gamma}\log[\beta(1+r_t)]}_{\text{intertemporal substitution}} + \underbrace{\frac{1+r_t}{2+r_t}\left(a_{i,t} + y_{i,t} + \frac{1}{1+r_t}y_{t+1}\right)}_{\text{permanent-income}} \underbrace{-\frac{1}{2+r_t}\frac{\gamma}{2}\hat{\sigma}^2_i}_{\text{estimated risk}}. \tag{36}$$

The first term on the right-hand side of equation (36) captures a typical intertemporal substitution motive stemming from real interest rate movements, while the second term captures the implications of the permanent-income hypothesis (i.e., consumption in every period is a fraction of lifetime wealth). Finally, the last term represents a shifter that depresses current consumption and captures the effects of *estimated* risk on precautionary savings. In particular, if the agent estimates higher values for income volatility (higher $\hat{\sigma}^2_i$), then she optimally increases her precautionary savings and reduces current consumption.

In addition, imposing $\mu_{t+1} = 1$ and $\sigma^2_{i,t+1} = \hat{\sigma}^2_i$ in equations (5) and (6), we can find the two-period version of the MPC out of wealth:

$$MPC^{\text{wealth}}_{i,t} = \mu_t = \frac{\partial c_{i,t}}{\partial a_{i,t}} = \frac{1+r_t}{2+r_t}, \tag{37}$$

and of the MPC out of transitory income shocks:[36]

$$MPC_{i,t} = \frac{\partial c_{i,t}}{\partial y_{i,t}} = \frac{1+r_t}{2+r_t} - \frac{1}{2+r_t}\frac{\gamma}{2}\frac{\partial \hat{\sigma}^2_i}{\partial y_{i,t}}. \tag{38}$$

It is clear from the previous expressions that the results from Proposition 1 still apply in the two-period model. In particular, if the marginal change in the estimated income volatility out of current earnings (i.e., $\partial\hat{\sigma}^2_i/\partial y_{i,t}$) is non-null and non-constant across the earnings distribution, then MPCs out of *income* differ across agents, are time-varying, and differ from the MPC out *wealth*. However, the mechanism that produces heterogeneity in MPCs here is different from the one in the complete information setup. Under incomplete information, changes in income are equivalent to changes in the *signal* received by each agent, who will therefore revise her own estimate of income volatility. Suppose, for instance, that in response to a marginal increase in her current earnings agent i revises downward her estimate of the income variance relatively more compared to agent j. Then agent i will perceive her future income as being relatively less *risky*, and will reduce her precautionary savings (increase her consumption) relatively more at the margin: therefore agent i will feature a

[35] Since the economy ends in the second period, agents optimally spend their resources entirely in $t+1$.
[36] Assume that $\hat{\sigma}^2_i : \mathbb{R} \to \mathbb{R}_+$ is differentiable with respect to income realizations $y_{i,t}$.

relatively higher MPC out of transitory income.

3.2.2 Aggregation and general equilibrium

The AD function of the two-period model can be obtained by imposing $\mu_{t+1} = 1$ and $\sigma^2_{i,t+1} = \hat{\sigma}^2_i$ to equation (9):

$$y_t = y_{t+1} - \frac{1}{\gamma}\log[\beta(1+r_t)] - \frac{\gamma}{2}\int_{[0,1]}\hat{\sigma}^2_i di + g_t - g_{t+1}. \tag{39}$$

Following the same proof as in Proposition 3, we can derive the analog of the three sufficient statistics that summarize the response of aggregate output to a demand shock.

(i) Expected marginal change in the estimate of income volatility out of earnings:

$$\underbrace{\mathbb{E}\left(\frac{\partial \hat{\sigma}^2_i}{\partial y_{i,t}}\right)}_{\text{individual risk effect}}.$$

(ii) Expected earnings sensitivity to aggregate output:

$$\underbrace{\mathbb{E}\left(\frac{\partial y_{i,t}}{\partial y_t}\right)}_{\text{income cyclicality effect}}.$$

(iii) Covariance between the marginal change in the estimate of income volatility and the earnings sensitivity to aggregate output:

$$\underbrace{Cov\left(\frac{\partial \hat{\sigma}^2_i}{\partial y_{i,t}}, \frac{\partial y_{i,t}}{\partial y_t}\right)}_{\text{risk-cyclicality covariance effect}}.$$

In particular, we can define Θ_t as the output multiplier out of a fully transitory government spending shock (assuming that the real interest rate is unchanged) as the derivative of current output to government spending:

$$\Theta_t := \frac{\partial y_t}{\partial g_t} = \left\{1 + \frac{\gamma}{2}\left[\mathbb{E}\left(\frac{\partial \hat{\sigma}^2_i}{\partial y_{i,t}}\right)\cdot\mathbb{E}\left(\frac{\partial y_{i,t}}{\partial y_t}\right) + Cov\left(\frac{\partial \hat{\sigma}^2_i}{\partial y_{i,t}}, \frac{\partial y_{i,t}}{\partial y_t}\right)\right]\right\}^{-1}. \tag{40}$$

Again, this multiplier is the same as in the infinite-horizon model, with $\mu_{t+1} = 1$. A similar multiplier could be obtained for a transitory monetary shock (e.g., an increase in the real interest rate):

$$\frac{\partial y_t}{\partial r_t} = -\Theta_t \cdot \frac{1}{\gamma(1+r_t)}. \tag{41}$$

As in the complete information framework, negative values for the covariance statistic and for the product between the average change in the estimate of variance and the average income sensitivity determine amplification in the response of output to demand shocks.

Consider a negative demand shock that has a first-round *direct* negative effect on aggregate output. Then, the reduction in output affects individual earnings in a heterogeneous fashion, depending on income sensitivities $(\partial y_{i,t}/\partial y_t)_{i\in[0,1]}$. As agents observe modifications in their signals, they revise their estimates of income volatility and change their consumption at the margin. For instance, if those agents featuring larger income sensitivities (i.e., those who lose relatively more from the contraction in output) are also those who revise their estimates upward marginally more, the decline in output will be amplified. In general, the interaction between revisions in estimates and income sensitivities generates an *indirect* effect of demand shocks on aggregate output.

3.3 A tractable case: Inverse-Gamma prior

A key advantage of the incomplete information setting as a foundation of heteroskedasticity is that it allows a flexible parameterization of the prior distribution of income volatility. In turn this allows us to easily obtain closed-form solutions for the posterior estimator $\hat{\sigma}_i^2$.

We assume that each agent $i \in [0,1]$ has a common Inverse-Gamma prior, with common hyper-parameters $\sigma_0^2 > 0$, $\nu_0 > 0$:

$$\sigma^2 \sim \text{Inv-Gamma}\Big(\frac{\nu_0}{2}, \frac{\nu_0 \sigma_0^2}{2}\Big). \tag{42}$$

After observing their income realization $y_{i,t}$, agents update the prior distribution following Bayes' Theorem and obtain the posterior distribution, which is itself an Inverse-Gamma with newly derived parameters:[37]

$$\sigma^2 \mid y_{i,t} \sim \text{Inv-Gamma}\Big(\frac{\nu_0 + 1}{2}, \frac{\nu_0 \sigma_0^2 + (y_{i,t} - y_t)^2}{2}\Big). \tag{43}$$

Then, each household estimates her own income volatility using the conditional expected value:

$$\hat{\sigma}_i^2 = \frac{\nu_0 \sigma_0^2 + (y_{i,t} - y_t)^2}{\nu_0 - 1}. \tag{44}$$

The intuition for (44) is standard in Bayesian statistics. The posterior estimator can be expressed as a weighted average of the prior mean[38] and the squared deviation of the income realization from

[37] See Appendix B.2.

[38] We define the prior mean as the expected value of the variance with respect to the prior distribution:

$$\text{prior mean} := \mathbb{E}(\sigma^2).$$

See Appendix B.2 for details on the moments of the Inverse-Gamma distribution.

its average, where the weights are given by the prior sample size and the actual sample size:[39]

$$\hat{\sigma}_i^2 = \left[\frac{\nu_0 - 2}{\nu_0 - 1} \cdot \underbrace{\frac{\nu_0 \sigma_0^2}{\nu_0 - 2}}_{\text{prior mean}} + \frac{1}{\nu_0 - 1} \cdot \underbrace{(y_{i,t} - y_t)^2}_{\text{squared deviation}}\right]. \tag{45}$$

In our model the actual sample size is equal to 1, since households observe a sample with only one (i.e., their own) income realization $y_{i,t}$. Instead, the prior sample size is $n_0 = \nu_0 - 2$. Therefore we can use the hyperparameter ν_0 to determine the prior sample size, while the prior mean is a function of both σ_0^2 and ν_0.

Let $\mathcal{Y}_{i,t} := y_{i,t} - y_t$ denote the deviation of realized individual income from its mean. Consider the expression of the posterior estimator in (44). What actually matters for the estimate of the income variance is not income itself, but its deviation from its known mean (that coincides with aggregate output). Hence, we can rewrite the expression for the government spending multiplier Θ_t as:

$$\Theta_t := \frac{\partial y_t}{\partial g_t} = \left\{1 + \frac{\gamma}{2}\left[\mathbb{E}\left(\frac{\partial \hat{\sigma}_i^2}{\partial \mathcal{Y}_{i,t}}\right) \cdot \mathbb{E}\left(\frac{\partial \mathcal{Y}_{i,t}}{\partial y_t}\right) + Cov\left(\frac{\partial \hat{\sigma}_i^2}{\partial \mathcal{Y}_{i,t}}, \frac{\partial \mathcal{Y}_{i,t}}{\partial y_t}\right)\right]\right\}^{-1}. \tag{46}$$

Then, given the estimator $\hat{\sigma}_i^2$ in (44), we can derive explicit formulas for the sufficient statistics of Proposition 3.

$$\mathbb{E}\left(\frac{\partial \hat{\sigma}_i^2}{\partial \mathcal{Y}_{i,t}}\right) = 0 \tag{47a}$$

$$\mathbb{E}\left(\frac{\partial \mathcal{Y}_{i,t}}{\partial y_t}\right) = 0 \tag{47b}$$

$$Cov\left(\frac{\partial \hat{\sigma}_i^2}{\partial \mathcal{Y}_{i,t}}, \frac{\partial \mathcal{Y}_{i,t}}{\partial y_t}\right) = \frac{2}{\nu_0 - 1} \int_{[0,1]} \mathcal{Y}_{i,t} \cdot \frac{\partial \mathcal{Y}_{i,t}}{\partial y_t} \, di. \tag{47c}$$

The fact that the average marginal change in the estimate of income volatility is zero stems from the symmetry of the Gaussian distribution. As apparent from the estimator in (44), given that symmetry, agents on the opposite sides of the income distribution, but equidistant from its mean, form the same estimate of income volatility. Therefore, the estimates depend, at the margin, on the deviations of income from its mean, $\mathcal{Y}_{i,t}$, that, on average, are null. More formally:

$$\frac{\partial \hat{\sigma}_i^2}{\partial \mathcal{Y}_{i,t}} = \frac{2}{\nu_0 - 1} \cdot \mathcal{Y}_{i,t} \implies \frac{2}{\nu_0 - 1} \int_{[0,1]} \mathcal{Y}_{i,t} \, di = 0.$$

Moreover, the second sufficient statistic (47b) is also zero:

$$\frac{\partial \mathcal{Y}_{i,t}}{\partial y_t} = \frac{\partial y_{i,t}}{\partial y_t} - 1 \implies \int_{[0,1]} \left(\frac{\partial y_{i,t}}{\partial y_t} - 1\right) di = 0.$$

[39] We assume that $\nu_0 > 2$ for the existence of the first moment of the prior distribution.

The covariance in (47c) is however different from zero:

$$Cov\left(\frac{\partial \hat{\sigma}_i^2}{\partial \mathcal{Y}_{i,t}}, \frac{\partial \mathcal{Y}_{i,t}}{\partial y_t}\right) = \frac{2}{\nu_0 - 1} \int_{[0,1]} \mathcal{Y}_{i,t} \cdot \left(\frac{\partial y_{i,t}}{\partial y_t} - 1\right) \, di = \frac{2}{\nu_0 - 1} \int_{[0,1]} (y_{i,t} - y_t) \cdot \frac{\partial y_{i,t}}{\partial y_t} \, di.$$

Its sign depends on the integral term and cannot be established without knowing the values of the individual income sensitivities $(\partial y_{i,t}/\partial y_t)_{i \in [0,1]}$.

3.3.1 Calibration in the simple case

We assume that the true variance of the income process is $\sigma^2 = 0.5$ (as in Section 2.5.5) and that average output is equal to 2.[40] Then, we assign a value of $r_t = 4\%$ to the real interest rate and of $\gamma = 3$ to the coefficient of absolute risk aversion (as in Section 2.5.5).

We can calibrate the hyperparameters by setting, first of all, the prior sample size parameter $\nu_0 > 2$ (recalling that the prior sample size is $n_0 = \nu_0 - 2$) then by choosing the parameter σ_0^2 to target the true value of the fundamental income variance. In particular, we want the prior mean to be equal to the fundamental volatility:[41]

$$\frac{\nu_0 \sigma_0^2}{\nu_0 - 2} = \sigma^2 \implies \sigma_0^2 = \sigma^2 \frac{\nu_0 - 2}{\nu_0}.$$

In our baseline calibration we set the prior sample size parameter equal to 10, that implies a value for σ_0^2 of 0.40. We simulate a panel of 1,000 agents, assigning to each of them a realization of the income process $y_{i,t} \sim \mathcal{N}(2, 0.5)$ and computing the associated estimates of income volatility according to the formula in (44).

In Figure 7a we plot the estimates of the variance of income $\hat{\sigma}_i^2$ as a function of income realizations $y_{i,t}$. We can observe that the estimates are U-shaped, as a consequence of the symmetry of the Gaussian density. Moreover, we can easily observe that the further the income realization is from its mean, the greater the estimate of income volatility, since agents interpret a large deviation from the mean of the income distribution as a signal for larger variance.

Moreover, taking the derivative of the volatility estimator (44), we can find a closed-form expression for the MPCs out of transitory income:

$$MPC_{i,t} = \frac{\partial c_{i,t}}{\partial y_{i,t}} = \frac{1 + r_t}{2 + r_t} - \frac{1}{2 + r_t} \frac{\gamma}{\nu_0 - 1} \cdot (y_{i,t} - y_t). \tag{48}$$

From equation (48) it is clear that MPCs are decreasing in realized income and, when $y_{i,t} < y_t$ ($y_{i,t} > y_t$), they are above (below) the MPC out of wealth.[42] Figure 7b plots the MPCs out of

[40] We choose this value in such a way to reduce significantly the probability that income realizations take negative values in our simulations.

[41] In this way, there is an anchoring of agents' pre-learning expectations to the true value of the fundamental variance.

[42] Recall that the MPC out of wealth is equal to $(1 + r_t)/(2 + r_t)$.

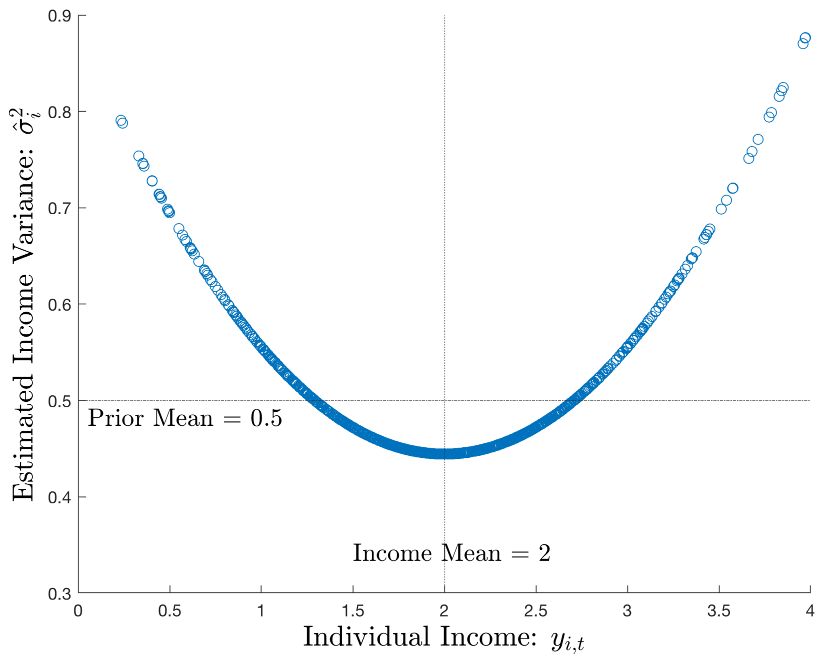

FIGURE 7a: HETEROGENEITY IN ESTIMATES OF INCOME VOLATILITY OVER THE INCOME DISTRIBUTION.

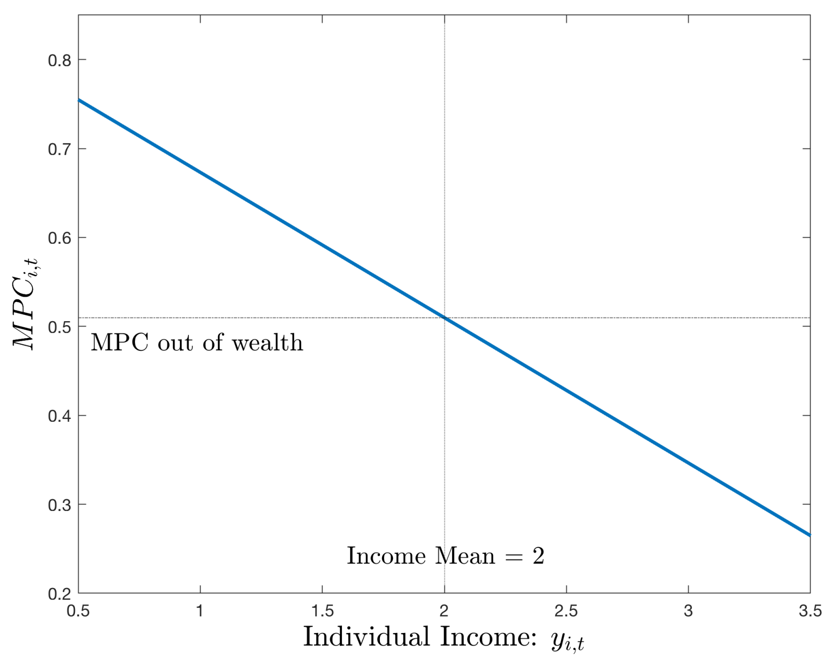

FIGURE 7b: HETEROGENEITY IN MPCS OUT OF TRANSITORY INCOME OVER THE INCOME DISTRIBUTION.

transitory income as a function of earnings, displaying a negative linear relation between MPCs and income. Therefore, our "Inverse-Gamma - Gaussian Bayesian" model is able to reproduce the empirical evidence whereby MPCs decline over the earnings distribution.

This result is an implication of the convexity of the variance estimator (44) in current income (in deviation from its mean). Intuitively, given the normality assumption on the income process (with common and known mean), and given the absence of any persistence in earnings, agents interpret any deviation of their income from its mean as purely transitory and anticipate that mean reversion might take place in the future. Since rich individuals anticipate that their income will revert to the mean in the future, they save a major part of it and have lower MPCs. The opposite is true for poorer agents, who anticipate higher future income and currently spend more.

Suppose for instance, and consistently with the empirical evidence in Patterson (2021), that MPCs correlate positively with income sensitivities, i.e., agents whose earnings are more elastic to aggregate output have higher MPCs. In our model, the positive sign of this correlation is generated by a negative covariance between marginal changes in volatility estimates and earnings sensitivities:

$$Cov\left(\frac{\partial \hat{\sigma}_i^2}{\partial \mathcal{Y}_{i,t}}, \frac{\partial \mathcal{Y}_{i,t}}{\partial y_t}\right) = \frac{2}{\nu_0 - 1} \int_{[0,1]} (y_{i,t} - y_t) \cdot \frac{\partial y_{i,t}}{\partial y_t} \; di \propto -\underbrace{Cov\left(MPC_{i,t}, \frac{\partial y_{i,t}}{\partial y_t}\right)}_{>0} < 0.$$

Under this assumption, poor individuals, who have higher MPCs, have relatively more sensitive earnings to aggregate output. In general equilibrium, amplification of demand shocks will, therefore, take place.

3.3.2 Robustness

We study the implications of different values of the hyperparameters of the prior distribution. In particular, we continue to assume that agents set those parameters in such a way to have their prior mean coinciding with the true value of the income variance. In particular, after assigning the prior sample size $n_0 = \nu_0 - 2$, they choose σ_0^2 to have:

$$\sigma_0^2 = \sigma^2 \frac{\nu_0 - 2}{\nu_0}.$$

We argue that the prior sample size parameter ν_0 regulates the strength of the belief that agents have in their prior. In fact, from the expression of the variance estimator (44), the greater ν_0, the larger the weight given to the prior mean in the Bayesian updating procedure:

$$\hat{\sigma}_i^2 = \left[\frac{\nu_0 - 2}{\nu_0 - 1} \cdot \underbrace{\frac{\nu_0 \sigma_0^2}{\nu_0 - 2}}_{\text{prior mean}} + \frac{1}{\nu_0 - 1} \cdot \underbrace{(y_{i,t} - y_t)^2}_{\text{squared deviation}}\right].$$

Figure 8 plots the distribution of the estimates of income volatility (as a function of earnings) for different values of the prior sample size hyperparameter.

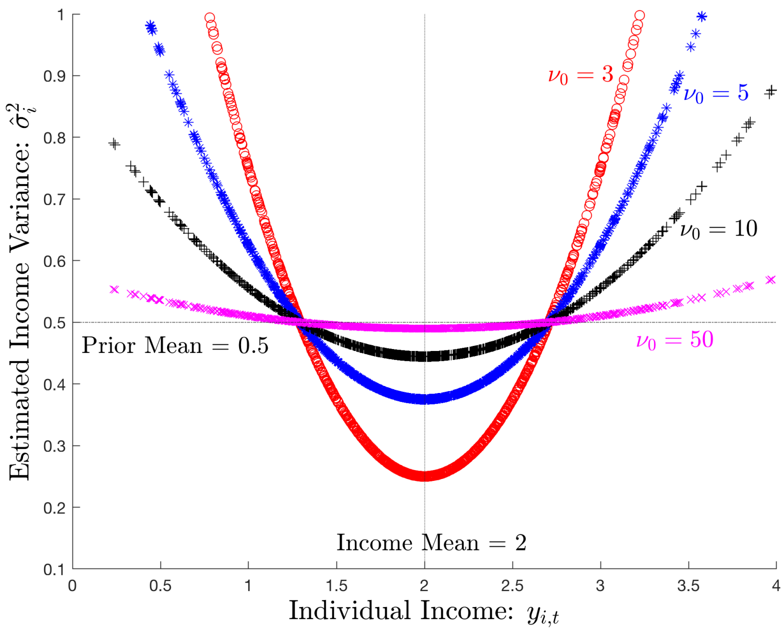

FIGURE 8: ESTIMATES OF FUTURE INCOME VOLATILITY FOR DIFFERENT PRIOR SAMPLE SIZES.

As the prior sample size increases, the estimates remain U-shaped (since estimates are quadratic in the deviations of the income realization from its mean), but become progressively flatter, converging to the prior mean (and coinciding with the fundamental variance $\sigma_t^2 = 0.5$). Intuitively, the agents give more weight to their prior, while the signal becomes progressively (as ν_0 increases) less relevant for the posterior estimator.[43] Hence, we can approximate the complete information model for large values of the prior sample size, i.e., $\nu_0 \to \infty$:

$$\hat{\sigma}_i^2 = \left[\frac{\nu_0 - 2}{\nu_0 - 1} \cdot \frac{\nu_0 \sigma_0^2}{\nu_0 - 2} + \frac{1}{\nu_0 - 1} \cdot (y_{i,t} - y_t)^2 \right] \to \sigma_0^2 = \sigma_t^2$$

where all agents know that income volatility coincides with the prior mean (i.e., with the true value of the fundamental variance).

[43] On the contrary, if we considered $\nu_0 = 2$, we would be using an improper prior (coinciding with the Jeffreys prior), that is an uninformative one. In this case, not plotted here since the prior mean does not exist, the prior distribution does not give any information to the agent and the variance (posterior) estimator would simply coincide with the squared income deviation from its mean:

$$\hat{\sigma}_i^2 = (y_{i,t} - y_t)^2.$$

4 Conclusion

We have studied the implications of cross-sectional heteroskedasticity in individual income in a heterogeneous agent New Keynesian model. Our setup is grounded in the empirical evidence from administrative micro data showing significant heterogeneity in the volatility of individual income over the earnings distribution.

Under complete information, we have assumed the existence of a structural relation linking lagged income to earnings volatility. In this setting, changes in earnings affect idiosyncratic risk and, in turn, consumption choices. We have then shown that such a relationship can be microfounded in a setting with incomplete information and Bayesian learning, whereby agents are not able to observe the true variance of their income process, but estimate it using their income realizations as private signals.

Unlike a framework with homoskedastic income (where the MPC is uniform across agents), our model is able to generate heterogeneity in MPCs out of transitory income (consistent with some empirical literature), and to account for idiosyncratic risk and market incompleteness in an analytically tractable way. Since in our model marginal changes in individual earnings shape income volatility at the margin, saving and consumption choices are also affected at the margin, implying that individual MPCs depend on individual income.

We have shown that, through indirect general equilibrium effects, heteroskedasticity affects the response of output to demand shocks, the transmission of fiscal and monetary policy, the (in)determinacy of equilibrium, and the FG puzzle. Those general equilibrium effects hinge on the interaction between the income distribution and changes in individual risk. Most importantly, those effects can be captured by a set of sufficient statistics which allow, in principle, to further discipline the model via a small set of moments observable in micro data. In general, those sufficient statistics are null both in a RA model and in a corresponding CARA framework with homoskedasticity.

We suggest that future research on the macroeconomic implications of cross-sectional heteroskedasticity in individual income should take (at least) three directions. First, empirical estimates of the relevant (three) sufficient statistics should be obtained from micro data. Second, those estimated moments should be used to discipline macroeconomic models with heteroskedasticity in individual income and to understand quantitatively the economic relevance of the heteroskedasticity channel that we propose. Third, the model with incomplete information could be integrated with more elaborated behavioral hypotheses to study how the misperception of the second moment of the income process interacts with agents' consumption and saving decisions in general equilibrium.

References

Acharya, S., E. Challe, and K. Dogra (2020, February). Optimal Monetary Policy According to HANK. CEPR Discussion Paper 14429, C.E.P.R. Discussion Papers.

Acharya, S. and K. Dogra (2020). Understanding HANK: Insights From a PRANK. *Econometrica 88*(3), 1113–1158.

Arellano, M. (2014, October). Uncertainty, Persistence, and Heterogeneity: A Panel Data Perspective. *Journal of the European Economic Association 12*(5), 1127–1153.

Arellano, M., R. Blundell, and S. Bonhomme (2017). Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework. *Econometrica 85*(3), 693–734.

Arellano, M. and S. Bonhomme (2017). Nonlinear Panel Data Methods for Dynamic Heterogeneous Agent Models. *Annual Review of Economics 9*(1), 471–496.

Ascari, G. and L. Rossi (2012, September). Trend Inflation and Firms Price-Setting: Rotemberg Versus Calvo. *The Economic Journal 122*(563), 1115–1141.

Auclert, A. (2019, June). Monetary Policy and the Redistribution Channel. *American Economic Review 109*(6), 2333–2367.

Auclert, A., M. Rognlie, and L. Straub (2018, September). The Intertemporal Keynesian Cross. Working Paper 25020. *National Bureau of Economic Research*.

Bayer, C., R. Luetticke, L. Pham-Dao, and V. Tjaden (2019). Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk. *Econometrica 87*(1), 255–290.

Bilbiie, F. O. (2008, May). Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic. *Journal of Economic Theory 140*(1), 162–196.

Bilbiie, F. O. (2018, January). Monetary Policy and Heterogeneity: An Analytical Framework. CEPR Discussion Papers 12601, C.E.P.R. Discussion Papers.

Bilbiie, F. O. (2020, October). The New Keynesian Cross. *Journal of Monetary Economics 114*, 90–108.

Bilbiie, F. O., T. Monacelli, and R. Perotti (2013, February). Public Debt and Redistribution with Borrowing Constraints. *The Economic Journal 123*(566), F64–F98.

Caballero, R. J. (1990, January). Consumption Puzzles and Precautionary Savings. *Journal of Monetary Economics 25*(1), 113–136.

Challe, E. (2020, April). Uninsured Unemployment Risk and Optimal Monetary Policy in a Zero-Liquidity Economy. *American Economic Journal: Macroeconomics 12*(2), 241–283.

Challe, E. and X. Ragot (2016, February). Precautionary Saving Over the Business Cycle. *The Economic Journal 126*(590), 135–164.

Chodorow-Reich, G., P. T. Nenov, and A. Simsek (2021, May). Stock Market Wealth and the Real Economy: A Local Labor Market Approach. *American Economic Review 111*(5), 1613–1657.

Debortoli, D. and J. Galí (2022, January). Idiosyncratic Income Risk and Aggregate Fluctuations. Working Paper 29704. *National Bureau of Economic Research*.

Del Negro, M., M. P. Giannoni, and C. Patterson (2015). The Forward Guidance Puzzle. Staff Reports 574, Federal Reserve Bank of New York.

Dixit, A. K. and J. E. Stiglitz (1977). Monopolistic Competition and Optimum Product Diversity. *The American Economic Review 67*(3), 297–308.

Guren, A. M., A. McKay, E. Nakamura, and J. Steinsson (2021, March). Housing Wealth Effects: The Long View. *The Review of Economic Studies 88*(2), 669–707.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021). What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Dynamics? *Econometrica 89*(5), 2303–2339.

Guvenen, F., S. Ozkan, and J. Song (2014, June). The Nature of Countercyclical Income Risk. *Journal of Political Economy 122*(3), 621–660.

Guvenen, F., S. Schulhofer-Wohl, J. Song, and M. Yogo (2017, May). Worker Betas: Five Facts about Systematic Earnings Risk. *American Economic Review 107*(5), 398–403.

Jappelli, T. and L. Pistaferri (2014, October). Fiscal Policy and MPC Heterogeneity. *American Economic Journal: Macroeconomics 6*(4), 107–136.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006, December). Household Expenditure and the Income Tax Rebates of 2001. *American Economic Review 96*(5), 1589–1610.

Kaplan, G., B. Moll, and G. L. Violante (2018, March). Monetary Policy According to HANK. *American Economic Review 108*(3), 697–743.

McKay, A., E. Nakamura, and J. Steinsson (2016, October). The Power of Forward Guidance Revisited. *American Economic Review 106*(10), 3133–3158.

McKay, A., E. Nakamura, and J. Steinsson (2017). The Discounted Euler Equation: A Note. *Economica 84*(336), 820–831.

Meghir, C. and L. Pistaferri (2004). Income Variance Dynamics and Heterogeneity. *Econometrica 72*(1), 1–32.

Patterson, C. (2021, February). The Matching Multiplier and the Amplification of Recessions. Working Paper. https://www.christinahydepatterson.com/_files/ugd/32299b_a868c382991f42d2876c593ad0537ee3.pdf.

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

Rotemberg, J. J. (1982, December). Sticky Prices in the United States. *Journal of Political Economy 90*(6), 1187–1211.

Sargent, T. J. and N. Wallace (1975). “Rational” Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule. *Journal of Political Economy 83*(2), 241–254.

Storesletten, K., C. I. Telmer, and A. Yaron (2004, June). Cyclical Dynamics in Idiosyncratic Labor Market Risk. *Journal of Political Economy 112*(3), 695–717.

Uhlig, H. (1988, August). A Law of Large Numbers for Large Economies. Working Paper 342, Federal Reserve Bank of Minneapolis, Research Department.

Wang, N. (2002, December). Consumption and Conditionally Heteroskedastic Income. SSRN Scholarly Paper ID 422323, Social Science Research Network, Rochester, NY.

Wang, N. (2003, June). Caballero Meets Bewley: The Permanent-Income Hypothesis in General Equilibrium. *American Economic Review 93*(3), 927–936.

Werning, I. (2015, August). Incomplete Markets and Aggregate Demand. Working paper 21448. *National Bureau of Economic Research*.

Wolf, C. K. (2021, August). Interest Rate Cuts vs. Stimulus Payments: An Equivalence Result. Working Paper 29193. *National Bureau of Economic Research*.

Woodford, M. (2011, January). Simple Analytics of the Government Expenditure Multiplier. *American Economic Journal: Macroeconomics 3*(1), 1–35.

APPENDIX

A A heteroskedastic New Keynesian model

This appendix contains the proofs of the results and the description of the supply-side of the heteroskedastic New Keynesian model of Section 2.

A.1 Proofs of Propositions 1, 2, 3

A.1.1 Proposition 1

Proof. Suppose that $c_{i,t}$ follows the process:

$$c_{i,t+1} = \phi_{i,t} c_{i,t} + \Gamma_{i,t} + v_{i,t+1}, \tag{A-1}$$

where $\phi_{i,t}$ and the drift $\Gamma_{i,t}$ are time- and individual-specific coefficients, and $v_{i,t+1}$ is the consumption process innovation in $t+1$.

- Firstly, substituting the guess (A-1) in the Euler equation (3) and taking the logs we obtain:

  $$-\gamma(1-\phi_{i,t})c_{i,t} = \log[\beta(1+r_t)] - \gamma\Gamma_{i,t} + \log \mathbb{E}_t[\exp\{-\gamma v_{i,t+1}\}].$$

  Then, $\phi_{i,t} = 1$. Otherwise, as explained by Caballero (1990), consumption would be entirely determined from the Euler equation, without any reference to the intertemporal budget constraint. Imposing $\phi_{i,t} = 1$, we can write the drift as a function of the real interest rate and the moments of the consumption process innovation:

  $$\Gamma_{i,t} = \frac{1}{\gamma}\log[\beta(1+r_t)] + \frac{1}{\gamma}\log \mathbb{E}_t[\exp\{-\gamma v_{i,t+1}\}]. \tag{A-2}$$

- Secondly, let us consider the intertemporal budget constraint (iBC):

  $$\sum_{k\geq 0} Q_{t,t+k} c_{i,t+k} = \sum_{k\geq 0} Q_{t,t+k} y_{i,t+k} + a_{i,t}, \tag{A-3}$$

  where $Q_{t,t+k} \coloneqq [(1+r_t)\cdot\ldots\cdot(1+r_{t+k-1})]^{-1}$ for $k \geq 1$ and $Q_{t,t} = 1$. Additionally, let us integrate backward the consumption process in (A-1) (with $\phi_{i,t} = 1$):

  $$c_{i,t+k} = c_{i,t} + \sum_{j=1}^{k}(\Gamma_{i,t+j-1} + v_{i,t+j}),$$

Substituting this expression in the iBC, after some manipulations, we obtain:

$$\begin{aligned}&\sum_{k\geq 0} Q_{t,t+k} c_{i,t} + \sum_{k\geq 1} Q_{t,t+k} \sum_{j=1}^{k} (\Gamma_{i,t+j-1} + v_{i,t+j}) = \\ &= a_{i,t} + \sum_{k\geq 0} Q_{t,t+k} \mathbb{E}_t(y_{i,t+k}) + \sum_{k\geq 1} Q_{t,t+k} \underbrace{[y_{i,t+k} - \mathbb{E}_t(y_{i,t+k})]}_{\epsilon_{i,t+k}}.\end{aligned} \tag{A-4}$$

- Taking the expectations of (A-4) conditional on time t information, we obtain the solution for consumption:

$$c_{i,t} = \mu_t a_{i,t} + \mu_t \left(y_{i,t} + \sum_{k\geq 1} Q_{t,t+k} y^d_{i,t+k} \right) - \tilde{\Gamma}_{i,t}, \tag{A-5}$$

  where μ_t is the MPC out of wealth, such that: $\mu_t := \left(\sum_{k\geq 0} Q_{t,t+k} \right)^{-1}$ and the constant $\tilde{\Gamma}_{i,t}$ is defined as:

$$\tilde{\Gamma}_{i,t} := \mu_t \sum_{k\geq 1} Q_{t,t+k} \sum_{j=1}^{k} \Gamma_{i,t+j-1}.$$

  Note that we can establish (as in Acharya and Dogra, 2020) the following recursion for the wealth MPC:

$$\mu_t^{-1} = 1 + Q_{t,t+1} \mu_{t+1}^{-1} \iff \mu_t = \frac{\mu_{t+1}(1 + r_t)}{1 + \mu_{t+1}(1 + r_t)}.$$

- To solve for the consumption process innovation as a function of the structural income innovation, we can combine equations (A-4) and (A-5). Then:[44]

$$\begin{aligned}&\sum_{k\geq 1} Q_{t,t+k} \left(\sum_{j=1}^{k} v_{i,t+j} - \epsilon_{i,t+k} \right) = 0 \iff {}^{45} \\ &\sum_{j\geq k} Q_{t,t+j} v_{i,t+k} - Q_{t,t+k} \epsilon_{i,t+k} = 0 \quad \forall k \geq 1 \implies\end{aligned}$$

$$v_{i,t+k} = \mu_{t+k} \epsilon_{i,t+k} \quad \forall k \geq 1.^{46} \tag{A-6}$$

[44] Note that the following result holds since we assumed that the entire path of variances is known (given the dependence on the deterministic components of lagged income). Otherwise, if the path of variances were stochastic, equation (A-5) would have to be modified to:

$$c_{i,t} = \mu_t a_{i,t} + \mu_t \left(y_{i,t} + \sum_{k\geq 1} Q_{t,t+k} y^d_{i,t+k} \right) - \mathbb{E}_t(\tilde{\Gamma}_{i,t}).$$

Then, combining the previous equation with equation (A-4), we would get, almost surely:

$$\sum_{k\geq 1} Q_{t,t+k} \left(\sum_{j=1}^{k} v_{i,t+j} - \epsilon_{i,t+k} \right) + \sum_{k\geq 1} Q_{t,t+k} \sum_{j=1}^{k} \left[\Gamma_{i,t+j-1} - \mathbb{E}_t(\Gamma_{i,t+j-1}) \right] = 0,$$

where $\Gamma_{i,t+j-1}$ depends on the moments of the income innovation process whose variance is stochastic.

[45] While the *if* part of this statement is trivial, the *only if* can be proved by contraposition, exploiting the time-independence of consumption and income innovations.

Substituting equation (A-6) back in equation (A-2) and using the moments of the log-Normal distribution,[47] we can finally obtain the formula for the drift of the consumption process.

$$\Gamma_{i,t} = \frac{1}{\gamma}\log[\beta(1+r_t)] + \frac{\gamma\mu_{t+1}^2}{2}\sigma_{i,t+1}^2. \tag{A-7}$$

To conclude, substituting the drift (A-7) and the consumption innovation (A-6) back in the initial guess, we obtain the stochastic difference equation that characterize the solution for the consumption process.

$$c_{i,t} = c_{i,t+1} \underbrace{-\frac{1}{\gamma}\log[\beta(1+r_t)] - \frac{\gamma\mu_{t+1}^2}{2}\sigma_{i,t+1}^2}_{=-\Gamma_{i,t}} \underbrace{-\mu_{t+1}\epsilon_{i,t+1}}_{=-v_{i,t+1}}.$$

■

A.1.2 Proposition 2

Proof. Let us define the MPC out of income as the partial derivative of the optimal consumption function to current income: $MPC_{i,t} := \frac{\partial c_{i,t}}{\partial y_{i,t}}$. To compute equation (6), we can differentiate equation (A-5):

$$\begin{aligned} MPC_{i,t} := \frac{\partial c_{i,t}}{\partial y_{i,t}} &= \mu_t - \frac{\partial \tilde{\Gamma}_{i,t}}{\partial y_{i,t}} = \\ &= \mu_t - \mu_t \sum_{k\geq 1} Q_{t,t+k}\frac{\partial \Gamma_{i,t}}{\partial y_{i,t}} = \\ &= \mu_t - \frac{\mu_t}{(1+r_t)\mu_{t+1}}\frac{\gamma\mu_{t+1}^2}{2}\frac{\partial\sigma_{i,t+1}^2}{\partial y_{i,t}} = \\ &= \mu_t - \frac{1}{1+\mu_{t+1}(1+r_t)}\frac{\gamma\mu_{t+1}^2}{2}\frac{\partial\sigma_{i,t+1}^2}{\partial y_{i,t}}. \end{aligned}$$

It is clear that if $\partial\sigma_{i,t+1}^2/\partial y_{i,t} \neq 0$, then the MPC out of wealth differs from the one out of transitory income: $\mu_t \neq MPC_{i,t}$. Moreover, if $\partial\sigma_{i,t+1}^2/\partial y_{i,t}$ is non-constant with respect to $y_{i,t}$, then the distribution of MPCs is non-degenerate.

■

A.1.3 Proposition 3

Proof. Consider a fully transitory government spending shock g_t (i.e., $g_t \neq 0$, $g_\tau = 0$ for $\tau > t$) and let the monetary authority follow the interest rate peg rule (i.e., $r_\tau = r$, with $v_\tau = 0$ for every $\tau \geq t$). Let

[46] Since $Q_{t,t+k}(\sum_{j\geq k} Q_{t,t+j})^{-1} = \mu_{t+k}$.

[47] Since $-\gamma\mu_{t+1}\epsilon_{i,t+1}|\mathcal{I}_t \sim \mathcal{N}(0,\gamma^2\mu_{t+1}^2\sigma_{i,t+1}^2)$, then $\mathbb{E}_t\big[\exp\{-\gamma\mu_{t+1}\epsilon_{i,t+1}\}\big] = \exp\Big\{\frac{\gamma^2\mu_{t+1}^2}{2}\sigma_{i,t+1}^2\Big\}$.

$\mathcal{H}:\mathbb{R}^2\to\mathbb{R}$ be the function linking aggregate output and spending (for a given $y_{t+1}, g_{t+1}, r_t, \mu_{t+1}$), defined by:

$$\mathcal{H}(y_t,g_t)=y_t-y_{t+1}+\frac{1}{\gamma}\log[\beta(1+r_t)]+\frac{\gamma\mu_{t+1}^2}{2}\int_{[0,1]}\sigma_{i,t+1}^2 di-g_t+g_{t+1}.$$

Notice that at the solution (y_t,g_t) to the equation $\mathcal{H}(y_t,g_t)=0$, the dynamic AD equation (9) holds. By the Implicit Function Theorem, if $\mathcal{H}\in\mathcal{C}^1$ (space of continuously differentiable functions) and $\mathcal{H}_y(y_t,g_t)\neq 0$,[48] then there exist two neighborhoods $\mathcal{U}(y_t)$ and $\mathcal{V}(g_t)$ and a function $\tilde{F}:\mathcal{V}(g_t)\to\mathcal{U}(y_t)$ such that:

$$y_t=\tilde{F}(g_t)\iff\mathcal{H}(y_t,g_t)=0\quad\forall g_t\in\mathcal{V}(g_t),$$

and:

$$\frac{\partial y_t}{\partial g_t}=\tilde{F}'(g_t)=-\frac{\mathcal{H}_g(y_t,g_t)}{\mathcal{H}_y(y_t,g_t)}.$$

Since:

$$\begin{aligned}\mathcal{H}_g(y_t,g_t)&=-1,\\ \mathcal{H}_y(y_t,g_t)&=1+\frac{\gamma\mu_{t+1}^2}{2}\frac{d}{dy_t}\int_{[0,1]}\sigma_{i,t+1}^2\, di,\end{aligned}$$

and assuming that we can exchange the integral and the derivative:[49]

$$\begin{aligned}\frac{d}{dy_t}\int_{[0,1]}\sigma_{i,t+1}^2\, di&=\int_{[0,1]}\left(\frac{\partial\sigma_{i,t+1}^2}{\partial y_{i,t}}\frac{\partial y_{i,t}}{\partial y_t}\right)di=\\ &=\int_{[0,1]}\frac{\partial\sigma_{i,t+1}^2}{\partial y_{i,t}}di\int_{[0,1]}\frac{\partial y_{i,t}}{\partial y_t}di+Cov\left(\frac{\partial\sigma_{i,t+1}^2}{\partial y_{i,t}},\frac{\partial y_{i,t}}{\partial y_t}\right),\end{aligned}$$

we obtain the output multiplier of government spending:

$$\Theta_t:=\frac{\partial y_t}{\partial g_t}=\left\{1+\frac{\gamma\mu_{t+1}^2}{2}\left[\mathbb{E}\left(\frac{\partial\sigma_{i,t+1}^2}{\partial y_{i,t}}\right)\cdot\mathbb{E}\left(\frac{\partial y_{i,t}}{\partial y_t}\right)+Cov\left(\frac{\partial\sigma_{i,t+1}^2}{\partial y_{i,t}},\frac{\partial y_{i,t}}{\partial y_t}\right)\right]\right\}^{-1}.$$

[48] With the notation H_y we mean the partial derivative of H to y.

[49] More formally, let $y_{i,t}=y_{i,t}(y_t)$ be a function of y_t, i.e., household earnings are a function of aggregate output. Then, express income volatility as a function of aggregate output and of the index denoting agent i: $\sigma_{i,t+1}^2=\sigma^2\big(y_{i,t}(y_t)\big)=\tilde{f}(y_t,i)$, where $\tilde{f}:\mathbb{R}\times[0,1]\to[0,\infty)$ is Borel measurable and integrable on $[0,1]$ for every y_t.

a. Assume that for every $i\in[0,1]$, the partial derivative $\partial\sigma_{i,t+1}^2/\partial y_t=\big(\partial\sigma_{i,t+1}^2/\partial y_{i,t}\big)\cdot(\partial y_{i,t}/\partial y_t)$ exists for every $y_t\in\mathbb{R}$.

b. Assume that there exists a Borel-measurable and integrable function $h:[0,1]\to\mathbb{R}$, such that $\left|\partial\sigma_{i,t+1}^2/\partial y_t\right|\leq h(i)$ for every $(y_t,i)\in\mathbb{R}\times[0,1]$.

Then:

$$\frac{d}{dy_t}\int_{[0,1]}\sigma_{i,t+1}^2 di=\int_{[0,1]}\frac{\partial\sigma_{i,t+1}^2}{\partial y_{i,t}}\frac{\partial y_{i,t}}{\partial y_t}di.$$

Following the same steps, we can derive the output multiplier out of a fully transitory interest rate shock:[50]

$$\frac{\partial y_t}{\partial r_t} = -\Theta_t \cdot \frac{1}{\gamma(1+r_t)}. \tag{A-8}$$

To conclude, since Θ_t depends on the three cross-sectional statistics described in Proposition 3, we can state that these are *sufficient* to characterize the IRFs of aggregate output to demand shocks.

■

A.2 Income process, supply-side, and government

We model the supply-side and the government-side of our heteroskedastic economy following Acharya and Dogra (2020).

A.2.1 Income process

We assume that labor productivity is the underlying source of households' idiosyncratic risk. In particular, this structure allows to model the income process in such a way to reproduce two key empirical facts: (1) earnings are highly persistent[51] and (2) the elasticity of individual to aggregate earnings is heterogeneous over the earnings distribution (Guvenen et al., 2017).

We assume that income $y_{i,t}$ is given by the sum of labor earnings, dividends ($d_{i,t}$) and government transfers ($T_{i,t}$). Labor earnings are, in turn, equal to the real wage (w_t) net of taxes (where τ_t is the tax rate) times the total number of hours worked. In particular, total hours are represented by the sum of an agent-specific productivity factor $l^d_{i,t}$ (that evolves deterministically over time) and of a stochastic component $l_{i,t}$.

We assume that the stochastic component of labor productivity $l_{i,t}$ follows a zero-mean Gaussian distribution, conditional on the previous period information set: $l_{i,t} \mid \mathcal{I}_{t-1} \sim \mathcal{N}(0, \sigma^2_{l,i,t})$. The volatility of this process is heterogeneous in the cross-section and depends on lagged deterministic earnings $\sigma^2_{l,i,t} \coloneqq \sigma^2(y^d_{i,t-1})$.[52] Under these assumptions, the part of labor earnings that is determined by the agent-specific productivity factor corresponds to what in Section 2 we call *deterministic income*: $y^d_{i,t} \coloneqq (1-\tau_t) \cdot w_t \cdot l^d_{i,t}$. Meanwhile, the *stochastic component* is: $\epsilon_{i,t} = (1-\tau) \cdot w_t \cdot l_{i,t}$. Therefore:

$$\begin{aligned} y_{i,t} &= \underbrace{(1-\tau_t) \cdot w_t \cdot (l^d_{i,t} + l_{i,t})}_{\text{labor earnings}} + d_{i,t} + T_{i,t} = \\ &= \underbrace{(1-\tau_t) \cdot w_t \cdot l^d_{i,t}}_{y^d_{i,t}} + \underbrace{(1-\tau) \cdot w_t \cdot l_{i,t}}_{\epsilon_{i,t}} + d_{i,t} + T_{i,t}. \end{aligned} \tag{A-9}$$

[50] Equivalently, a shock to v_t.

[51] For instance, Storesletten et al. (2004) estimate a first autocorrelation coefficient of 0.95.

[52] This assumption captures the idea that, for example, richer agents might be exposed to a full set of experiences and opportunities that make them more certain about their future labor productivity.

Finally, we assume that, in every period, average hours are equal to 1: $\int_{[0,1]} l^d_{i,t}\, di = 1$. Under this specification for the income process we can make some considerations.

- The presence of the factor $l^d_{i,t}$ implies that earnings have a persistent nature. Suppose, henceforth, that $l^d_{i,t} = H(y_t, l^d_{i,-1}) = l^d_{i,-1} + h_i(y_t)$ where h_i is an individual-specific function of aggregate output. Under this formulation, $l^d_{i,t}$ is decomposed into an individual productivity *fixed effect* $l^d_{i,-1}$ and a time effect $h_i(y_t)$ (through the function h_i).
- Since that $l^d_{i,t}$ changes over the business cycle, parametrizing $l^d_{i,t}$ as an appropriate function of y_t (i.e., parametrizing the function h_i), we are able to reproduce the *worker betas* (Guvenen et al., 2017), i.e., the elasticities of individual to aggregate earnings. Guvenen et al. (2017) show that *worker betas* are heterogeneous and U-shaped over the earnings distribution. In our setting, we can compute *worker betas* as:[53]

$$\beta_{i,w} = \frac{\partial y^{labor}_{i,t}}{\partial y_t} \cdot \frac{y_t}{y^{labor}_{i,t}} = \frac{\partial w_t}{\partial y_t} \cdot \frac{y_t}{w_t} + \frac{\partial l^d_{i,t}}{\partial y_t} \cdot \frac{y_t}{l^d_{i,t} + l_{i,t}}.$$

- Our income process is akin to a traditional *permanent/transitory* structural income process. The transitory component coincides with our stochastic component $\epsilon_{i,t}$ and the permanent one coincides with our deterministic component $y^d_{i,t}$, that is characterized by the presence of the fixed effect $l^d_{i,-1}$.

To simplify the analysis, assume henceforth uniform taxes and dividends: $\tau_t = \tau$, $d_{i,t} = d_t$, and $T_{i,t} = T_t$.

A.2.2 Technology

There is a continuum of intermediate firms, of measure 1, that produce a variety j of the intermediate good $x_t(j)$, according to the constant return to scale technology:

$$x_t(j) = z_t m_t(j)^\alpha n_t(j)^{1-\alpha}, \tag{A-10}$$

with $\alpha \in (0,1)$. $m_t(j) = (\int_{[0,1]} m_t(j,k)^{\frac{\epsilon-1}{\epsilon}} dk)^{\frac{\epsilon}{\epsilon-1}}$ is the Dixit and Stiglitz (1977) aggregator of the varieties of the intermediate goods that are re-used in the production process ($m_t(j,k)$ represents the variety k purchased by the firm j). Labor employed by firm j is denoted by $n_t(j)$ and z_t is the common total factor productivity (TFP).

Given the constant elasticity of substitution (CES) aggregator for the intermediate good m_t, the demand function for variety k by firm j is:

$$m_t(j,k) = \left(\frac{p_t(k)}{p_t}\right)^{-\epsilon} m_t(j),$$

[53] Let us ignore dividends and transfers and consider only labor earnings $y^{labor}_{i,t} = (1-\tau_t)\cdot w_t \cdot (l^d_{i,t} + l_{i,t})$.

with $p_t := (\int_{[0,1]} p_t(k)^{1-\epsilon} dk)^{\frac{1}{1-\epsilon}}$ being the aggregate price index. Then, total demand for variety k of the intermediate good is given by:

$$m_t(k) := \int_{[0,1]} m_t(j,k)dj = \left(\frac{p_t(k)}{p_t}\right)^{-\epsilon} \underbrace{\int_{[0,1]} m_t(j)dj}_{m_t:=} = \left(\frac{p_t(k)}{p_t}\right)^{-\epsilon} m_t.$$

Finally, there is a perfectly competitive firm producing a final good y_t according to the CES technology $y_t = (\int_{[0,1]} y_t(k)^{\frac{\epsilon-1}{\epsilon}} dk)^{\frac{\epsilon}{\epsilon-1}}$. Then, the demand function from the final good producer for each variety k is:

$$y_t(k) = \left(\frac{p_t(k)}{p_t}\right)^{-\epsilon} y_t.$$

Summarizing, good $x_t(k)$ (produced by the intermediate good firm k) is purchased in part by other intermediate good firms ($m_t(k)$), and in part by the final good firm ($y_t(k)$):

$$x_t(k) = m_t(k) + y_t(k) = \left(\frac{p_t(k)}{p_t}\right)^{-\epsilon} \underbrace{(m_t + y_t)}_{x_t:=} = \left(\frac{p_t(k)}{p_t}\right)^{-\epsilon} x_t. \tag{A-11}$$

Therefore, y_t can be also interpreted as net output: $y_t = x_t - m_t$.

A.2.3 Cost minimization

Each intermediate producer purchases the aggregated intermediate goods and hires workers, solving the cost minimization problem (expressed in real terms):

$$\begin{aligned} \mathcal{C}_t(x_t(j)) := \min_{(m_t(j), n_t(j))} & \left[\frac{p_t(j)}{p_t} m_t(j) + w_t n_t(j)\right] \\ s.t. \qquad & x_t(j) \leq z_t m_t(j)^{\alpha} n_t(j)^{1-\alpha}. \end{aligned} \tag{A-12}$$

We get the following optimality conditions, considering a symmetric equilibrium[54] and imposing the labor market clearing, i.e., $\int_{[0,1]} n_t(j)dj = \int_{[0,1]} (l^d_{i,t} + l_{i,t})di = 1$:

$$x_t = z_t m_t^{\alpha}, \tag{A-13a}$$

$$w_t = \frac{1-\alpha}{\alpha}\left(\frac{x_t}{z_t}\right)^{\frac{1}{\alpha}}, \tag{A-13b}$$

$$\mathcal{C}_t(x_t) = \frac{1}{\alpha}\left(\frac{x_t}{z_t}\right)^{\frac{1}{\alpha}}, \tag{A-13c}$$

where (A-13a) is the production function, (A-13b) is the optimality condition (MRTS = real wage) and (A-13c) is the (real) cost function.

[54]The symmetry is implied by the price stickiness à la Rotemberg (1982), that we assume here. See Section A.2.4.

Finally, we can determine an equation linking aggregate output and the real wage, by combining equations (A-13a), (A-13b) and the definition of net output.

$$y_t = z_t \cdot \left(\frac{\alpha}{1-\alpha}\right)^{\alpha} \cdot w_t^{\alpha} - \frac{\alpha}{1-\alpha} \cdot w_t. \tag{A-14}$$

We can, then, determine aggregate real dividends:

$$d_t = x_t - \mathcal{C}_t(x_t) = x_t - m_t - w_t = y_t - w_t. \tag{A-15}$$

A.2.4 Nominal frictions

We model nominal rigidities following Rotemberg (1982). Each intermediate good firm sets prices in every period paying a quadratic convex adjustment cost (rebated lump-sum and uniformly to households)[55] and solves the following recursive maximization problem:

$$V_t^f(p_{t-1}(j)) = \max_{p_t(j)} \left\{ \left[\frac{p_t(j)}{p_t} x_t(j) - \mathcal{C}_t(x_t(j)) - \frac{\psi}{2}\left(\frac{p_t(j)}{p_{t-1}(j)} - 1\right)^2 x_t \right] + \frac{1}{1+r_t} V_{t+1}^f(p_t(j)) \right\}$$
$$s.t \quad x_t(j) = \left(\frac{p_t(j)}{p_t}\right)^{-\epsilon} x_t,$$

where ψ regulates the size of the adjustment cost (hence, the degree of price rigidity in the economy) and $V_t^f(\cdot)$ is the value function of the firm optimization problem.

From the optimality conditions, exploiting the symmetry of the equilibrium, we obtain the non-linear PC:

$$\psi(1+\pi_t)\pi_t = 1 - \epsilon(1 - \mathcal{MC}_t(x_t)) + \psi \cdot \frac{1}{1+r_t}(1+\pi_{t+1})\pi_{t+1}\frac{x_{t+1}}{x_t}, \tag{A-16}$$

where $\pi_t := \frac{p_t}{p_{t-1}} - 1$ and $\mathcal{MC}_t(x_t) := \frac{\partial \mathcal{C}_t}{\partial x_t} = \alpha^{-2} z_t^{-1/\alpha} x_t^{\frac{1-\alpha}{\alpha}}$.

Linearizing equation (A-16) around the zero-inflation steady state (setting $z_t = z$ and normalizing $y = 1$) described in Section 2.5.1, we obtain the linearized PC of equation (19c):

$$\hat{\pi}_t = \tilde{\beta} \cdot \hat{\pi}_{t+1} + k \cdot \hat{y}_t,$$

where the *hat* notation denotes variables in deviation from their steady-state value, $\tilde{\beta} := \frac{1}{1+r}$, and

[55] Following Ascari and Rossi (2012), footnote 17, this means that the true households' budget constraint should be:

$$c_{i,t} + \frac{a_{i,t+1}}{1+r_t} = y_{i,t} + a_{i,t} + \frac{\psi}{2}\left(\frac{p_t(j)}{p_{t-1}(j)} - 1\right)^2 x_t.$$

This assumption implies that in equilibrium there is no waste from price adjustment costs and $y_t = c_t + g_t$.

the slope of the PC is:[56]

$$k := \frac{\epsilon}{\psi}\frac{1-\alpha}{\alpha^2}\frac{1}{x}\left(\alpha z^{\frac{1}{\alpha}} x^{1-\frac{1}{\alpha}} - 1\right)^{-1}.$$

A.2.5 Government

The government purchases the final good as g_t and pays lump-sum transfers T_t to households. In addition, it can raise proportional labor income taxes $\tau_t w_t$ and finances its deficits with government debt (in real terms) b_t. The intra-temporal government budget constraint is:

$$\frac{b_{t+1}}{1+r_t} = b_t - s_t, \tag{A-17}$$

where s_t is primary surplus:

$$s_t = \tau_t w_t - g_t - T_t. \tag{A-18}$$

A.2.6 Equilibrium

Markets clear in general equilibrium.

- Financial market clearing: $a_t = \int_{[0,1]} a_{i,t}\ di = b_t$.
- Labor market clearing: $n_t = \int_{[0,1]} n_t(j)\ dj = \int_{[0,1]} (l_{i,t}^d + l_{i,t})\ di = 1$.
- Goods market clearing for each intermediate good: $x_t(k) = y_t(k) + m_t(k) \implies x_t = y_t + m_t$.
- Final good market clearing: $y_t = c_t + g_t$ where $c_t = \int_{[0,1]} c_{i,t}\ di$.

Notice that we are integrating with respect to the Lebesgue measure. In particular, let $\bar{y}_t$ be average household earnings, defined as $\bar{y}_t := \int_{[0,1]} y_{i,t}\ di = (1-\tau_t)w_t + d_t + T_t$.[57], then (exploiting the definition of $\bar{y}_t$, d_t and the financial market clearing) we can prove that the following relations hold in general equilibrium:

$$\bar{y}_t + s_t = w_t + d_t - g_t,$$
$$y_t = s_t + \bar{y}_t + g_t.$$

[56] This expression can be obtained by combining the linearized versions of equations (A-13a) and $y_t = x_t - m_t$ (since $\hat{y}_t = x \cdot \hat{x}_t - m \cdot \hat{m}_t$ and $\hat{x}_t = \alpha \cdot \hat{m}_t$, with the steady-state value of x solving $x - (x/z)^{1/\alpha} = 1$):

$$\hat{x}_t = \left[x - \frac{1}{\alpha}\left(\frac{x}{z}\right)^{\frac{1}{\alpha}}\right]^{-1} \hat{y}_t.$$

The expression above can then be used to substitute $\hat{x}_t$ with $\hat{y}_t$ in:

$$\frac{\epsilon}{\psi}\mathcal{MC}_t'(x) \cdot x \cdot \hat{x}_t = \frac{\epsilon}{\psi}\frac{1-\alpha}{\alpha^3}\frac{1}{z^{\frac{1}{\alpha}}} x^{\frac{1-\alpha}{\alpha}} \hat{x}_t.$$

[57] Since, by a continuous version of the LLN, $\int_{[0,1]} l_{i,t}\ di = 0$. See Uhlig (1988), Proposition 1.

A.3 Steady state

In this section we prove the statements of Section 2.5.1, where we describe the steady state of the model. Recall that we set $g = 0$ and normalize aggregate consumption and output to 1 in the steady state. The real wage is determined from equation (A-14),[58] given the parameters $\alpha \in (0,1), z > 0$. In addition, the deterministic component of labor productivity follows: $l^d_{i,t} = \bar{l}_{i,-1} + h_i(y_t)$ with $h_i(y) = 0$ (where y is steady-state output).

Lemma A.1. *Let us assume that the initial fixed effects $l^d_{i,-1}$ are i.i.d. draws from the cumulative distribution function (CDF) $\bar{F} : \mathbb{R} \to [0,1]$. Then, $F : \mathbb{R} \to [0,1]$, the CDF describing the stationary distribution of (deterministic) income, exists in the steady state.*

Proof. Let $y = 1$ and the real wage w be determined from (A-14). The expression for deterministic income is:

$$y^d_i = (1-\tau)w \cdot l^d_i = (1-\tau)w \cdot l^d_{i,-1},$$

since in the steady state we take the permanent component of labor productivity to be equal to the initial extraction of the fixed effect: $l^d_i = l^d_{i,-1}$. Since y^d_i is a scalar transformation of $l^d_{i,-1}$, for every $l^d_{i,-1}$, let $F : \mathbb{R} \to [0,1]$ be defined as:

$$F(y^d_i) = \bar{F}(l^d_{i,-1}) \quad \text{with} \quad y^d_i = (1-\tau)w \cdot l^d_{i,-1}.$$

Then, since $\bar{F}$ is right-continuous and non-decreasing (as it is a CDF) and y^d_i is a linear transformation of $l^d_{i,-1}$, then F is also right-continuous and non-decreasing. Therefore, it can be proved that F is the CDF of a unique probability measure on $\mathbb{R}$, describing the distribution of deterministic income.

■

Lemma A.2. *Consider the measure space $([0,1], \mathcal{B}([0,1]), \nu)$, where ν is the Lebesgue measure, and the measurable space $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ where $\mathcal{B}$ denotes the Borel sigma-algebra. Let us assume that the function $y^d : [0,1] \to \mathbb{R}$ ($i \mapsto y^d_i$) is measurable. Then, the CDF of deterministic income, $F : \mathbb{R} \to [0,1]$, represents the Lebesgue-Stieltjes measure of the push-forward of ν through y^d.*

Therefore, we can equivalently compute the average volatility weighting either by the CDF of deterministic income F or by the Lebesgue measure ν:

$$\bar{\sigma}^2 := \int_{\mathbb{R}} \sigma^2(y) \; dF(y) = \int_{[0,1]} \sigma^2(y^d_i) \; \nu(di).$$

[58] Notice that this equation admits more than one solution. We consider the lowest possible solution for the real wage.

Proof. By definition[59] of *push-forward* measure,[60] that we denote by $y^d\#\nu$, we have that:

$$(y^d\#\nu)(B) = \nu\Big(\underbrace{(y^d)^{-1}(B)}_{\substack{\text{counter-image} \\ \text{of } B \text{ through } y^d}} \Big) \quad \text{for every } B \in \mathcal{B}(\mathbb{R}).$$

Notice that $y^d\#\nu$ is a measure on the target space $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$. We want to show that this measure admits a unique CDF representation, coinciding with F. Consider the *semi-algebra* defined by the semi-intervals of the kind $(-\infty, y]$. Then, we have:

$$\begin{aligned}(y^d\#\nu)\left((-\infty, y]\right) =&\nu\Big((y^d)^{-1}\left((-\infty, y]\right)\Big) = \\ =&\nu\Big(\{i \in [0,1] : y_i^d \leq y\}\Big) = \\ =& \int_{[0,1]} \mathbb{I}\left(y_i^d \in (-\infty, y]\right)\nu(di) = {}^{61} \\ =&F(y).\end{aligned}$$

The first equality is by definition of push-forward measure and the second one is by definition of counter-image. The third one follows, instead, from the definition of the Lebesgue integral and from the fact that $\mathbb{I}\left(y_i^d \in (-\infty, y]\right) = 1 \iff i \in \{i : y_i^d \leq y\}$. Finally, the last equality comes from a continuous version of the Law of Large Number.[62] In fact, let $Z_i := \mathbb{I}\left(y_i^d \in (-\infty, y]\right)$ and notice that the variables of the collection $(Z_i)_{[0,1]}$ are identically distributed and pairwise independent, with mean $\mathbb{E}(Z_i) = \mathbb{P}\left(y_i^d \in (-\infty, y]\right) = F(y)$.

From the previous equalities, it follows that $y^d\#\nu$ is a probability measure (since $(y^d\#\nu)(\mathbb{R}) = 1$) and, on the intervals of the kind $(-\infty, y]$, it coincides with F. Then, by Carathéodory extension theorem, F must be the unique representation of $y^d\#\nu$.[63]

Finally, it can be shown that the variance function $\sigma^2 : \mathbb{R} \to \mathbb{R}_+$ is integrable with respect to F *if and only if* the composition $\sigma^2 \circ y^d : [0,1] \to R_+$ is integrable with respect to ν. Moreover:

$$\int_{\mathbb{R}} \sigma^2(y)\ dF(y) = \int_{[0,1]} \sigma^2(y_i^d)\ \nu(di).^{64}$$

■

Lemma A.3. *Existence and uniqueness of the real interest rate. If $\bar{\sigma}^2 > 0$, the solution r to equation (16) exists and is unique.*

[59] Since the function y^d is measurable, the definition below is meaningful.

[60] The push-forward measure allows to define a measure on a target measurable space – here: $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ – starting from a measure on an original space – here: ν on $([0,1], \mathcal{B}([0,1]))$ – through a function that maps from the original to the target space – here: y^d.

[61] $\mathbb{I}$ is an indicator function.

[62] See Uhlig (1988), Proposition 1.

[63] More precisely: the probability measure represented by F (that is unique) must coincide with $y^d\#\nu$. In other words, $y^d\#\nu$ is the probability measure describing the stationary distribution of deterministic income.

[64] In the paper we have denoted the integration with respect to the Lebesgue measure with di instead of $\nu(di)$.

Proof. We rewrite here the equation determining the real interest rate in the steady state:

$$\left(\frac{1+r}{r}\right)^2 \log[\beta\cdot(1+r)] = -\frac{\gamma^2}{2}\bar{\sigma}^2.$$

Notice that in equation (16) since $\bar{\sigma}^2 > 0$ then $\beta(1+r) < 1$ (a standard result in the incomplete markets literature to bound the solution to the income fluctuation problem). Then, note that the LHS of equation (16) is continuous and strictly increasing in r for $r \in (0, 1/\beta - 1]$. Moreover, since the LHS $\to -\infty$ when $r \to 0$ and LHS $\to 0$ when $r \to 1/\beta - 1$, then there exists a unique solution r to equation (16).

■

A.4 Equilibrium determinacy

Assume that monetary policy follows the Taylor rule, equation (12). Rewriting the linearized model - equations (19) - in matrix form,

$$\begin{pmatrix} \hat{y}_{t+1} \\ \hat{\mu}_{t+1} \\ \hat{\pi}_{t+1} \end{pmatrix} = \underbrace{\begin{bmatrix} \Theta^{-1} + \frac{k}{\tilde{\beta}}(\gamma^{-1} - \Lambda) & \frac{\Lambda}{\tilde{\beta}} & (\gamma^{-1} - \Lambda)\left(\phi_\pi - \frac{1}{\tilde{\beta}}\right) \\ -\frac{k}{\tilde{\beta}} & \frac{1}{\tilde{\beta}} & \frac{1}{\tilde{\beta}} - \phi_\pi \\ -\frac{k}{\tilde{\beta}} & 0 & \frac{1}{\tilde{\beta}} \end{bmatrix}}_{\Omega :=} \begin{pmatrix} \hat{y}_t \\ \hat{\mu}_t \\ \hat{\pi}_t \end{pmatrix} + \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \Delta\hat{g}_{t+1} + \begin{bmatrix} \gamma^{-1} - \Lambda \\ -1 \\ 0 \end{bmatrix} v_t,$$

determinacy requires that the matrix Ω has all its three eigenvalues outside the unit circle since the three variables $\hat{y}_t, \hat{\mu}_t, \hat{\pi}_t$ are forward-looking. We can obtain the characteristic polynomial of the matrix Ω as $\mathcal{P}(z) = \det(\Omega - I_3 \cdot z)$.[65]

$$\mathcal{P}(z) = -z^3 + \frac{2\tilde{\beta} + \tilde{\beta}^2\Theta^{-1} + k\tilde{\beta}\cdot(\gamma^{-1} - \Lambda)}{\tilde{\beta}^2}\cdot z^2 - \\ -\frac{1 + 2\tilde{\beta}\Theta^{-1} + k\gamma^{-1} + \tilde{\beta}k\phi_\pi\cdot(\gamma^{-1} - \Lambda)}{\tilde{\beta}^2}\cdot z + \frac{\Theta^{-1} + k\gamma^{-1}\phi_\pi}{\tilde{\beta}^2}.$$

The eigenvalues of the matrix solve the equation $\mathcal{P}(z^*) = 0$. Necessary and sufficient conditions for equilibrium determinacy are that $|z_i^*| > 1$ for $i = 1, 2, 3$.

We can find a necessary condition for the eigenvalues to be outside the unit circle under the following parametric assumption.

Assumption 5. *Assume that $\gamma^{-1} \geq \Lambda$.*[66]

[65] I_3 denotes the 3×3 identity matrix.
[66] Under our calibrations (see Section 2.5.5) assumption 5 holds.

Proposition 4. *Under assumption 5, the following condition:*

$$\phi_\pi > 1 + \frac{1}{k}\left[\frac{\gamma \cdot (1-\tilde{\beta})^2}{1-\tilde{\beta}+\tilde{\beta}\cdot\gamma\Lambda}\right]\left(1-\frac{1}{\Theta}\right) =: \bar{\phi}_\pi, \tag{A-19}$$

is necessary for equilibrium determinacy since it is necessary[67] *for the three eigenvalues to be outside the unit circle.*

Proof. The polynomial $\mathcal{P}(z)$ is continuous over the real line. Moreover, $\mathcal{P}(z) \to +\infty$ as $z \to -\infty$ and $\mathcal{P}(z) \to -\infty$ as $z \to +\infty$. Under assumption 5, for every $z \le 0$, $\mathcal{P}(z) > 0$, therefore the root of the polynomial must be positive.
$\mathcal{P}(1) > 0$ is a necessary condition for the roots to be outside the unit circle. By contraposition, if $\mathcal{P}(1) \le 0$, then either $\mathcal{P}(1) = 0$ or $\mathcal{P}(1) < 0$. In latter case, the Intermediate Value Theorem[68] implies that at least one root z_i^* must be between $(0, 1)$.
Therefore, imposing $\mathcal{P}(1) > 0$ and rearranging, we get condition (A-19).

■

Notice that when $\Theta = 1$, as in the RA and homoskedastic models, this necessary condition collapses to the well-known *Taylor principle*: $\phi_\pi > 1$. In addition, since the RHS of condition (A-19) – $\bar{\phi}_\pi$ – is increasing in Θ, the more demand shocks are amplified, the stronger the reaction of monetary policy must be to stabilize the economy.

A.5 Multipliers with persistent shocks

Applying the method of undetermined coefficient, as explained in Section 2.5.4, we can obtain the multipliers of fiscal and monetary policy when shocks are persistent. For the *fiscal shock*, we have:

$$\begin{aligned}
\psi_y^g &= \frac{\gamma\Theta \cdot (1-\rho_g)\cdot(1-\tilde{\beta}\rho_g)^2}{\gamma\cdot(1-\tilde{\beta}\rho_g)^2\cdot(1-\rho_g\Theta) + \Theta\cdot[1-\tilde{\beta}\rho_g(1-\Lambda\gamma)]\cdot(\phi_\pi-\rho_g)\cdot k} > 0,\\
\psi_\mu^g &= \frac{(\phi_\pi-\rho_g)\cdot\tilde{\beta}k}{(1-\tilde{\beta}\rho_g)^2}\cdot\psi_y^g > 0,\\
\psi_\pi^g &= \frac{k}{1-\tilde{\beta}\rho_g}\cdot\psi_y^g > 0,
\end{aligned} \tag{A-20}$$

[67] Acharya and Dogra (2020) show that under additional assumptions on the discount factor being sufficiently large and Θ not being too large, condition (A-19) is also sufficient for all the eigenvalues to be outside the unit circle. See Acharya and Dogra (2020), Appendix B.3.

[68] $\mathcal{P}(z)$ is continuous over the interval $[0, 1]$ with $\mathcal{P}(0) > 0$ and $\mathcal{P}(1) < 0$.

and for the *monetary shock*:

$$\psi_y^v = -\frac{\Theta \cdot [1-\tilde{\beta}\rho_v(1-\Lambda\gamma)] \cdot (1-\tilde{\beta}\rho_v)}{\gamma \cdot (1-\tilde{\beta}\rho_v)^2 \cdot (1-\rho_v\Theta) + \Theta \cdot [1-\tilde{\beta}\rho_v(1-\Lambda\gamma)] \cdot (\phi_\pi - \rho_v) \cdot k} < 0,$$
$$\psi_\mu^v = \frac{\tilde{\beta}}{1-\tilde{\beta}\rho_v} + \frac{(\phi_\pi - \rho_v) \cdot \tilde{\beta}k}{(1-\tilde{\beta}\rho_v)^2} \cdot \psi_y^v > 0, \tag{A-21}$$
$$\psi_\pi^v = \frac{k}{1-\tilde{\beta}\rho_v} \cdot \psi_y^v < 0.$$

The multipliers for i.i.d. shocks, as in Section 2.5.4, can be obtained imposing $\rho_g = \rho_v = 0$. The same qualitative implications for the effects of monetary and fiscal shocks (that we have in the i.i.d. case) can be derived in the more general case of persistent shocks. In particular, notice that higher values for Θ imply larger multipliers for output and inflation, in absolute value.

B Incomplete Information

B.1 Two-period setup, derivations

As in the model with complete information, we consider a CARA-utility economy with a continuum of households, of measure 1 on the interval $[0,1]$ indexed by $i \in [0,1]$.

B.1.1 Consumer problem

Now, the economy runs for two periods. In the first period, denoted by t, each household can purchase a risk-free (real) bond $a_{i,t+1}$ (paying one unit of good in $t+1$) at the price $1/(1+r_t)$. Additionally, the household can consume a perishable consumption good $c_{i,t}$, using its total income $y_{i,t}$ net of a uniform lump-sum tax T_t,[69] and its initial wealth $a_{i,t}$. Then in the second period, denoted by $t+1$, the household consumes its resources entirely. Households maximize their intertemporal utility – discounting with a factor $\beta \in (0,1)$ – choosing consumption and savings subject to two intratemporal budget constraints.

$$\max_{(c_{i,t}, c_{i,t+1}, a_{i,t+1})} \quad u(c_{i,t}) + \beta \mathbb{E}_{i,t}[u(c_{i,t+1})]$$
$$s.t \quad c_{i,t} + \frac{a_{i,t+1}}{1+r_t} = y_{i,t} - T_t + a_{i,t},$$
$$c_{i,t+1} = y_{i,t+1} - T_{t+1} + a_{i,t+1}.$$

[69] We assume that the government purchases the final good g_t, g_{t+1} and raises taxes with a balanced budget in every period ($\tau = t, t+1$):

$$T_\tau = g_\tau.$$

We could equivalently allow the government to issue bonds in the first period, but this will not change the equilibrium of the model and the results that we obtain.

Note that the expectation operator is conditional on agent i's information set. Since agents do not observe the fundamental variance of the income process and they face uncertainty over future income $y_{i,t+1}$, they compute their future expectations with respect to the Gaussian distribution:

$$y_{i,t+1} \mid \mathcal{I}_{i,t} \sim \mathcal{N}(y_{t+1}, \hat{\sigma}_i^2), \tag{B-1}$$

where the relevant variance is the point estimate $\hat{\sigma}_i^2$ obtained in (35) through Bayesian learning.

Consider the usual CARA-exponential utility function: $u(c) = -\frac{1}{\gamma}\exp\{-\gamma c\}$. Then, the Euler equation of the utility maximization problem is:

$$\exp\{-\gamma c_{i,t}\} = \beta(1+r_t) \cdot \mathbb{E}_{i,t}\big[\exp\{-\gamma c_{i,t+1}\}\big]. \tag{B-2}$$

Equation (B-2), together with the two budget constraints, characterize the solution of the problem.

It is possible to find an explicit solution for consumption by taking the logs of equation (B-2), using the moments of the log-normal distribution, and the two budget constraints to substitute for $c_{i,t+1}$ and $a_{i,t+1}$:

$$c_{i,t} = -\frac{1}{(2+r_t)\gamma}\log[\beta(1+r_t)] + \frac{1+r_t}{2+r_t}\left[a_{i,t} + y_{i,t} - T_t + \frac{1}{1+r_t}(y_{t+1} - T_{t+1})\right] - \frac{1}{2+r_t}\frac{\gamma}{2}\hat{\sigma}_i^2. \tag{B-3}$$

The above equation represents a two-period version of the consumption function obtained in the complete information infinite-horizon model (equation (4)). The key difference is that the precautionary saving motive on the right hand side features the *estimated* variance of the income process rather than the true variance.

MPCs out of transitory income shocks and out of wealth can be obtained by taking derivatives of equation (B-3). They are, respectively:

$$MPC_{i,t} = \frac{\partial c_{i,t}}{\partial y_{i,t}} = \frac{1+r_t}{2+r_t} - \frac{1}{2+r_t}\frac{\gamma}{2}\frac{\partial \hat{\sigma}_i^2}{\partial y_{i,t}},$$

and:

$$MPC_{i,t}^{\text{wealth}} = \frac{\partial c_{i,t}}{\partial a_{i,t}} = \frac{1+r_t}{2+r_t}. \tag{B-4}$$

B.1.2 Aggregation and general equilibrium

Let us aggregate the consumption function (B-3) by integrating over the unit interval with respect to the Lebesgue measure to obtain an aggregate demand function:

$$c_t = -\frac{1}{(2+r_t)\gamma}\log[\beta(1+r_t)] + \frac{1+r_t}{2+r_t}(y_t - g_t) + \frac{1}{2+r_t}(y_{t+1} - g_{t+1}) - \frac{1}{2+r_t}\frac{\gamma}{2}\int_{[0,1]}\hat{\sigma}_i^2 di. \tag{B-5}$$

Notice that we obtained the previous equation (B-5) by imposing the bond market clearing:

$$\int_{[0,1]} a_{i,t} \; di = 0,$$

using the government budget constraints ($T_\tau = g_\tau$ for $\tau = t, t+1$), and by assuming that individual income integrates to aggregate output:

$$\int_{[0,1]} y_{i,t} \; di = y_t.$$

In what follows we postulate that output is endogenous. Moreover, we also assume that when a shock hits the economy, movements in aggregate output shift the entire income distribution with individual earnings sensitivities that are modeled as heterogeneous. More formally, assume that, for every agent $i \in [0,1]$:

$$y_{i,t} = h_i(y_t), \tag{B-6}$$

where $h : \mathbb{R} \to \mathbb{R}$ is some real-valued function, that we assume differentiable and integrable. Hence, the sensitivity of individual to aggregate income is:

$$\frac{\partial y_{i,t}}{\partial y_t} = h_i'(y_t).\text{[70]} \tag{B-7}$$

This last set of assumptions and notations is convenient for the purpose of our analysis, since we want to study how the income distribution interacts with revisions in the estimates of income volatility in a general equilibrium framework. Finally, imposing the goods market clearing – $y_t = c_t + g_t$ – we obtain the two-period version of the dynamic AD equation (9):

$$y_t = y_{t+1} - \frac{1}{\gamma} \log[\beta(1+r_t)] - \frac{\gamma}{2} \int_{[0,1]} \hat{\sigma}_i^2 di + g_t - g_{t+1}. \tag{B-8}$$

B.2 Bayesian updating with Inverse-Gamma

Consider the following model, where the sample $X = (X_1, ..., X_n)$ is i.i.d. from the Gaussian distribution, conditional on the variance σ^2: $X_i \mid \sigma^2 \sim \mathcal{N}(0, \sigma^2)$. Moreover, the prior on σ^2 is an Inverse-Gamma (IG) distribution with hyperparameters $\alpha > 0$, $\beta > 0$, that corresponds to the conjugate prior for the variance of the Gaussian distribution.

$$X_i \mid \sigma^2 \sim \mathcal{N}(0, \sigma^2),$$
$$\sigma^2 \sim \text{Inv-Gamma}(\alpha, \beta).$$

[70] For instance, we can model income realizations as the sum of two components:

$$y_{i,t} = h_i' \cdot y_t + \epsilon_{i,t},$$

where $\epsilon_{i,t} \sim \mathcal{N}\left((1 - h_i') \cdot y_t, \sigma^2\right)$ and h_i' is the sensitivity of individual to aggregate income. In this way, changes in aggregate output affect individual earnings in a deterministic way. Notice that $y_{i,t} \sim \mathcal{N}\left(y_t, \sigma^2\right)$.

The IG has the following density function $q(\cdot)$:

$$q(\sigma^2) = \frac{\beta^\alpha}{\Gamma(\alpha)} \left(\sigma^2\right)^{-\alpha-1} \exp\left(-\frac{\beta}{\sigma^2}\right),$$

where $\Gamma(\cdot)$ is the Gamma function. The expected value and the variance of the IG distribution are, respectively: $\mathbb{E}(\sigma^2) = \frac{\beta}{\alpha-1}$ and $\text{Var}(\sigma^2) = \frac{\beta^2}{(\alpha-1)^2(\alpha-2)}$.[71]

The likelihood function of the sample $x = (x_1, ..., x_n)$ can be rewritten as:

$$q(x|\sigma^2) \propto \left(\sigma^2\right)^{-n/2} \exp\left(-\frac{n}{2\sigma^2}\frac{1}{n}\sum_{i=1}^{n} x_i^2\right).$$

Let us define $w^2 := \frac{1}{n}\sum_{i=1}^{n} x_i^2$ and consider the Bayesian updating stage given the sample x:

$$q(\sigma^2|x) \propto q(x|\sigma^2) \cdot q(\sigma^2) \propto \left(\sigma^2\right)^{-n/2} \exp\left(-\frac{n}{2\sigma^2}w^2\right) \cdot \left(\sigma^2\right)^{-\alpha-1} \exp\left(-\frac{\beta}{\sigma^2}\right) =$$
$$= \left(\sigma^2\right)^{-\alpha-n/2-1} \exp\left(-\frac{\beta + nw^2/2}{\sigma^2}\right).$$

Therefore, the posterior follows an IG distribution:

$$\sigma^2 \mid X = x \sim \text{Inv-Gamma}\left(\alpha + \frac{n}{2}, \beta + \frac{nw^2}{2}\right).$$

Finally, notice that in our tractable example of Section 3.3, $\alpha = \nu_0/2$ and $\beta = \nu_0\sigma_0^2/2$, while $n = 1$ and $w^2 = (y_{i,t} - y_t)^2$.

[71] The existence of these moments requires $\alpha > 1$ and $\alpha > 2$, respectively.