Heterogeneity and Aggregate Fluctuations: Insights from TANK Models[1]

A Proposal for the NBER Macroeconomics Annual Conference

Davide Debortoli[2] and Jordi Galí[3]

August 2023

[Do Not Quote or Circulate without Permission]

The dominance of the representative household paradigm in macroeconomics has been challenged in recent years by a growing literature that argues that such an assumption, while convenient on tractability grounds, may be less innocuous than one may think, even when the focus is on understanding aggregate fluctuations and their interaction with macro policies. An emerging class of models, often referred to as HANK (for Heterogeneous Agent New Keynesian), have been used to explore the implications of heterogeneity for economic fluctuations and the design of policy rules.[4] Heterogeneity in those models is usually associated with the presence of idiosyncratic income shocks experienced by households. Those shocks cannot be insured against due to incomplete financial markets, as well as borrowing constraints. Those features of the household block are then combined with a supply block characterized by the kind of nominal rigidities and consequent monetary non-neutralities that are the hallmark of New Keynesian (NK) models.

The research proposed here seeks to advance our knowledge on the implications of heterogeneity for aggregate fluctuations and their interaction with monetary policy. More specifically, we want to study the role of the different propagation channels found in HANK models in accounting for the latter's differential implications relative to the standard NK model with a representative agent (RANK). Our ultimate goal is twofold. Firstly, we want to understand the mechanisms through which heterogeneity affects aggregate fluctuations and the transmission of shocks in HANK models. As discussed below, we plan to do so by comparing the properties of a range of nested HANK models, a strategy that allows us to isolate the role played by different elements found in conventional versions of those models. Secondly, we want to investigate whether there are simple tractable models that can capture reasonably well the mechanisms of HANK models that we identify as most relevant.

Throughout, our analysis relies heavily on the distinction between hand-to-mouth and unconstrained households. That distinction arises endogenously in conventional HANK model at any point in time. It is also central to Two-Agent New Keynesian (TANK) models, though in the latter it is introduced in a rather stark way.[5] While the tractability and transparency of TANK models is generally viewed as an advantage relative to conventional HANK models, it is not clear that they can capture well the aggregate implications of the richer HANK models. This motivates a key objective of our analysis, namely, the evaluation of the ability of TANK models to approximate the aggregate equilibrium dynamics generated by HANK models.

To be clear, we are not the first to compare the properties of a baseline HANK model to those of simpler, more tractable frameworks.[6] The key difference lies in the particular approach we use to doing so, which relies on the differential behavior of hand-to-mouth and unconstrained households. It is in that sense that analyze the properties of HANK models through the lens of a TANK framework, thus motivating our title. We believe this provides an interesting and useful perspective to understanding the mechanisms at work in HANK models.

Next we describe the main elements of our approach.

[1] The present proposal shares the spirit (though it is very different in the details) of an earlier working paper of ours titled "Monetary Policy with Heterogeneous Agents: Insights from TANK Models," dated July 2018. That paper was presented at several conferences and drew a good amount of attention (it has 272 Google citations as of today) but it was never submitted to any journal and was instead put aside in order to work on a separate but related project (Debortoli and Galí (2023)) which we viewed as a logical preliminary step.

[2] ICREA-UPF, CREi and BSE

[3] CREi, UPF and BSE

[4] See, e.g. Kaplan et al (2018) and Auclert et al. (2023) for examples of such models.

[5] See, e.g., Galí et al. (2007) and Bilbiie (2008). The idea of partitioning households between hand-to-mouth and unconstrained can be traced back to Campbell and Mankiw (1989).

[6] See, e.g. Auclert, Rognlie and Straub (2023)

1 A Kit of HANK Models

In the present section we describe a variety of HANK models, which differ in terms of the assumptions made on their household block. The supply side, described further below, is common to all the models.

1.1 Households

Next we describe the household block of a variety of HANK models. All the models considered assume a continuum of infinite-lived households, indexed by $j \in [0,1]$. Each household seeks to maximize utility $\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t(j), \mathcal{N}_t(j); Z_t)$, where $C_t(j) \equiv \left(\int_0^1 C_t(i,j)^{1-\frac{1}{\epsilon}} di \right)^{\frac{\epsilon}{\epsilon-1}}$ is an index of the quantities consumed of the different available goods, $\mathcal{N}_t(j)$ denotes work hours, and Z_t is an exogenous preference shifter following a stationary process . We assume $U(C, \mathcal{N}; Z) \equiv \left(\frac{C^{1-\sigma}-1}{1-\sigma} - \frac{\mathcal{N}^{1+\varphi}}{1+\varphi} \right) Z$.
A second common feature of all the models considered is the coexistence of two types of assets. The first asset is a one-period bond, which yields a riskless real return. Bonds are assumed to be *fully liquid*, i.e. they can be bought and sold in a competitive market with no transaction costs. In our baseline specification, they are assumed to be in zero net supply. The second type of asset are shares in an equity fund that holds the stocks of all firms in the economy. In the models considered in the present proposal, those shares are *fully illiquid*, i.e. they cannot be traded. Each period they yield a payoff to their holders, equal to firms' aggregate dividends. The previous two-asset structure captures in a simple way a feature of many HANK models in the literature. [7]
In addition to the previous features, all the models considered also share the same supply block, which is kept as close as possible to that of the standard NK model. As described in detail below, an implication of that block is that work hours are demand-determined, and distributed uniformly across households. Accordingly, $\mathcal{N}_t(j) = N_t$ for all $j \in [0,1]$, an assumption that impose from now onward.
Across the models considered, the household block differs along three dimensions. Firstly, households may or may not be identical ex-ante. Secondly, households may or may not experience (uninsurable) idiosyncratic income shocks. Thirdly, households' borrowing constraints may or may not be occasionally binding. We next discuss these differences in more detail.

1.1.1 HANK models with Ex-Ante Identical Households

A first class of models assume that *all* households are identical ex-ante. In that case, all households have access to the bond market where they can borrow or lend at a (gross) riskless real rate R_t. They also hold an equal number of shares in the equity fund. The period budget constraint for household j is given by

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

where $B_t(j)$ denotes bond holdings and $Y_t(j)$ non-interest income. The latter is given by

$$Y_t(j) = W_t N_t \Xi_t(j) + D_t(j) \tag{2}$$

where W_t is the wage (per efficiency unit of labor), $\Xi_t(j) \equiv \exp\{\zeta_t(j)\}$ is an idiosyncratic shock to the household productivity (i.e. the number of efficiency units supplied per work hour), and $D_t(j)$ are the dividends paid by the equity fund. Shares in the latter are nontradable and, under the assumption of no ex-ante heterogeneity, they are held uniformly across households, implying $D_t(j) = D_t$. We assume $\Xi_t(j)$ is independent across households, and has a time-invariant cross-sectional distribution with mean $\mathbb{E}\{\Xi_t(j)\} = 1$.
Given the absence of a complete set of state-contingent securities, households cannot fully insure against idiosyncratic income shocks. Yet, they can use bonds to provide partial insurance, by saving (borrowing) when they experience positive (negative) shocks and thus helping smooth individual consumption. That mechanism is, however, limited by the existence of a borrowing constraint for each household of the form

$$B_t(j) \geq \underline{B} = -\Psi Y \tag{3}$$

[7] See, e.g., Kaplan et al. (2018) and Auclert et al. (2023). A future extension of our proposed analysis will relax the assumption of fully illiquid stocks, by allowing trade in those assets subject to convex adjustment costs. That extension should help us understand the role played by changes in the degree of (il)liquidity in shaping the equilibrium behavior of HANK economies.

where $\underline{B} \leq 0$ is the borrowing limit and Ψ is the maximum debt ratio (relative to steady state output). In the previous environment, the intertemporal optimality condition associated with the household's problem is given by:

$$Z_t C_t(j)^{-\sigma} \geq \beta R_t \mathbb{E}_t\{Z_{t+1} C_{t+1}(j)^{-\sigma}\} \tag{4}$$

which holds with strict equality if $B_t(j) > \underline{B}$. Under the (baseline) assumption that $\underline{B}$ is above the natural debt limit, a measure of households λ_t^W will be against their borrowing constraint in any given period t. Let $\mathcal{W}_t \subset [0,1]$ denote that subset. Thus for $j \in \mathcal{W}_t$ consumption will be given by

$$C_t(j) = Y_t(j) + \underline{B} - R_{t-1} B_{t-1}(j) \tag{5}$$

implying a unit marginal propensity to consume (MPC) out of income or, in other words, *hand-to-mouth* (HtM) behavior. Note, however, that such behavior is circumstantial, i.e. not intrinsic to any specific households. Furthermore, in the above environment households displaying HtM behavior will simultaneously be holding a large amount of illiquid assets (by assumption, as large as unconstrained households). Accordingly, and following the terminology in Kaplan et al. (2018) we henceforth refer to them as *wealthy hand-to-mouth*.

On the other hand, for the measure $1 - \lambda_t^W$ households satisfying $B_t(j) > \underline{B}$, Euler equation (4) will hold with equality. We refer to those households simply as *unconstrained*, and represent them by the set $\mathcal{U}_t \subset [0,1]$.

We view the general household block described above as incorporating the key elements of the corresponding block in standard HANK models. Accordingly, we henceforth refer to the models that embed it simply as HANK, with no further qualifications.

We highlight two particular cases that are nested in the previous framework. First, when initial holdings of bonds are uniformly distributed ($B_{-1}(j) = 0$ for all $j \in [0,1]$) and there are no idiosyncratic income shocks ($\Xi_t(j) = 1$ for all $j \in [0,1]$ and $t = 0, 1, 2, ..$), the environment described above corresponds to that in the familiar NK model with a representative household (RANK).

Secondly, when $\underline{B}$ equals the natural debt limit, our assumption on utility implies that borrowing constraint (3) never binds and, as a result, (4) always holds with equality for all households. This case, analyzed in Debortoli and Galí (2023), is of special interest since it allows us to insulate the role played by binding borrowing constraints in shaping the economy's aggregate equilibrium dynamics. We label a version of the HANK model with that feature as HANK-NBC (for nonbinding borrowing constraints).

1.1.2 HANK Models with Ex-Ante Household Heterogeneity

In parallel to the models discussed above, we also analyze a class of HANK models with a specific type of ex-ante heterogeneity, namely, the existence of time-invariant set of households $\mathcal{P} \subset [0,1]$, with measure λ^P, that hold no assets or liabilities of any type (liquid or illiquid), and just consume their current labor income. Those households *always* behave in an HtM fashion. While the previous assumption is standard in so-called TANK models, in the class of models analyzed here those households will generally coexist with a measure λ_t^W of wealthy HtM households. To distinguish them from the latter, and following again the terminology in Kaplan et al. (2018), we henceforth refer to households in $\mathcal{P}$ as the *poor hand-to-mouth*.

Letting C_t^P denote average consumption of poor HtM households, and assuming that their idiosyncratic income shocks follow the same distribution as for all other households, we have:[8]

$$\begin{aligned} C_t^P &\equiv \frac{1}{\lambda^P} \int_{i \in \mathcal{P}} C_t(j) dj \\ &= W_t N_t \frac{1}{\lambda^P} \int_{i \in \mathcal{P}} \Xi_t(j) dj \\ &= W_t N_t \end{aligned}$$

We refer to versions of the HANK models that incorporate the previous ex-ante heterogeneity (and, hence, a positive measure of poor HtM households) as HANK+ models. In those models the behavior of wealthy HtM and unconstrained households follows exactly the description in the previous section.

[8] The assumption of a common distribution of idiosyncratic income shocks may not be viewed as realistic, since the average wage of households with no assets or liabilities is likely to be higher than that of households participating in financial markets. We plan to allow for such differences in future extensions of our work.

In HANK+ models one can define the set encompassing *all* HtM households in any given period t as $\mathcal{H}_t = \mathcal{W}_t \cup \mathcal{P}$, with measure given by $\lambda_t^H = \lambda_t^W + \lambda^P$. As with conventional HANK models, two particular cases nested among the set of HANK+ models are worth singling out. The first particular case corresponds to a model where $\underline{B}$ equals the natural debt limit, so that the borrowing constraint is never binding for households participating in the bond market. Accordingly, $\lambda_t^W = 0$ and $\lambda_t^H = \lambda^P$ for all t, i.e. no wealthy HtM households exist. We refer to this variant of the HANK+ model as HANK-NBC+.

A second particular case of interest emerges if we assume away idiosyncratic income shocks, i.e. if we set $\Xi_t(j) = 1$ for all t and $j \in [0,1]$.[9] That case corresponds to the environment in the standard TANK model (see, e.g., Galí et al. (2007) and Bilbiie (2008)), which is characterized by two types of households, (poor) hand-to-mouth and unconstrained, with no heterogeneity (ex-ante or ex-post) within each group.

Table 1 summarizes the different models considered, based on the assumption pertaining the household block.

Table 1. A Kit of HANK Models

	No ex-ante heterogeneity	*Ex-ante heterogeneity*
No idiosyncratic shocks	RANK	TANK
Idiosyncratic shocks + nonbinding BCs	HANK-NBC	HANK-NBC+
Idiosyncratic shocks + binding BCs	HANK	HANK+

1.2 Firms

The supply side of our model is kept as simple as possible. In particular we make assumptions that guarantee that it is not affected by the presence of heterogeneity. This allows us to focus on the impact of the latter on aggregate demand (which coincides with aggregate consumption in our simple model).

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

$$Y_t(i) = A_t N_t(i) \tag{6}$$

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

$$\Pi_t (\Pi_t - 1) = \mathbb{E}_t \left\{ \Lambda_{t,t+1} \left(\frac{Y_{t+1}}{Y_t} \right) \Pi_{t+1} (\Pi_{t+1} - 1) \right\} + \frac{\epsilon}{\xi} \left(\frac{1}{\mathcal{M}_t} - \frac{1}{\mathcal{M}} \right) \tag{7}$$

where $\Pi_t \equiv P_t / P_{t-1}$ is (gross) price inflation rate, $\mathcal{M}_t = A_t / W_t$ is the price markup, with $\mathcal{M} \equiv \epsilon/(\epsilon - 1) > 1$ being its desired (or flexible price) counterpart. Aggregate profits are given by $D_t = Y_t \Delta(\Pi_t) - W_t N_t$ where $\Delta(\Pi_t) \equiv 1 - (\xi/2)(\Pi_t - 1)^2$.

We assume a wage schedule given by

$$W_t = \mathcal{M}_w C_t^\sigma N_t^\varphi \tag{8}$$

where $C_t \equiv \int_0^1 C_t(j) dj$ and $N_t \equiv \int_0^1 N_t(i) di$ denote aggregate consumption and employment, respectively. Note that $C_t^\sigma N_t^\varphi$ can be interpreted as an average marginal rate of substitution, with $\mathcal{M}_w > 1$ thus playing the role of an average wage markup, determined by workers' market power.[10] Given the wage, firms determine the quantity of labor hired, which is distributed uniformly across all households.[11]

[9] Equivalently, we could assume the presence of a complete set of state-contingent securities available to unconstrained households, allowing the latter to fully insure against their idiosyncratic income shocks.

[10] Strictly speaking, this interpretation is valid only if $\sigma = 1$.

[11] Under our assumptions, if $\mathcal{M}_w \gg 1$, all households will be willing to supply the amount of labor demanded by firms, as long as shocks are not too large.

Combining equations (8)-(7), and taking a first-order approximation around the zero-inflation steady state yields the inflation equation

$$\pi_t = \beta \mathbb{E}_t\{\pi_{t+1}\} - \lambda \widehat{\mu}_t \tag{9}$$

where $\widehat{\mu}_t \equiv \log(\mathcal{M}_t/\mathcal{M})$ and $\lambda \equiv (\epsilon - 1)/\xi$. Noting that $\mathcal{M}_t = A_t/W_t$ and using the wage schedule above we obtain

$$\widehat{\mu}_t = -(\sigma + \varphi)\widehat{y}_t + (1 + \varphi)a_t \tag{10}$$

where $\widehat{y}_t \equiv \log(Y_t/Y)$. Combining (9) and (10) yields a version of the familiar New Keynesian Phillips curve

$$\pi_t = \beta \mathbb{E}_t\{\pi_{t+1}\} + \kappa \widehat{y}_t - \delta a_t \tag{11}$$

where $\kappa \equiv \lambda(\sigma + \varphi)$ and $\delta \equiv \lambda(1 + \varphi)$.[12] Note that equation (11) is invariant to the presence of household heterogeneity and the nature of the latter. In particular, it holds across all HANK models discussed above.

1.3 Monetary Policy

Regarding monetary policy, we assume that the central bank controls the real interest rate $r_t \equiv \log R_t$. In our baseline specification, we assume that the central bank keeps the real interest rate constant in the face of aggregate shocks to preferences and technology. That approach allows us to determine the response of aggregate variable to different shocks in a way that is not mediated by any particular assumption regarding the monetary policy rule, which would generally lead to different paths of the real rate across environments associated with different responses of aggregate variables.[13] When considering monetary policy shocks, we assume instead that the real rate follows an exogenous $AR(1)$ process $\widehat{r}_t = \rho_r \widehat{r}_{t-1} + \varepsilon_{m,t}$. However, as a robustness exercise, we also analyze a variant of our model where the central bank follows a Taylor-type rule for the nominal interest rate, to confirm that our main qualitative findings remain unaltered.

2 Inspecting the Mechanism

A central objective of our research is to understanding the channels through which different dimensions of heterogeneity found in HANK models end up affecting aggregate consumption and, as a result, aggregate output and inflation.[14] In order to help us in that task, next we derive an (approximate) aggregate consumption function which holds in the most general model in our kit of HANK models: the HANK+. The remaining models can be viewed as being nested in that model, and can be analyzed as particular cases of the latter, by switching off different channels.
Consider to begin with the following decomposition of aggregate consumption:

$$C_t = \lambda^P C_t^P + \lambda_t^W C_t^W + (1 - \lambda_t^H) C_t^U \tag{12}$$

where $C_t^P \equiv \frac{1}{\lambda^P}\int_{i\in\mathcal{P}} C_t(j)dj$, $C_t^W \equiv \frac{1}{\lambda_t^W}\int_{j\in\mathcal{W}_t} C_t(j)dj$ and $C_t^U \equiv \frac{1}{\lambda_t^U}\int_{j\in\mathcal{U}_t} C_t(j)dj$ denote the average consumption of poor HtM, wealthy HtM and unconstrained households, respectively, while $\lambda_t^H \equiv \lambda^P + \lambda_t^W$ denotes the combined measure of HtM households (poor and wealthy) in period t. A first order approximation to the the previous decomposition around the stochastic steady state yields

$$\widehat{c}_t = \lambda^P \widehat{c}_t^P + \lambda^W \widehat{c}_t^W + (1 - \lambda^H)\widehat{c}_t^U - \Gamma \widehat{\lambda}_t \tag{13}$$

where $\widehat{c}_t \equiv \log(C_t/C)$, $\widehat{c}_t^i \equiv (C_t^i - C^i)/C$ for $i \in \{P, W, U\}$, $\widehat{\lambda}_t \equiv \lambda_t - \lambda$, and $\Gamma \equiv (C^U - C^W)/C$, with variables without time subscripts denoting the corresponding mean values in the (stochastic) steady

[12] Alternatively, we could have written the NKPC in terms of the output gap

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

where $\widetilde{y}_t \equiv y_t - y_t^n$ denotes the gap between (log) output y_t and its natural (i.e. flexible price) counterpart $y_t^n \equiv \frac{1+\varphi}{\sigma+\varphi} a_t - \frac{\mu}{\sigma+\varphi}$

[13] A similar approach is followed by Woodford (2011) when studying the size of the fiscal multiplier in a New Keynesian model.

[14] The spirit of this exercise is similar to Werning (2015), though the specifics of our approach are different.

state (i.e. $X \equiv \mathbb{E}\{X_t\}$). Note that Γ is a measure of steady state consumption inequality between unconstrained and wealthy HtM households. The term $\Gamma\widehat{\lambda}_t$ thus captures the (first-order) impact of changes in the fraction of constrained households on aggregate consumption, resulting from the difference in average consumption between the two types. Next we analyze each of the components of (12) and (13) in some detail.

Average consumption of poor HtM households is given by

$$C_t^P = W_t N_t = \frac{Y_t}{\mathcal{M}_t} \tag{14}$$

Equivalently, and up to a first-order approximation

$$\begin{aligned} \widehat{c}_t^P &= \frac{1}{\mathcal{M}}(\widehat{y}_t - \widehat{\mu}_t) \\ &= \frac{1}{\mathcal{M}}(1+(\sigma+\varphi))\widehat{y}_t - \frac{1}{\mathcal{M}}(1+\varphi)a_t \end{aligned} \tag{15}$$

where the second equality follows from (10). The (individual) marginal propensity to consume out of *aggregate* output for poor HtM households, conditional on a given average markup, is given by $\vartheta^P \equiv \frac{\partial \widehat{c}_t^P}{\partial \widehat{y}_t} = \frac{1}{\mathcal{M}} < 1$ The countercyclicality of the markup in response to non-technology shocks, combined with the fact that labor income is the only source of income for poor HtM households, amplifies the overall response of those households' consumption to changes in aggregate output. As discussed below, this will be a source of a strengthened output multiplier for demand shocks.

Average consumption for wealthy HtM households can be obtained by integrating (5) over all $j \in \mathcal{W}_t$, which yields

$$C_t^W = Y_t^W + \underline{B} - R_{t-1}B_{t-1|t}^W \tag{16}$$

where $Y_t^W \equiv \frac{1}{\lambda_t^W}\int_{j\in\mathcal{W}_t} Y_t(j)dj$ is average non-interest income and $B_{t-1|t}^W \equiv \frac{1}{\lambda_t^W}\int_{j\in\mathcal{W}_t} B_{t-1}(j)dj$ denotes (lagged) average net bond holdings, both referring to wealthy HtM households in period t. Note that we can write

$$\begin{aligned} Y_t^W &\equiv \frac{1}{\lambda_t^W}\int_{j\in\mathcal{W}_t} Y_t(j)dj \\ &= W_t N_t \Theta_t + \frac{1}{1-\lambda^P} D_t \end{aligned} \tag{17}$$

where $\Theta_t \equiv \frac{1}{\lambda_t^W}\int_{j\in\mathcal{W}_t} \Xi_t(j)dj$ is the average productivity of currently wealthy HtM households. Under standard assumptions, we would expect $\Theta_t < 1$ due to a selection effect, to the extent that households with lower productivity realizations are the ones who are morel likely to hit the borrowing limit. Thus, combining (16) and (17), together with the fact that $W_t N_t = \frac{Y_y}{\mathcal{M}_t}$ and $D_t = \left(\Delta(\Pi_t) - \frac{1}{\mathcal{M}_t}\right) Y_t$, we can write

$$C_t^W = \left[\frac{\Theta_t}{\mathcal{M}_t} + \frac{1}{1-\lambda^P}\left(\Delta(\Pi_t) - \frac{1}{\mathcal{M}_t}\right)\right] Y_t + \underline{B} - R_{t-1}B_{t-1|t}^W \tag{18}$$

Thus, we see that while consumption of the wealthy HtM is, ceteris paribus, increasing in aggregate output, its relation with the price markup is, in principle, ambiguous, since an increase in the latter lowers wages but raises dividends, both of which are a source of income for wealthy HtM households. A first order approximation of the previous expression around the (zero inflation) steady state yields:

$$\widehat{c}_t^W = \vartheta^W \widehat{y}_t + \frac{1-(1-\lambda^P)\Theta}{\mathcal{M}(1-\lambda^P)}\widehat{\mu}_t + \frac{1}{\mathcal{M}}\widehat{\Theta}_t - R(\widehat{b}_{t-1|t}^K + b_-^K \widehat{r}_{t-1}) \tag{19}$$

where $\vartheta^W \equiv \frac{\partial \widehat{c}_t^W}{\partial \widehat{y}_t} = \frac{\Theta}{\mathcal{M}} + \frac{1}{1-\lambda^P}\left(1-\frac{1}{\mathcal{M}}\right)$, $\widehat{b}_{t-1|t}^K \equiv \frac{B_{t-1|t}^K - B_-^K}{Y}$, $B_-^K \equiv \mathbb{E}\{B_{t-1|t}^K\}$ and $b_{--}^K \equiv \frac{B_-^K}{Y}$. The first term on the right hand side captures the fact that an increase in aggregate output increases the income of constrained households, ceteris paribus. The (individual) MPC out of aggregate output for wealthy HtM households, ϑ^W, will generally differ from that of the poor HtM, with the sign of the difference being in principle ambiguous. The reason for this is that wealthy HtM households have two sources of non-interest income, both of which respond to changes in aggregate output. Their labor

income increases by $\frac{\Theta}{\mathcal{M}}$ per each unit increase in aggregate output, which is less than $\frac{\Theta}{\mathcal{M}}$, the marginal increase in average labor income, to the extent that $\Theta < 1$ due to the selection effect discussed above. This is offset (at least partly, and possibly more than fully) by the fact that their dividend income also increases in response to a unit increase in aggregate output, by an amount $\frac{1}{1-\lambda^P}\left(1-\frac{1}{\mathcal{M}}\right)$ The second term in (18) confirms that a change in the price markup has an ambiguous effect on wealthy HtM consumption. Yet, to the extent that $(1-\lambda^P)\Theta < 1$, a plausible condition associated with a composition of income sufficiently tilted towards dividends, an increase in the markup will raise the consumption of wealthy HtM households. The third term in (19) captures the role of changes in the average productivity of wealthy HtM households. Finally, the fourth term in (19) reflects the negative impact of increases in the debt service ratio on wealthy HtM households' cash-on-hand. The debt service ratio depends in turn on their debt ratio in the previous period, $\widehat{b}^K_{t-1|t}$, as well as on the (lagged) real interest rate, $\widehat{r}_{t-1}$. The latter effect thus captures a *direct* channel, independent of intertemporal substitution, through which monetary policy affects consumption, i.e. through changes in the interest payments of indebted HtM households.

Finally, and as shown in Debortoli and Galí (2023), average consumption for unconstrained households satisfies the augmented Euler equation

$$\widehat{c}^U_t = \mathbb{E}_t\{\widehat{c}^U_{t+1}\} - \gamma^U\left(\frac{1}{\sigma}\widehat{r}_t + \frac{1}{\sigma}\mathbb{E}_t\{\Delta z_{t+1}\} + \frac{(\sigma+1)}{2}\widehat{v}_t\right) \tag{20}$$

where $\gamma^U \equiv \frac{C^U}{C} = 1 - \frac{\lambda^P}{\mathcal{M}} - \lambda^W\left[\frac{\Theta}{\mathcal{M}} + \frac{1}{1-\lambda^H}\left(1-\frac{1}{\mathcal{M}}\right) + \underline{b} - Rb^W_{--}\right]$, and $\widehat{v}_t = v_t - v$, with

$$v_t \equiv \frac{1}{\lambda^U}\int_{j\in\mathcal{U}_t}\frac{C_t(j)}{C_t}v_t(j)dj \tag{21}$$

and where $v_t(j) \equiv var_t\{c_{t+1}(j)\}$ is a measure of risk regarding household j's one period ahead (log) consumption.[15] Note that v_t is a *consumption-weighted average of individual consumption risk*, which captures the impact of precautionary savings on unconstrained households's consumption. Henceforth we refer to v_t as the *risk shifter*. In the absence of idiosyncratic income risk, v_t is of second order relative to aggregate variables and can thus be ignored in a first-order approximation.[16]

Finally, we can substitute (15) and (19) into (13) to yield:

$$\widehat{c}_t = \vartheta\widehat{y}_t + \chi\widehat{\mu}_t + \frac{\lambda^W}{\mathcal{M}}\widehat{\Theta}_t - \lambda^W R(\widehat{b}^W_{t-1|t} + b^W_-\widehat{r}_t) - \Gamma\widehat{\lambda}_t + (1-\lambda^H)\widehat{c}^U_t \tag{22}$$

where

$$\begin{aligned}\vartheta &\equiv \lambda^P\vartheta^P + \lambda^W\vartheta^W \\ &= \frac{\lambda^P}{\mathcal{M}} + \lambda^W\left[\frac{\Theta}{\mathcal{M}} + \frac{1}{1-\lambda^P}\left(1-\frac{1}{\mathcal{M}}\right)\right]\end{aligned}$$

$$\chi \equiv \frac{\lambda^W[1-(1-\lambda^P)\Theta]}{\mathcal{M}(1-\lambda^P)} - \frac{\lambda^P}{\mathcal{M}}$$

where $\widehat{c}^U_t$ satisfies Euler equation (20). Coefficient ϑ can be interpreted as the aggregate MPC of HtM households (out of aggregate output). It is unambiguously positive and assumed to be less than one. Coefficient χ captures the effect of changes in price markups on consumption. Its sign is generally ambiguous, depending on the relative measures of wealthy vs. poor HtM households.

[15] Note that:

$$\begin{aligned}\gamma^U &\equiv \frac{C^U}{C} \\ &= 1 - \frac{\lambda^P C^P}{C} - \frac{\lambda^W C^W}{C} \\ &= 1 - \frac{\lambda^P}{\mathcal{M}} - \lambda^W\left[\frac{\Theta}{\mathcal{M}} + \frac{1}{1-\lambda^H}\left(1-\frac{1}{\mathcal{M}}\right) + \underline{b} - Rb^K_{--}\right]\end{aligned}$$

[16] See Debortoli and Galí (2023) for a discussion.

The different terms on the right hand side of (22) are associated with the different channels through which aggregate shocks affect aggregate consumption in our most general HANK model (i.e., HANK+). Each of the models listed in Table 1 are associated with an aggregate consumption function that is nested in the above specification.

Thus, for instance, in the (baseline) HANK model we have $\lambda^P = 0$ and $\lambda^H = \lambda^W$. All the terms on the right hand side of (22) are preserved, though with suitable redefined coefficients. In particular, $\vartheta \equiv \lambda^W \left[\frac{\Theta}{\mathcal{M}} + \left(1 - \frac{1}{\mathcal{M}}\right)\right]$, $\chi \equiv \frac{\lambda^W(1-\Theta)}{\mathcal{M}}$ and $\gamma^U = 1 - \lambda^W \left[\frac{\Theta}{\mathcal{M}} + \frac{1}{1-\lambda^W}\left(1 - \frac{1}{\mathcal{M}}\right) + \underline{b} - Rb^W_{--}\right]$ in that case.

The HANK-NBC model studied by Debortoli and Galí (2023) implies $\lambda^P = \lambda^W = \lambda^H = \widehat{\lambda}_t = \widehat{\Theta}_t = 0$, and $\gamma^U = 1$. Accordingly, $\widehat{c}_t = \widehat{c}^U_t$, with aggregate consumption thus satisfying the augmented Euler equation

$$\widehat{c}_t = \mathbb{E}_t\{\widehat{c}_{t+1}\} - \frac{1}{\sigma}\widehat{r}_t - \frac{1}{\sigma}\mathbb{E}_t\{\Delta z_{t+1}\} - \frac{\sigma+1}{2}\widehat{v}_t$$

which differs from the Euler equation for aggregate consumption in the RANK model only through the presence of the term $\frac{\sigma+1}{2}\widehat{v}_t$, which is absent in the latter.

On the other hand, in the baseline TANK model we have $\lambda^W = \widehat{\lambda}_t = \widehat{\Theta}_t = \widehat{v}_t = 0$ and $\lambda^H = \lambda^P$. Accordingly, $\vartheta = \frac{\lambda^P}{\mathcal{M}}$, $\chi = -\frac{\lambda^P}{\mathcal{M}}$ and $\gamma^U = 1 - \frac{\lambda^P}{\mathcal{M}}$, so the aggregate consumption equation simplifies to:

$$\widehat{c}_t = \frac{\lambda^P}{\mathcal{M}}\widehat{y}_t - \frac{\lambda^P}{\mathcal{M}}\widehat{\mu}_t + \left(1 - \frac{\lambda^P}{\mathcal{M}}\right)\widehat{c}^U_t$$

with

$$\widehat{c}^U_t = \mathbb{E}_t\{\widehat{c}^U_{t+1}\} - \left(1 - \frac{\lambda^P}{\mathcal{M}}\right)\left[\frac{1}{\sigma}\widehat{r}_t + \frac{1}{\sigma}\mathbb{E}_t\{\Delta z_{t+1}\}\right]$$

An advantage of the TANK model, relative to existing HANK models, is that its equilibrium can be solved analytically, which facilitates the understanding of the forces shaping the economy's response to different shocks or the analysis of optimal monetary policy, among other considerations. By contrast, in existing HANK models, the presence of idiosyncratic shocks, combined with the absence of a complete set of state-contingent securities implies a non-degenerate distribution of consumption and bond holdings at each point in time. The distribution of bond holdings is an (infinite-dimensional) endogenous state variable of the model, which prevents solving for the latter's equilibrium by simple analytical methods, requiring instead the use of numerical methods. That requirement makes it harder to understand its workings and to study their normative implications.

Thus, the key question at hand here is how much are we bound to lose by using a simple TANK (or even RANK) model instead of their richer HANK counterparts. We plan to address that question by comparing the properties of the different models using the framework developed above and more specifically, by evaluating the importance of each of the terms in the aggregate consumption function (22), some of which are missing in the simpler models. This is ultimately a quantitative issue and the focus of our analysis.

3 An Illustrative Simulation (Preliminary)

To provide an (incomplete) illustration of the intended exercise, we consider a baseline calibration of the model economies described above. In particular, each period is assumed to correspond to a quarter. We set the coefficient of risk aversion $\sigma = 1$, which corresponds to log utility. We set the discount factor $\beta = 0.9937$, which implies a real risk-free rate of 2 percent (in annual terms) in the steady state.

Following Auclert et. al. (2021), we calibrate the process for the idiosyncratic income shocks using the Rouwenhorst method in order to match the volatility and persistence of an *AR(1)* process $\zeta_t(j) = \rho_\zeta \zeta_{t-1}(j) + \varepsilon^\zeta_t(j)$, where $\zeta_t(j) \sim N(0, \sigma_\zeta\sqrt{1-\rho^2_\zeta})$, with $\rho_\zeta = 0.966$ and $\sigma_\zeta = 0.5$. In addition, we set the (inverse) Frisch elasticity of substitution to unity, $\varphi = 1$. Also, we set the elasticity of substitution among good varieties $\epsilon = 11$, which implies an average price markup of about 10 percent and the price adjustment cost parameter ξ so that the resulting slope of the Phillips Curve is $\kappa = 0.1$. Finally, we set the borrowing limit $\underline{B} = -0.2$ so that the fraction of hand-to-mouth households in steady state is $\lambda = 0.3$.

Figure 1 shows the response of aggregate consumption to an expansionary monetary policy shock, in the form of a 25 basis point decrease on impact in the real interest rate (1 p.p. in annualized terms), followed by a gradual reversion to its steady state value according to an $AR(1)$ process with autoregressive coefficient $\rho_r = 0.5$.The figure reports the responses of aggregate consumption (which is equal to output) for four economies: RANK, TANK, HANK-NBC and HANK.
Consistently with previous results in the literature, we find that the response of consumption in the HANK economy (red line) is much larger (almost three times) than in the RANK economy (black line), thus suggesting that the presence of idiosyncratic risk and a binding borrowing constraint leads to a substantial amplification of the effects of monetary policy shocks. However, most of the amplification effects are well approximated by a simple TANK eeconomy (blue line), which features an equal share of hand-to-mouth households as the HANK economy does in the steady state, but completely abstracts from the presence of idiosyncratic risk. The reason behind this result can be better understood looking at the response of the HANK-NBC economy (red line with circles), which is almost identical to its counterpart in the RANK economy. This suggests that the presence of idiosyncratic risk, play a limited role for the transmission of monetary shocks, when the borrowing constraint is not binding, a central finding in Debortoli and Galí (2023).. In other words, our analysis suggests that the key propagation channel found in the HANK economy is related to the presence of hand-to-mouth households, a channel that is captured reasonably well by a simple TANK economy, with no idiosyncratic risk. The remaining channels uncovered in our analysis above seem to have a negligible quantitative role, at least when aggregated.

4 Steps Ahead

There are three key steps that we plan to undertake based on the framework spelled out above:

- Develop variants of the HANK model where the degree of (il)liquidity can be parameterized.
- Analysis and comparison of the impulse responses of aggregate variables to different shocks (monetary policy, technology, discount factor) for calibrated versions of the different models considered
- Quantitative analysis of the different components of the aggregate consumption function, corresponding to different mechanisms at work in the different models, conditional on each type of shock
- Attempt to give a (possibly tentative) answer to the question that originally motivated our research, namely, under what assumptions does a properly calibrated TANK model approximate well the equilibrium dynamics of a HANK model?

References.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie and Ludwig Straub (2021): "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica*, 89(5), 2375-2408.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2023): "The Intertemporal Keynesian Cross," unpublished manuscript.

Bilbiie, Florin (2008): "Limited Asset Market Participation, Monetary Policy and (Inverted) Aggregate Demand Logic," *Journal of Economic Theory*, 140(1), 162-196.

Campbell, John Y. and N. Gregory Mankiw (1989): "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence,", *NBER Macroeconomics Annual,* 4, 185-216.

Galí, Jordi (2015): *Monetary Policy, Inflation and the Business Cycle*, Second edition, Princeton University Press.

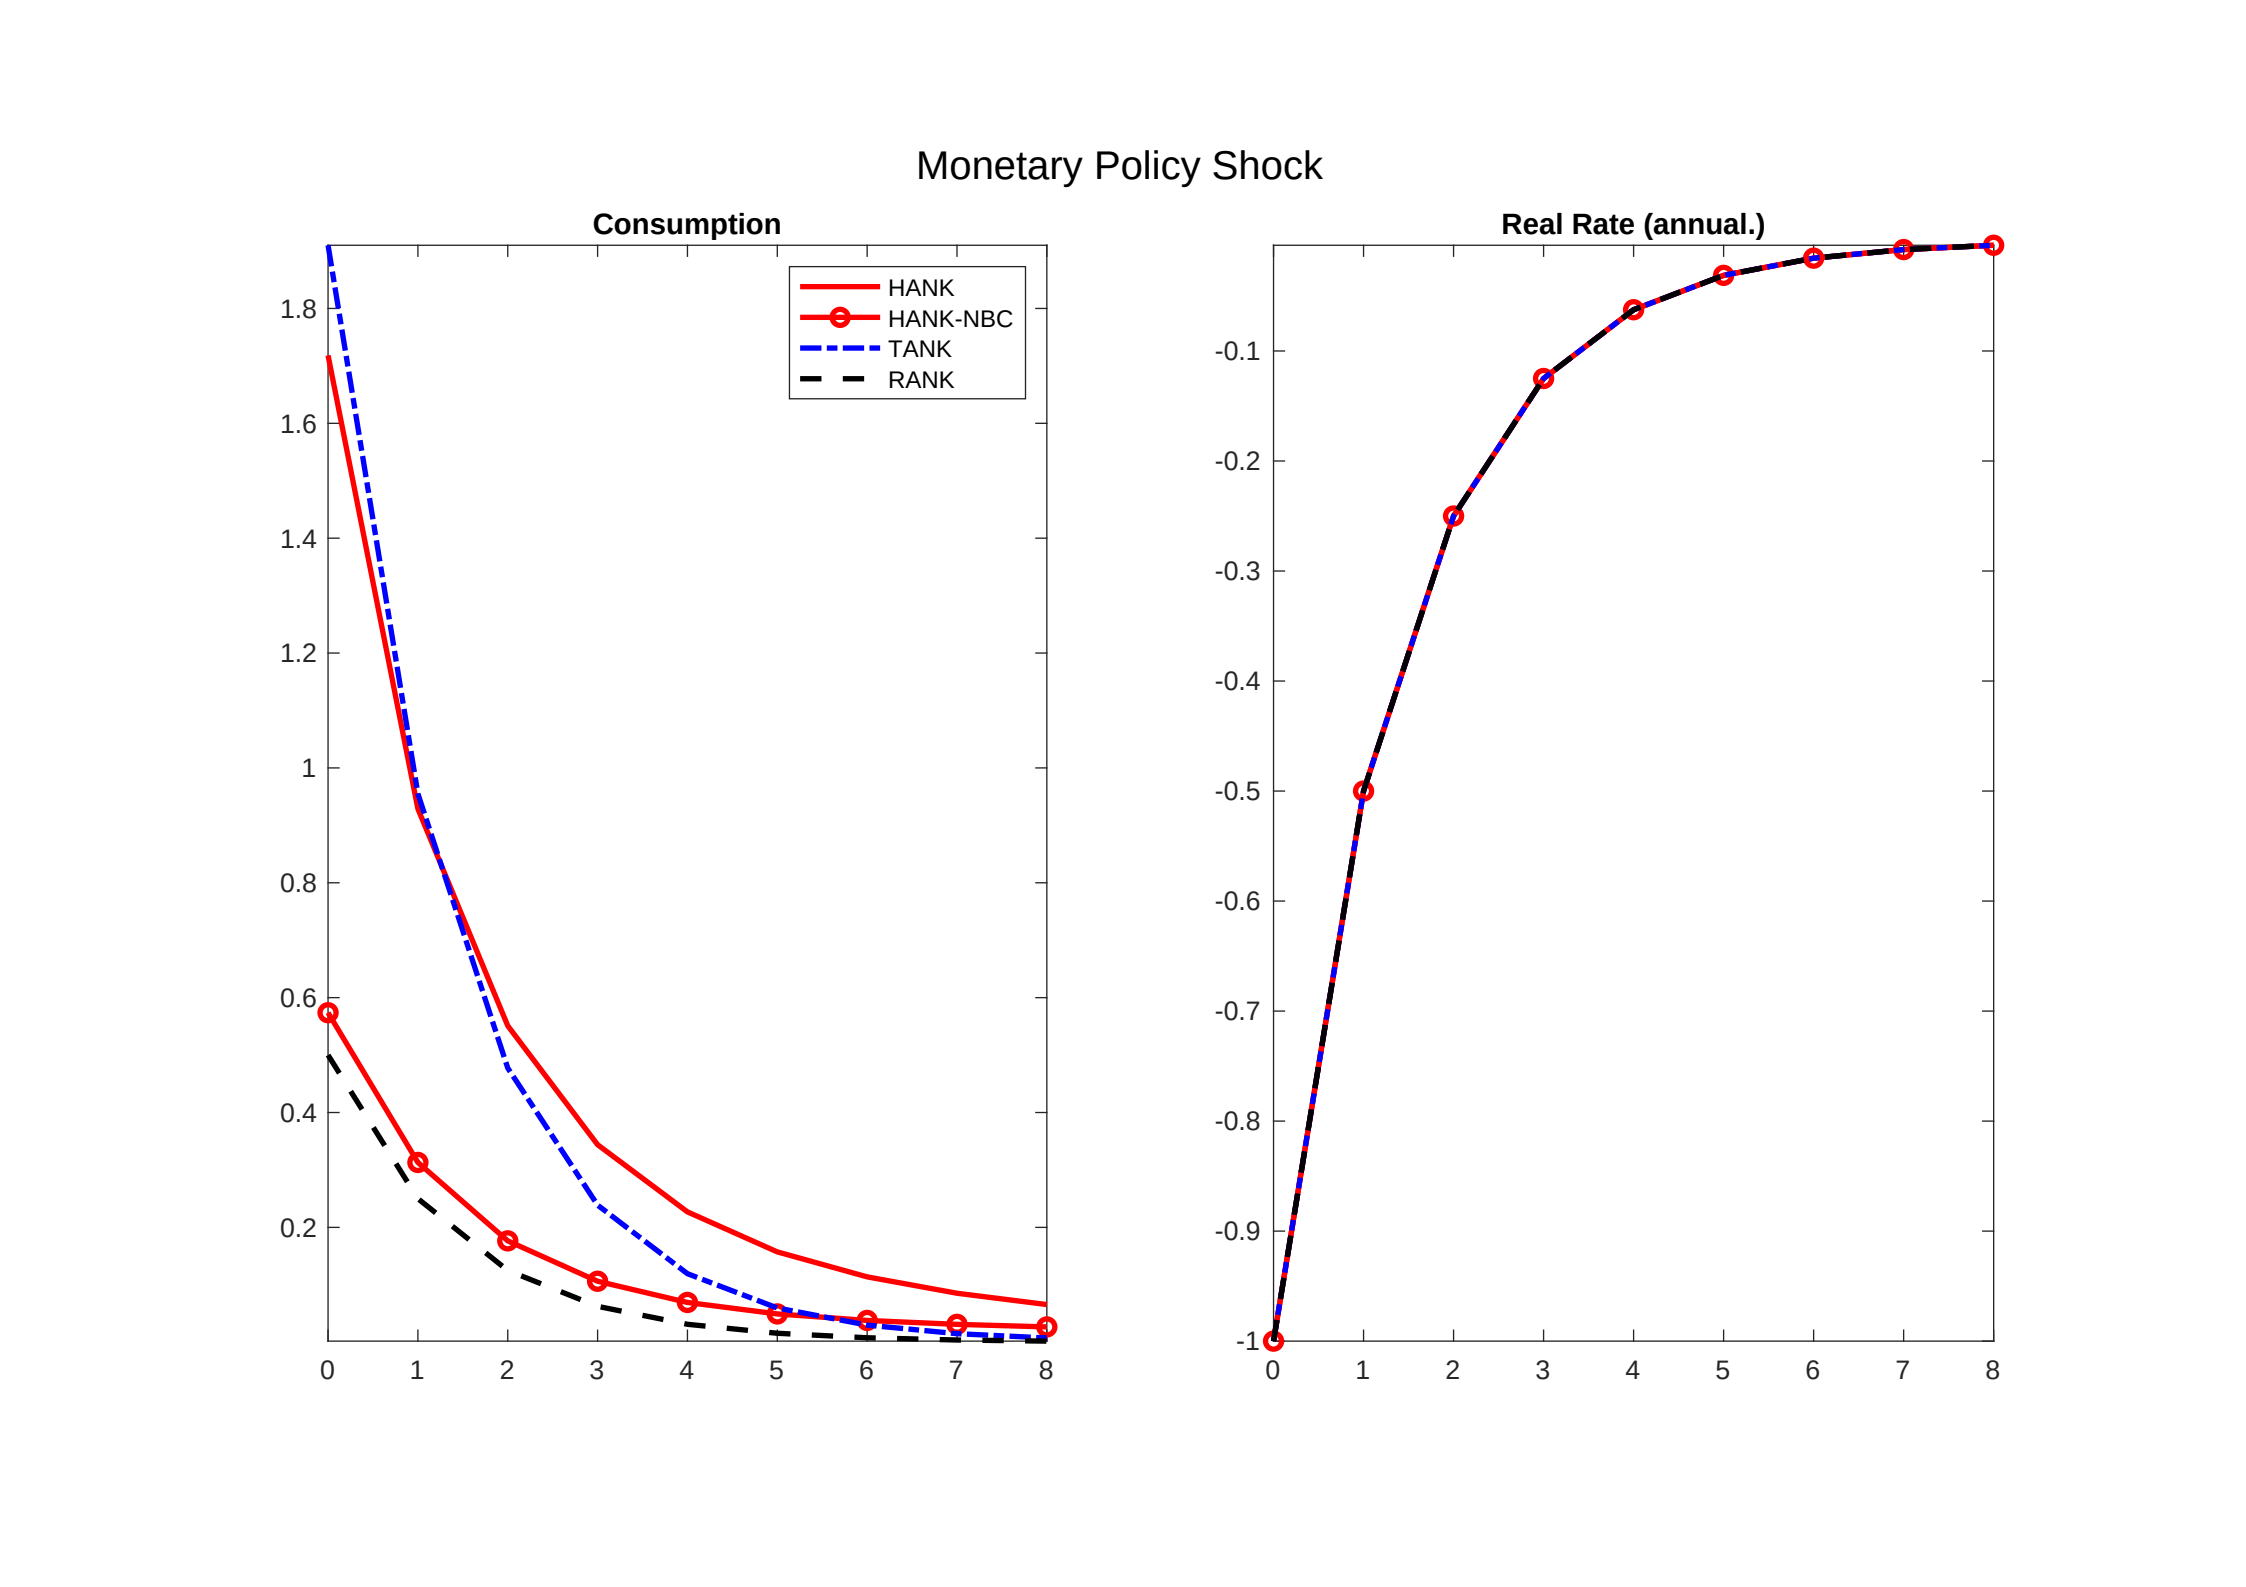
Galí, Jordi (2018): "The State of New Keynesian Economics: A Partial Assessment," *Journal of Economic Perspectives* 32(3), 87-112.

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

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

Werning, Iván (2015): "Incomplete Markets and Aggregate Demand," NBER Working Paper No. 21448.

Woodford, Michael (2011): "Simple analytics of the government expenditure multiplier," *American Economic Journal: Macroeconomics*, 3(1), 1-35.


Monetary Policy Shock
Consumption
HANK
HANK-NBC
TANK
RANK
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
1
2
3
4
5
6
7
8
Real Rate (annual.)
-0.1
-0.2
-0.3
-0.4
-0.5
-0.6
-0.7
-0.8
-0.9
-1
0
1
2
3
4
5
6
7
8