

NBER WORKING PAPER SERIES

INEQUALITY AND THE ZERO LOWER BOUND

Jesús Fernández-Villaverde
Joël Marbet
Galo Nuño
Omar Rachedi

Working Paper 31282
<http://www.nber.org/papers/w31282>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2023

Jesús gratefully acknowledges funding from the Fundación Ramón Areces. We thank Guido Ascari, Olivier Blanchard, Anmol Bhandari, Edouard Challe, Jordi Galí, Julien Matheron, Taisuke Nakata, Anton Nakov, Sebastian Schmidt, Gianluca Violante, Pablo Winant, and participants at numerous seminars and conferences for useful comments and suggestions. The views expressed in this paper are those of the authors and do not necessarily represent the views of the Banco de España, the Eurosystem, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Jesús Fernández-Villaverde, Joël Marbet, Galo Nuño, and Omar Rachedi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Inequality and the Zero Lower Bound

Jesús Fernández-Villaverde, Joël Marbet, Galo Nuño, and Omar Rachedi

NBER Working Paper No. 31282

May 2023

JEL No. D31,E12,E21,E31,E43,E52,E58

ABSTRACT

This paper studies how household inequality shapes the effects of the zero lower bound (ZLB) on nominal interest rates on aggregate dynamics. To do so, we consider a heterogeneous agent New Keynesian (HANK) model with an occasionally binding ZLB and solve for its fully non-linear stochastic equilibrium using a novel neural network algorithm. In this setting, changes in the monetary policy stance influence households' precautionary savings by altering the frequency of ZLB events. As a result, the model features monetary policy non-neutrality in the long run. The degree of long-run non-neutrality, i.e., by how much monetary policy shifts real rates in the ergodic distribution of the model, can be substantial when we combine low inflation targets and high levels of wealth inequality.

Jesús Fernández-Villaverde
Department of Economics
University of Pennsylvania
The Ronald O. Perelman Center
for Political Science and Economics
133 South 36th Street Suite 150
Philadelphia, PA 19104
and CEPR
and also NBER
jesusfv@econ.upenn.edu

Joël Marbet
CEMFI
Calle Casado del Alisal, 5
Madrid 28005
Spanien
joel.marbet@cemfi.edu.es

Galo Nuño
Banco de España
Alcalá 48,
28014 Madrid
Spain
galo.nuno@bde.es

Omar Rachedi
Esade Business School
Avinguda de la Torre Blanca, 59
Sant Cugat, Barcelona
Spain
omar.rachedi@yahoo.com

1 Introduction

This paper examines the interplay between household inequality and the zero lower bound (ZLB) of nominal interest rates and their impact on the effectiveness of monetary policy. To accomplish this, we integrate two distinct lines of research. First, we delve into the analysis of heterogeneous agent New Keynesian (HANK) models, which have emerged as a popular framework for exploring the relationship between the transmission mechanism of monetary policy and income and wealth inequality (McKay et al., 2016, Kaplan et al., 2018, and Auclert, 2019). However, HANK models tend to either abstract from the ZLB or incorporate it only with perfect foresight (McKay et al., 2016 and Guerrieri and Lorenzoni, 2017). Second, we draw upon a body of literature that employs representative-agent New Keynesian (RANK) models to emphasize the adverse consequences of the ZLB (Christiano et al., 2011, Coibion et al., 2012, and Andrade et al., 2019). These studies underscore that the ZLB gives rise to deflationary spirals, severely constraining central banks' ability to effectively accommodate negative shocks to the economy.

More concretely, we build and globally solve a HANK model with a ZLB to demonstrate how household inequality exacerbates the challenges arising from the ZLB. The primary channel driving our results is that household inequality increases the demand for precautionary savings, which lowers the real interest rate. Since the nominal interest rate is just the real interest rate plus inflation (which, as we will argue below, is also lower due to the ZLB), the average nominal rates are lower, providing less scope for the central bank to stabilize the economy. Furthermore, any factor that changes idiosyncratic risk, such as higher job volatility, will affect the real rates and shift the effectiveness of monetary policy.

Our ZLB-HANK economy features a continuum of ex-ante homogeneous households, which are ex-post heterogeneous due to uninsurable idiosyncratic labor earning risk in the spirit of Bewley (1980), Huggett (1993), and Aiyagari (1994). Households can smooth the effect of labor earning risk on their consumption via borrowing. However, their credit capacity is limited by a borrowing constraint. We explicitly account for the risk of hitting the ZLB in households' expectations, as in Fernández-Villaverde et al. (2015a). On the production side, there is a continuum of intermediate-good firms that produce using labor and set prices subject to adjustment costs à la Rotemberg (1982). The model also features a central bank that sets the nominal interest rate according to a Taylor rule subject to the ZLB constraint and a fiscal authority that taxes labor income. Finally, we consider demand shocks as the source of aggregate uncertainty.

We solve our model non-linearly by introducing a novel neural-network algorithm that allows us to compute the stochastic equilibrium dynamics of a ZLB-HANK economy. Rather

than using standard computation methods for HANK models based on linearization (e.g., [Ahn et al., 2018](#), and [Auclert et al., 2021](#)), we build on the global approach of [Fernández-Villaverde et al. \(2023\)](#). We demonstrate how neural networks can successfully approximate the non-linear laws of motion of the economy’s states generated by the ZLB. We calibrate the model to replicate labor earnings and wealth dispersion, average marginal propensity to consume, and the frequency of ZLB occurrences observed in the U.S. economy after 1945.

In terms of substantive findings, our first result is that the ZLB moves the ergodic distribution of inflation, the nominal and real interest rates, and aggregate consumption downward; that is, the ZLB means that we will have more and more severe episodes of deflation and lower consumption. This result is unsurprising and also appears in RANK models with a ZLB. The intuition is simple: the central bank cannot fully accommodate large negative demand shocks by lowering nominal rates below zero. In contrast, when the negative demand shocks are small, the central bank still has space to accommodate, even with the ZLB.

Our second result is that the negative effects of the ZLB fall disproportionately on wealth-poor households. Since these households depend heavily on wages, a deep recession exacerbated by the ZLB reduces their income relatively more than that of wealth-rich households, which have interest income from their bond holdings. Also, since wealth-poor households have few or negative assets, they cannot rely on their bond holdings to smooth consumption. Therefore, the consumption of wealth-poor households drops relatively more than that of wealth-rich households.

Our third result is that household heterogeneity limits the central bank’s space to lower nominal rates, making spells at the ZLB more likely. As in [Aiyagari \(1994\)](#), household heterogeneity leads to more precautionary savings even without aggregate shocks and, thus, lower real interest rates, which also results in lower nominal rates.

Our fourth result is that the ZLB makes the stochastic steady state (SSS) of the HANK model separate from the deterministic steady state (DSS).¹ The key mechanism is that agents understand the presence of a “deflationary bias” in the economy: the central bank can accommodate large positive demand shocks, as the nominal interest rate can go up as much as needed, but not large negative demand shocks, as the ZLB becomes binding. Thus, the economy will suffer more periods of deflation than in the absence of the ZLB, and the central bank will undershoot its inflation target. Agents respond to this deflationary bias by

¹The DSS is a point where all aggregate variables are constant and there are no aggregate shocks, but agents still face idiosyncratic risk. In contrast, the SSS (or risky steady state, e.g., [Coeurdacier et al., 2011](#)) is a point at which all aggregate variables are constant and the realization of the aggregate demand shock is zero, although agents are aware that non-zero realizations can come in the future. Also, agents still face idiosyncratic risk. The SSS is an important concept because it often provides a better summary of the ergodic distribution of non-linear models such as ours than the DSS.

increasing their precautionary savings, which lowers real rates: higher demand for savings given the same supply of bonds in our economy translates into higher prices for bonds, i.e., lower rates. While this deflationary bias is also present in ZLB-RANK models (e.g., [Adam and Billi, 2007](#), [Nakov, 2008](#), [Hills et al., 2019](#), and [Bianchi et al., 2021](#)), we show that heterogeneity makes it more acute.

Since our third and fourth results highlight the importance of household heterogeneity and the ZLB, we demonstrate that an increase in household heterogeneity (e.g., more idiosyncratic labor risk) or factors that make the ZLB more likely (e.g., a lower inflation target) will increase precautionary savings in the economy.

This observation has a sharp implication: monetary policy is not neutral in the long run in our ZLB-HANK economy, and the Fisher equation depends on the central bank’s stance. More specifically, if we denote the steady-state nominal interest rate, real rate, inflation rate, and inflation target by i , r , π , and $\tilde{\pi}$, respectively, then $i(\tilde{\pi}) = r(\tilde{\pi}) + \pi(\tilde{\pi})$, such that $dr/d\tilde{\pi} > 0$.

In our baseline calibration, household heterogeneity has a small impact on the long-run Fisher equation. To illustrate, when we adjust the inflation target from 4% to 1.7%, representing the average inflation rates between 1980 and 1999 and from 2000 onward, respectively, the decline in the real interest rate in the SSS of our model is 16 basis points (bps). In comparison, the real rate falls 12 bps in the ZLB-RANK. In other words, heterogeneity only adds 4 additional bps of reduction in the real rate. More generally, this finding for the ZLB-HANK and ZLB-RANK economies aligns well with the empirical literature that provides nuanced evidence for long-run monetary neutrality (see, among many others, [Lucas, 1980](#), or [King and Watson, 1997](#)).

However, we demonstrate that non-neutrality becomes empirically relevant at high levels of inequality. We consider a reduction in the inflation target from 4% to 1.7% as above, but now occurring jointly with a rise in the wealth Gini index by 3 percentage points, the change in wealth inequality observed in the U.S. in the early 2000s. In this case, our model predicts a fall in the real rate of 36 bps, or 18% of the 200 bps reduction in the real rate estimated by [Del Negro et al. \(2017\)](#) and [Fiorentini et al. \(2018\)](#). The interaction of rising inequality and frequent ZLB events rationalizes the evidence in [Hillenbrand \(2021\)](#) and [Bianchi et al. \(2022\)](#) showing that monetary policy is partially responsible for the secular real rate decline. Thus, we provide a novel explanation for the fall in real rates that complements additional explanations such as population aging ([Carvalho et al., 2016](#), and [Aksoy et al., 2019](#)), the surge in convenience yields ([Krishnamurthy and Vissing-Jorgensen, 2012](#)), or the fall in productivity growth. In summary, our model encompasses situations where long-run neutrality of monetary policy approximately holds and situations where it does not.

Finally, we show that our model replicates the finding of [Long and Summers \(1986\)](#) that higher nominal rigidities ameliorate the deflationary spirals. Thus, we can interpret the deflationary bias in our model as coming from households' precautionary savings against negative demand shocks, not from the pricing behavior of firms.

Our paper connects with several previous papers. [Ascari \(2004\)](#) and [Ascari and Sbordone \(2014\)](#) show that, in the generalized New Keynesian model, high inflation levels lower output by creating more price dispersion. Our mechanism is different (and we pick a Rotemberg pricing protocol to emphasize this): our results work through precautionary savings. [Auclert and Rognlie \(2020\)](#) analyze the link between inequality and aggregate demand. They find that an increase in income inequality in the DSS can lead to a substantial decline in consumption, output, and the real rate, especially since they assume that the ZLB always constrains the central bank. In a similar spirit, [McKay and Reis \(2016\)](#) show that the presence of the ZLB matters for the effects of automatic stabilizers. Instead, our paper examines how inequality alters the ex-ante incentives to accumulate precautionary savings when households face an occasionally binding ZLB.

Our work also aligns with the literature focused on solving the full stochastic dynamics of HANK economies. For example, in independent work, [Gorodnichenko et al. \(2021\)](#) explore how stochastic volatility in aggregate productivity affects household consumption-saving decisions. [Schaab \(2020\)](#) proposes an alternative projection method to analyze how the interaction between macro and micro volatility can drive the economy into a downward aggregate demand spiral, leading to a liquidity trap. [Kase et al. \(2022\)](#) employ neural networks to estimate HANK economies subject to the ZLB constraint. From the methodological contribution, our paper builds on [Fernández-Villaverde et al. \(2023\)](#), who pioneered using neural networks to solve heterogeneous agent models globally. Our approach relates to [Maliar et al. \(2021\)](#) and [Azinovic et al. \(2022\)](#), who also employ deep learning to solve globally heterogeneous agent models with aggregate shocks.

The rest of the paper is organized as follows. Section 2 presents our ZLB-HANK model, which we calibrate in Section 3. Section 4 briefly presents our solution algorithm, which is discussed in more detail in Appendix A. Section 5 discusses our quantitative findings. Section 6 concludes.

2 Model

We postulate an economy populated by heterogeneous households, firms subject to nominal price rigidities, a central bank that follows a Taylor rule subject to the ZLB constraint, and a fiscal authority.

Households: There is a unit measure of ex-ante identical households indexed by $i \in [0, 1]$. Households face idiosyncratic and aggregate risk. The idiosyncratic labor earning shock $s_{i,t} \in \{s_m\}_{m=1}^M$ determines the efficiency unit of hours supplied by each household. The shock follows a Markov chain with normalized average realization $\int s_{it} di = 1$. The aggregate shock ξ_t is a preference shifter that evolves as an AR(1) process in logs, $\log \xi_t = \rho_\xi \log \xi_{t-1} + \zeta_t$, where $\rho_\xi \in (0, 1)$ and $\zeta_t \sim \mathcal{N}(0, \omega_\xi)$. Our use of a preference shifter follows [Krugman \(1998\)](#), [Eggertsson et al. \(2003\)](#), and [Eggertsson and Krugman \(2012\)](#), who show that shifts in households' preferences are a powerful driver of ZLB events. These preference shocks are a reduced form for any variation in uncertainty, fiscal policy, or credit market tightness that could alter households' risk appetite.

Households choose consumption $c_{i,t}$, bonds $b_{i,t}$ and labor services $h_{i,t}$ to maximize their life time expected discounted GHH utility:

$$\begin{aligned} & \max_{\{c_{i,t}, b_{i,t}, h_{i,t}\}_{t=0}^\infty} \mathbb{E}_0 \sum_{t=0}^\infty \beta^t \xi_t \frac{1}{1-\sigma} \left(c_{i,t} - \chi \frac{h_{i,t}^{1+\nu}}{1+\nu} \right)^{1-\sigma} \\ \text{s.t. } & c_{i,t} + b_{i,t} = (1-\tau) w_t s_{i,t} h_{i,t} + \Pi_t s_{i,t} + \frac{R_{t-1}}{\pi_t} b_{i,t-1}, \\ & b_{i,t} \geq \underline{b}, \end{aligned} \tag{1}$$

$$\tag{2}$$

where β is the time discount parameter, σ controls risk aversion, ν is the inverse of the Frisch elasticity of labor supply, and χ is the disutility of labor. We specify a GHH utility to eliminate any distraction in our results caused by wealth effects.

The households' optimization problem is subject to the budget constraint (1) and the borrowing constraint (2). The budget constraint posits that households finance consumption expenditures with firm profits Π_t , which are rebated to households depending on their idiosyncratic labor productivity, and labor earnings, $w_t s_{i,t} h_{i,t}$, where w_t denotes the real wage. Labor earnings are subject to a tax rate τ . Households also trade one-period non-contingent bonds, which yield the gross nominal return R_{t-1} . We refer to the gross real return on bonds as $r_t = \frac{R_{t-1}}{\pi_t}$, which divides R_{t-1} by gross inflation $\pi_t = \frac{P_t}{P_{t-1}}$. The borrowing constraint implies that the exogenous bound \underline{b} limits households' bond position.

Firms: A final good is produced by a representative final good firm and sold in a perfectly competitive final good market. The final good firm manufactures the final good Y_t by bundling together a continuum of intermediate inputs, $y_{j,t}$, indexed by $j \in [0, 1]$, by means of a CES production function:

$$Y_t = \left(\int_0^1 y_{j,t}^{\frac{\varepsilon-1}{\varepsilon}} dj \right)^{\frac{\varepsilon}{\varepsilon-1}},$$

where ε is the elasticity of substitution between the different intermediate goods. This CES aggregator yields an iso-elastic demand for intermediate good j , $y_{j,t} = \left(\frac{p_{j,t}}{P_t}\right)^{-\varepsilon} Y_t$, where $p_{j,t}$ is the price of variety j , and P_t is the aggregate price level. These two variables are linked through the equation $P_t = \left(\int_0^1 p_{j,t}^{1-\varepsilon} dj\right)^{\frac{1}{1-\varepsilon}}$.

Each intermediate good firm j produces its variety using labor $l_{j,t}$ and a linear technology $y_{j,t} = l_{j,t}$. Thus, marginal costs $m_{j,t}$ equal wages, $m_{j,t} = w_t$, and marginal costs and hired labor are equalized across producers, that is, $m_{j,t} = m_t$ and $l_{j,t} = l_t$, for all $j \in [0, 1]$.

The nominal rigidities in the model arise because intermediate good firms face price adjustment costs à la Rotemberg (1982). We follow Bayer et al. (2019) in the specification of the adjustment costs $\Theta_{j,t}$:

$$\Theta_{j,t} \equiv \Theta \left(\frac{p_{j,t}}{p_{j,t-1}} \right) = \frac{\theta}{2} \left[\log \left(\frac{p_{j,t}}{p_{j,t-1} \times \tilde{\pi}} \right) \right]^2 Y_t,$$

except that we extend their setting to the case of trend inflation $\tilde{\pi}$ targeted by the central bank, possibly different from zero. The parameter θ controls price stickiness. When $\theta = 0$, the economy has flexible prices and collapses to a standard neoclassical heterogeneous agent model with imperfect competition on the supply side. We choose Rotemberg costs adjusted for trend inflation instead of Calvo pricing because we will be conducting exercises where we modify the inflation target or the level of nominal rigidities. In these experiments, we want to exclude any effects caused by varying levels of price dispersion among firms, a mechanism unrelated to the main focus of our paper.

Given the price adjustment costs, the problem of intermediate good producers is to choose a sequence of prices $\{p_{j,t}\}_{t \geq 0}$ to maximize the expected discounted stream of profits net of the adjustment costs:

$$\mathbb{E}_t \sum_{k=t}^{\infty} \beta^k \left[\Pi_k(p_{j,k}) - \Theta \left(\frac{p_{j,k}}{p_{j,k-1}} \right) \right],$$

where $\Pi_k(p_{j,k}) = \left(\frac{p_{j,k}}{P_k} - m_k\right) \left(\frac{p_{j,k}}{P_k}\right)^{-\varepsilon} Y_k$ is the profits of the intermediate good firm.

In the pricing protocol, we follow Hagedorn et al. (2019) by assuming that the price adjustment costs are virtual. While the adjustment costs affect firms' optimal pricing decisions, they do not result in any transfer of real resources (e.g., they are effort costs). Accordingly, $\Pi_k(p_{j,t})$ is also the resources rebated back to each household according to its idiosyncratic productivity level. In such a way, we are consistent with the empirical observation that earnings-rich households tend to receive a disproportionately larger share of firm profits. As with the Rotemberg pricing protocol, virtual adjustment costs simplify the interpretation of the results as we change other model parameters.

The solution to the problem of the intermediate good firms yields a New Keynesian Phillips curve:

$$\log\left(\frac{\pi_t}{\tilde{\pi}}\right) = \beta\mathbb{E}_t\left[\log\left(\frac{\pi_{t+1}}{\tilde{\pi}}\right)\frac{Y_{t+1}}{Y_t}\right] + \frac{\varepsilon}{\theta}\left(m_t - \frac{\varepsilon - 1}{\varepsilon}\right).$$

Government: The government comprises a central bank and a fiscal authority. The central bank sets the gross nominal interest rate R_t according to the Taylor rule:

$$R_t = \max\left\{1, \tilde{R}\left(\frac{\pi_t}{\tilde{\pi}}\right)^{\phi_\pi}\left(\frac{Y_t}{\tilde{Y}}\right)^{\phi_y}\right\},$$

where \tilde{R} is the gross nominal rate and \tilde{Y} denotes the level of output, both at the DSS of the economy. The central bank sets R_t by reacting to changes in the inflation rate from its target –the parameter ϕ_π determines the strength with which this happens– and in output from its DSS level –the parameter ϕ_y pins down the strength of this second channel, unless the ZLB constraint, $R_t \geq 1$ is binding, in which case it sets $R_t = 1$.

The fiscal authority raises labor-earning taxes on the households to finance a fixed amount of outstanding debt \tilde{B} , such that the government budget constraint is satisfied $\int_0^1 \tau_t w_t s_{i,t} h_{i,t} di = (r_t - 1)\tilde{B}$. Given \tilde{B} and the equilibrium interest rate, the tax rate τ_t is set to clear the budget constraint.

Market clearing: Since the definition of competitive equilibrium for this economy is standard, we omit it for space considerations. Nonetheless, this competitive equilibrium implies three market-clearing conditions. First, the labor market clears so that the total efficiency units of hours provided by households equal the labor services demanded by the intermediate good firms: $\int_0^1 l_{j,t} dj = \int_0^1 s_{i,t} h_{i,t} di$. Second, the bond market clears so that the overall bond positions of households equal the outstanding government bonds issued by the fiscal authority: $\tilde{B} = \int_0^1 b_{i,t} di$. Finally, the aggregate resource constraint posits that total output equals the total value added by the intermediate good firms and the households' total consumption: $Y_t = \int_0^1 l_{j,t} dj = \int_0^1 c_{i,t} di$.

3 Calibration

We calibrate the model to U.S. quarterly data. Table 1 summarizes the parameterization of the model and the chosen targets.

We set the gross inflation target of the central bank to $\tilde{\pi} = \exp(0.02/4)$, so that the annual inflation target is 2%. For the Taylor rule, we set the sensitivity of the nominal rate to output deviations from the DSS to $\phi_y = 0.1$, and that to changes in inflation to $\phi_\pi = 2.5$.

Table 1: Baseline Parameterization

Parameter		Value	Target/Source
Panel A. Aggregate Risk			
ρ_ξ	AR coefficient of process for ξ	0.6	Bianchi et al. (2021)
ω_ξ	Standard deviation of ξ shock	0.0113	10% ZLB frequency
Panel B. Idiosyncratic Risk			
ρ_s	AR coefficient of process for s_t	0.94	10% Average MPC
ω_s	Standard deviation of s_t shock	0.075	30% Borrowers
\underline{b}	Borrowing limit	-0.58	Two-month average labor income
Panel C. Preferences			
β	Discount factor	0.9954	1.5% real interest rate in the DSS
σ	Risk aversion	1	Standard value
$1/\nu$	Frisch elasticity of labor supply	1	Standard value
χ	Disutility of labor	0.8707	Labor supply equals 1 in the DSS
Panel D. Production			
ε	Demand elasticity	7.67	15% price markup
θ	Rotemberg price adjustment cost	117	12-month duration of prices
Panel E. Central bank			
$\tilde{\pi}$	Inflation Target	$\exp(0.02/4)$	2% Annual inflation target
ϕ_π	Taylor rule coefficient on inflation	2.5	Standard value
ϕ_y	Taylor rule coefficient on output	0.1	Standard value
Panel F. Fiscal authority			
\tilde{B}	Government outstanding bonds	1	Total liquid assets = 25% annual GDP

We set $\beta = 0.9954$ such that the real interest rate in the DSS of the model equals 1.5%, a value in the ballpark of the estimates provided by [Del Negro et al. \(2017\)](#) and [Fiorentini et al. \(2018\)](#) and consistent with the November 2018 FOMC’s Summary of Economic Projections. In our quantitative results, we will show that different values of $\tilde{\pi}$ imply a large variation in the level of real interest rates in the SSS. This happens even if all the economies share the same value of real interest rates in the DSS.

We calibrate the preference shifter by making our model consistent with the 10% frequency of ZLB episodes observed in the U.S. economy in the post-war period ([Coibion et al., 2016](#)) by first setting $\rho_\xi = 0.6$, in line with the parameterization of [Bianchi et al. \(2021\)](#), and then fixing $\omega_\xi = 0.0113$.

Regarding the idiosyncratic risk, we first set the borrowing limit to $\underline{b} = -0.58$, equal to two months’ average wages. Then, we calibrate the labor earning process so that the model reproduces a share of borrowers of 30% ([Kaplan et al., 2014](#)), and an average MPC of

around 10%, which is at the lower end of the estimates provided by the literature (Johnson et al., 2006, Parker et al., 2013, and Broda and Parker, 2014). In such a way, we deliver quantitative results that are cautious regarding the importance of heterogeneity. These targets are matched by picking $\rho_s = 0.94$ and $\omega_s = 0.075$. We then convert the AR(1) process into a three-point Markov chain using Rouwenhorst’s (1995) method. These choices lead to a Gini index of wealth of 0.81, which replicates the value of the Gini index that Kuhn and Ríos-Rull (2016) find for the U.S. economy in the early 2000s.

We calibrate the remaining parameters of the households as follows. We fix $\sigma = 1$ and the Frisch elasticity to $1/\nu = 1$. This value is slightly higher than that proposed by Chetty et al. (2013) based on the elasticity of labor supply estimated at the micro level, but it is nonetheless at the lower end of the values typically considered by the literature on the transmission of monetary policy. Finally, we normalize the disutility of labor to $\chi = 0.8707$ such that the aggregate value of the efficiency units of hours equals one in the DSS.

On the production side, we set the demand elasticity across different intermediate goods to $\epsilon = 7.67$, such that the price markup is 15%, in line with the values commonly used in the New Keynesian literature (Christiano et al., 2005). The price adjustment parameter, $\theta = 117$, implies an equivalent Calvo-price protocol with a 12-month average price duration, the conventional value used in the literature.

Regarding the fiscal authority, we follow the calibration strategy of McKay et al. (2016) and fix $\tilde{B} = 0.25\%$ of annual GDP, in line with the estimate of liquid wealth in the U.S. economy derived by Kaplan et al. (2018).

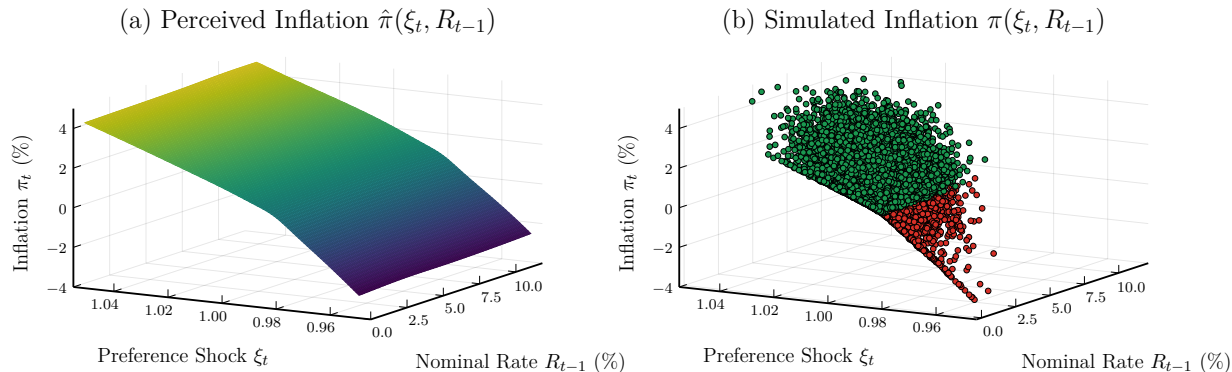
4 Solution Algorithm

One of the key contributions of this paper is the solution of a non-linear HANK model, which enables us to evaluate the stochastic dynamics of an economy with price rigidities, heterogeneous agents, and a ZLB.

To solve for the competitive equilibrium, we adopt the approach proposed by Fernández-Villaverde et al. (2023), who demonstrate the effectiveness of using neural networks to approximate the non-linear law of motion of household asset distributions, the key computational challenge of heterogeneous agent models. While we still assume that households forecast aggregate dynamics using a finite set of moments, as in Krusell and Smith (1998), we replace their state-dependent log-linear law of motion with a neural network that can capture the full non-linear dynamics of the economy. In our case, we have found that the best moment for households to keep track of is the level of the nominal interest rate (recall that the mean of the bond distribution is constant and, hence, cannot be used for that purpose). The nominal interest rate level summarizes the distribution’s shape since it encodes

the market clearing condition given the households' savings decisions. The nominal rate in HANK is also required to determine the interest rate income of households and, thus, becomes a state variable. In comparison, in RANK models, the nominal rate is not a state variable because the interest rate income is equal to the taxes paid by the representative agent and, therefore, drops from the decision rules.

Figure 1: Non-Linearity due to the ZLB.



In Appendix A, we provide a detailed description of our algorithm, but some important results are highlighted here. Specifically, Figure 1 illustrates how our algorithm uncovers a significant non-linearity in both the perceived and the simulated inflation rate for low values of the preference shifter ξ_t . When the ξ_t is sufficiently low, the nominal rate hits the ZLB and we have deflation (see Section 5).

Figure 2: Nowcast Errors for Inflation.

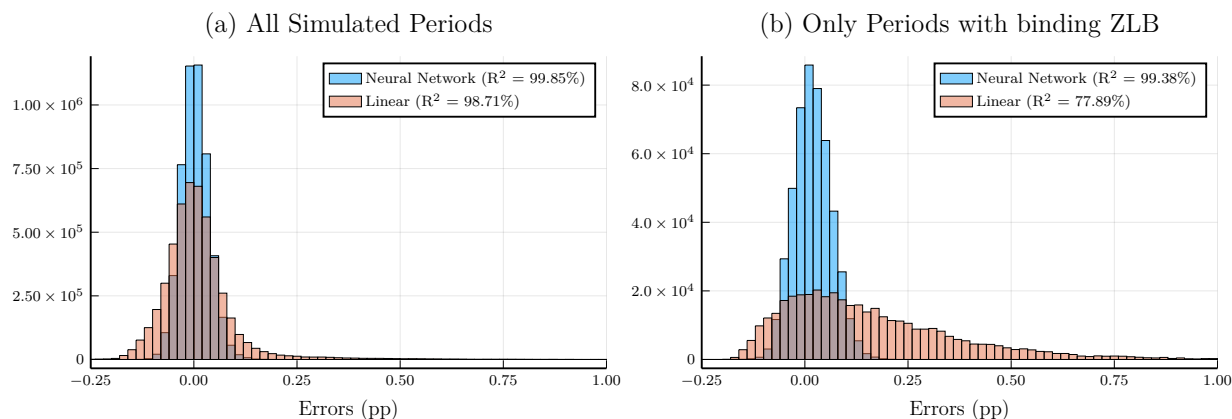


Figure 2 compares the nowcast errors on the inflation dynamics generated by our neural network approach with those implied by a naive application of the [Krusell and Smith \(1998\)](#) method, in which we predict the inflation rate as a log-linear function of the state variables

(see Appendix A for the corresponding formulae). Panel (a) reports the errors for all simulated periods: the neural network approach increases the fit of households' expectations by raising the R^2 from 98.71% to 99.85%. This improvement is particularly noteworthy in the right-hand part of the distribution of errors, which corresponds to periods in which inflation tends to be lower than expected due to the deflationary spirals triggered by a ZLB event. Panel (b) proves that our approach correctly captures the non-linear dynamics at the ZLB. While the linear approach provides an inferior fit for inflation dynamics when the ZLB constraint binds, with an R^2 of 77.89%, the neural network algorithm yields an R^2 of 99.38%, a much more satisfactory fit.

5 Results

This section reports how the presence of the ZLB affects household and aggregate variables in a HANK economy. We will uncover how the differences in the dynamics between our ZLB-HANK and the standard HANK model without the ZLB depend on both the inflation target and the level of wealth inequality. To begin, we compare the responses of aggregate variables and household decisions to aggregate demand shocks in both economies. We then examine how the ZLB alters the SSS of our model. Finally, we gauge the effects of varying the central bank's inflation target and the level of wealth inequality on aggregate variables.

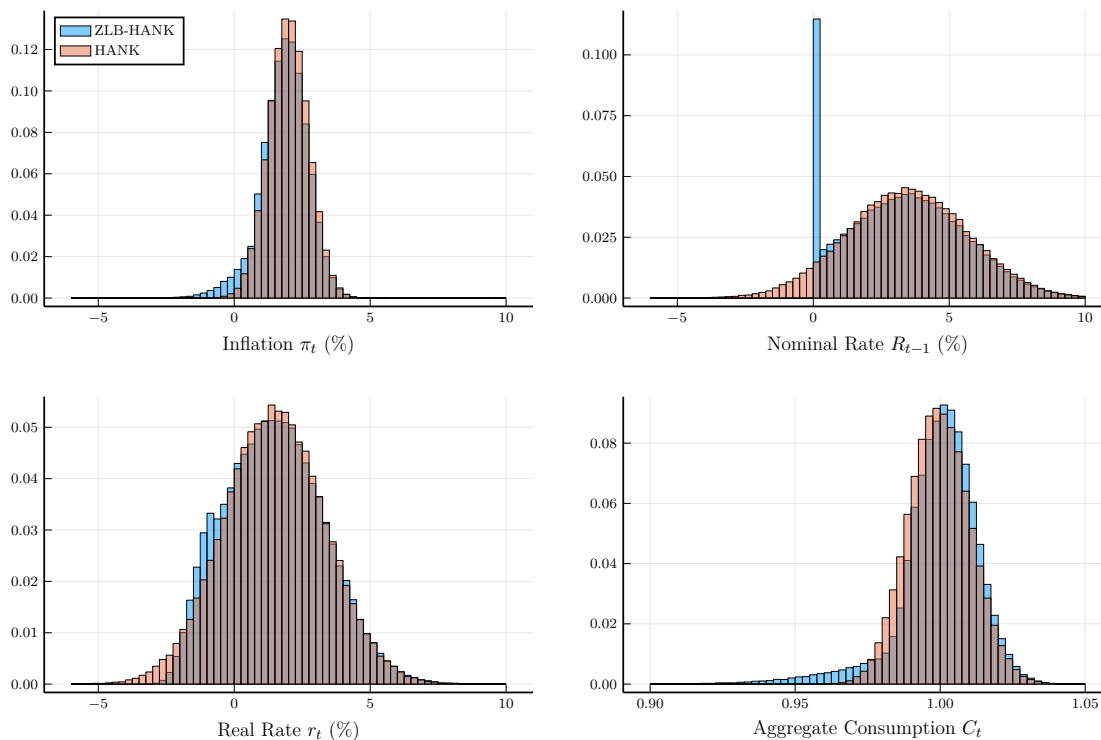
5.1 The Macroeconomics of the ZLB

As our paper is the first to explicitly consider the non-linearity generated by the ZLB constraint when studying the stochastic dynamics of a HANK model, it is important to understand the implications of this non-linearity for the model's dynamics. An intuitive way to appreciate this point is by comparing the ergodic distribution of the aggregate variables in our ZLB-HANK model with those of a HANK economy with the same calibration as in Section 3 except that the ZLB is absent from the Taylor rule that the central bank follows.

Figure 3 reports the ergodic distribution of inflation, the nominal interest rate, the real interest rate, and aggregate consumption in the ZLB-HANK and HANK economies. The graph shows how the presence of the ZLB skews the dynamics of the model to the left, except for the real rate, where the left tail gets truncated. These are the cases in which the ZLB constrains the nominal interest rate, and the economy experiences a sharp drop in aggregate consumption amidst a deflationary spiral and real rates that are too high. These dynamics are absent in the HANK economy.

Figure 4 evaluates the effect of the non-linearity of the ZLB on the dynamics of inflation and output by plotting the impulse-response functions (IRFs) at the SSS to small and large

Figure 3: Ergodic Distributions.



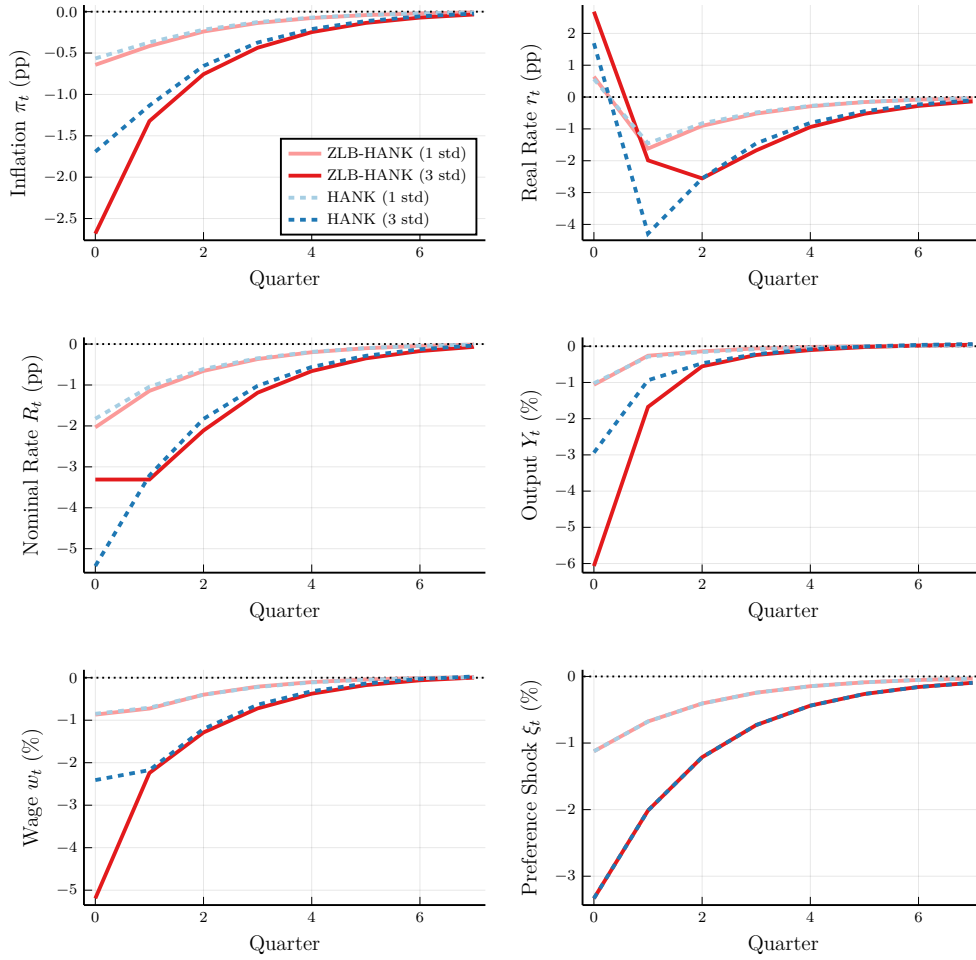
demand shocks in the ZLB-HANK and the HANK economies.² Specifically, we consider a small demand shock (a one-standard-deviation innovation) and a large demand shock (a three-standard-deviation innovation). When the shock is small, the ZLB is not binding in either economy, and the IRFs roughly coincide.³ However, when the shock is large, the ZLB generates stark differences in the responses of the two economies. In the ZLB-HANK economy, the nominal interest rate drops to zero, while in the HANK economy, it drops below -5% . As a result, the ZLB-HANK economy experiences a much sharper decline in both inflation and output.

Figure 4 emphasizes the crucial role played by the ZLB in amplifying the negative effects of large negative demand shocks on the economy when the shocks are too severe to be accommodated by a sufficiently large reduction in the nominal interest rate. Notice also how real wages fall significantly more at the ZLB, whereas real interest rates remain higher (as we have deflation but the nominal rate is stuck at zero). We discuss next how these changes in real prices have important redistributive consequences.

²Notice that our model is non-linear. Hence, the IRFs are state-dependent and the size and sign of the shock matter for their shape.

³Agents in the ZLB-HANK economy understand that the ZLB might bind in the future and, therefore, their current behavior is slightly different from that of agents in the HANK economy even outside of the ZLB. We will revisit this point below.

Figure 4: Impulse Response Functions.




5.2 The Microeconomics of the ZLB

This subsection demonstrates that the ZLB significantly impacts the distributional consequences of negative demand shocks over the business cycle. Specifically, we analyze the differential impact response of total income to a three-standard-deviation negative demand shock that arrives when the economy is at the SSS in the ZLB-HANK and HANK economies (the same aggregate shock as in the previous subsection). Figure 5 displays the changes in total income at impact (represented by diamonds) and decomposes them into contributions from changes in taxes, interest earnings, wages, and profits for households in the 10th and 99th percentile of the wealth distribution, across the three realizations of the labor shock. We also report the aggregate relative response of income.

Figure 5 provides two valuable insights. First, in the ZLB-HANK economy, income falls more sharply than in the HANK economy across all households, regardless of their labor earnings and position in the wealth distribution. This is a direct consequence of the

Figure 5: Individual Income Responses - ZLB-HANK vs. HANK.




Note: The graph reports the income response for each of the three labor-earning realizations at the 10th wealth percentile and the 99th wealth percentile, and then each bar decomposes the total income response to changes in taxes, interest revenues, wages, and profits.

central bank’s inability to accommodate the negative demand shock fully. Second, the ZLB constraint amplifies the decline in total income for wealth-poor households relative to wealth-rich individuals. This is because poorer households rely more on wages, which experience a sharp decline, and have either low or negative bond holdings (in our calibration, households at the 10th percentile have negative bond holdings). Since the real rates are higher because the ZLB prevents nominal rates from falling below zero despite deflation, wealth-rich households earn significantly more from their bonds. In contrast, households with low bond holdings earn little more, and those with negative holdings have to pay a higher real interest.

The ZLB also significantly affects the distributional consequences of negative demand shocks in terms of consumption. Figure 6 reports the differences in the impact response of consumption between the ZLB-HANK and HANK economies to the same shock as above, for households in the 10th and 99th percentiles of the wealth distribution, across the three realizations of the labor shock, as well as the differential response of aggregate consumption. Consumption drops more in the ZLB-HANK economy than in the HANK economy, a drop that is amplified for households at the lower end of the wealth distribution, which lack the assets to smooth consumption and cannot borrow much more or at all (8% of households are at the borrowing constraint at the SSS). More concretely, the presence of the ZLB amplifies the drop in consumption of wealth-poor individuals by about 3.5 percentage points.

Figure 6: Consumption Responses - ZLB-HANK vs. HANK.



Note: The graph reports the consumption response for each of the three labor-earning realizations at the 10th wealth percentile and the 99th wealth percentile.

Our results highlight the importance of considering the ZLB in analyzing the distributional consequences of business cycle fluctuations in a HANK economy. At the aggregate level, recessions become more severe after a large negative demand shock due to the deflationary spiral that arises when the central bank cannot fully accommodate the shock. At the individual level, the burden of recessions is even greater for wealth-poor households, resulting in increased consumption, income, and wealth inequality during severe downturns.

5.3 Why Does Heterogeneity Matter for Monetary Policy?

Standard models tend to predict that structural parameters pin down the real interest rate in the DSS, so there is no role for monetary policy to influence it. Although our model features a similar result in the DSS, it generates a relationship between changes in the inflation target and changes in real interest rates in the SSS.

For instance, in the DSS of our model, the inflation rate is 2%, which coincides with the central bank's inflation target. Instead, at the SSS, the level of inflation is lower, 1.90%. This happens because households internalize the possible occurrence of sizeable negative demand shocks that may bring the economy to the ZLB constraint, triggering deflationary episodes. In the following paragraphs, we will highlight how our model implies that changes in the level of the central bank's inflation target move the real interest rates in the SSS even if they do not do so in the DSS.

To show how heterogeneity matters for monetary policy, we start by comparing the DSS and SSS in the ZLB-HANK and HANK economies we used before, as well as two versions of our model with the same calibration as in Section 3 except that we shut down idiosyncratic labor risk to get a representative household, the ZLB-RANK (where the ZLB still exists) and RANK (no ZLB) economies.

Table 2: Comparison of DSS and SSS in ZLB-HANK, HANK, and ZLB-RANK.

Variable	ZLB-HANK		HANK		ZLB-RANK		RANK	
	DSS	SSS	DSS	SSS	DSS	SSS	DSS	SSS
Inflation	2.00%	1.90%	2.00%	1.99%	2.00%	1.91%	2.00%	1.99%
Nominal Rate	3.50%	3.31%	3.50%	3.47%	3.84%	3.67%	3.84%	3.82%
Real Rate	1.50%	1.41%	1.50%	1.48%	1.84%	1.75%	1.84%	1.82%
(Shadow) ZLB Frequency	-	9.72%	-	(5.88%)	-	8.36%	-	(5.09%)

Table 2 reports the results of our comparison exercise, showing the DSS and SSS values for inflation, the nominal interest rate, the real interest rate, and output, along with the frequency of ZLB events in the ergodic distribution of each model. For the HANK and RANK economies, we report the shadow frequency and duration of ZLB events, defined as any period where the nominal interest rate equals or falls below zero.

Both the ZLB-HANK and HANK economies exhibit identical values for all macroeconomic variables at the DSS since the ZLB is not binding at that point. However, the SSS values of the two economies diverge. Specifically, the HANK economy displays SSS values for inflation, the nominal rate, and the real rate that are virtually identical to those at the DSS, differing by only 1 bp, 3 bps, and 2 bps, respectively. In contrast, introducing the ZLB into the ZLB-HANK economy reduces SSS values for inflation, the nominal rate, and the real rate by approximately 10 bps, 19 bps, and 9 bps, respectively.

So, how does the introduction of the ZLB explain these differences? The ZLB changes the behavior of households and firms at the SSS, even if it is not binding at that point. Agents understand that a demand shock could push the economy toward the ZLB in the future and react preemptively to it. After all, the ZLB-HANK economy spends 9.72% of quarters at the ZLB (our calibration target was 10%). In particular, households increase their precautionary savings to ensure a buffer of savings (or reduce their borrowing amounts) to smooth their consumption stream in those recessions where the nominal rate hits zero. This effect is particularly salient for wealth-poor households, which we saw above suffer disproportionately

more from spells at the ZLB. Higher precautionary savings exert downward pressure on the real interest rate level, reducing the central bank’s room for maneuvering of the nominal rates and making the ex-post realization of the ZLB events even more likely.

In comparison, the shadow frequency of ZLB events in the HANK economy is lower at 5.88% since the central bank can accommodate negative demand shocks with aggressive reductions of the nominal interest rate more effectively, and households have a smaller need for precautionary behavior.

How does the variation between ZLB-HANK and HANK compare to that between ZLB-RANK and RANK? Precautionary savings still mean that the SSS of the ZLB-RANK diverges from its DSS. However, the divergence is smaller than between ZLB-HANK and HANK. For instance, the frequency of ZLB spells goes from a (shadow) 5.09% in RANK to 8.36% in ZLB-RANK (a difference of 3.27%) while it goes from a (shadow) 5.88% in HANK to 9.72% in ZLB-HANK (a difference of 3.84%). Furthermore, as we will document later, this divergence will become more acute when we look at a situation with lower inflation targets and higher wealth heterogeneity.

Table 3: Decomposition Exercise.

	Real Rate	Nominal Rate	Inflation
ZLB-RANK DSS	1.84%	3.84%	2.00%
ZLB-HANK SSS	1.41%	3.31%	1.90%
(i) Total	0.43pp	0.53pp	0.10pp
ZLB-RANK DSS	1.84%	3.84%	2.00%
ZLB-HANK DSS	1.50%	3.50%	2.00%
(ii) Precautionary Savings Idiosyncratic Risk	0.34pp	0.34pp	0.0pp
(i)-(ii) Deflationary Bias	0.09pp	0.19pp	0.10pp

Table 3 decomposes the differences in the real rate, the nominal rate, and inflation: the effect of precautionary savings induced by idiosyncratic labor risk and a deflationary bias. To explain this decomposition, let us concentrate on the difference in the real interest rate between the ZLB-RANK DSS, 1.84%, and the ZLB-HANK SSS, 1.41%, which is 43 bps lower in the latter economy (nearly a quarter of the level). Recall that, in the ZLB-RANK economy, the real interest rate at the DSS is uniquely pinned down by the time discount

factor β :

$$1.84\% = \left(\frac{1}{\beta = 0.9954} \right)^4.$$

Next, we compute the real interest rate at the ZLB-HANK DSS, 1.50%. Since there is no aggregate risk at the DSS, the ZLB is never binding, and β is still 0.9954. The effect on the real rate comes exclusively from the precautionary savings induced by idiosyncratic labor risk (Aiyagari, 1994). Thus, this precautionary savings channel accounts for 34 bps of the 43 bps of the total difference.

Then, we can return to the difference between the ZLB-HANK DSS and SSS from Table 3, which is 9 bps. This additional difference is caused by the precautionary savings induced by aggregate risk, in particular, the possibility of hitting the ZLB and its deflationary spiral, and the changes in the pricing decisions of firms. We call this channel the “deflationary bias.” Households save more to avoid the negative consequences of the ZLB for consumption smoothing, which reduces the real and nominal interest rates in the ergodic distribution, but also makes it more likely that a negative demand shock will push the economy to the ZLB and trigger a deflationary spiral. Similarly, firms bias their price changes downward to avoid paying large price adjustment costs when deflation hits the economy (see, for a similar mechanism, Fernández-Villaverde et al., 2015b). This deflationary bias also exists in ZLB-RANK economies because the representative agent also wants to ensure consumption against spells at the ZLB and firms change their pricing decisions. See the columns for ZLB-HANK and HANK in Table 2, which confirm previous findings by Adam and Billi (2007), Nakov (2008), Hills et al. (2019), and Bianchi et al. (2021).

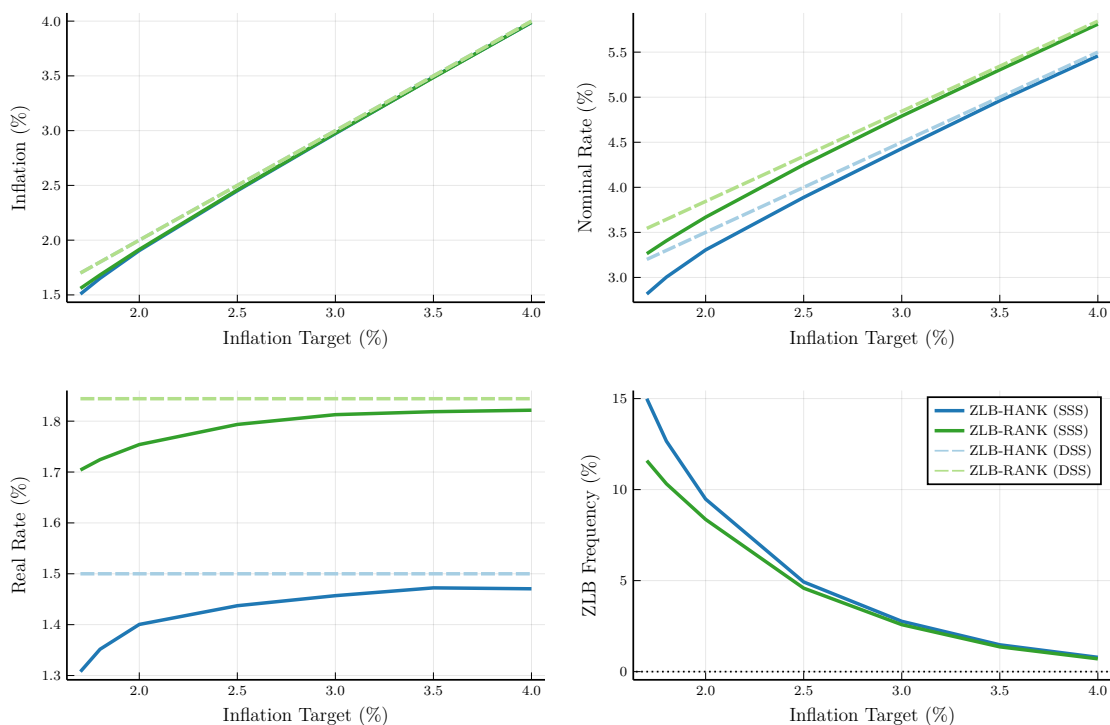
What makes our exercise novel is to show how, by lowering the real and nominal interest rates, household heterogeneity reduces the scope of central banks to accommodate negative demand shocks when the ZLB is present more than what occurs in the ZLB-RANK economy. Indeed, the ZLB-HANK economy has a frequency of ZLB events of 9.72%, while this frequency is 8.36% in the ZLB-RANK case. Thus, our paper adds a new layer to the insight into the relationship between real rates and household heterogeneity. This result does not emerge in the standard HANK literature, in which the drop in the real rate due to precautionary savings is immaterial for aggregate dynamics because the ZLB is not present.

More pointedly, our analysis shows that the real rate in the ZLB-HANK economy is endogenously determined by the interaction of the ZLB and households’ heterogeneity. Thus, factors affecting the probability of hitting the ZLB (e.g., the inflation target) or determining household heterogeneity (e.g., the variance of the idiosyncratic risk) will move the real interest rates in the long run. We move to explore these possibilities now.

5.4 Inflation Target and Real Interest Rates

We just saw how, in the ZLB-HANK economy, the presence of spells at the ZLB changes the real interest rate even when we are not at the ZLB. This observation raises an intriguing possibility. Since the central bank can alter the frequency of the ZLB spells by modifying its inflation target, it affects the level of real interest rates through the change in the households' savings behavior. In other words, monetary policy is not neutral, even in the long run. More specifically, the model features a long-run Fisher equation, $i(\tilde{\pi}) = r(\tilde{\pi}) + \pi(\tilde{\pi})$, in which the real rate in the SSS depends on the central bank's inflation target $\tilde{\pi}$. In this setting, a higher inflation target raises the SSS level of the real rate, that is, $dr/d\tilde{\pi} > 0$.

Figure 7: DSS/SSS in ZLB-RANK and ZLB-HANK as a Function of the Inflation Target.



To establish this result, we compare the real interest rate level in different model economies, which differ only in their inflation target $\tilde{\pi}$. Figure 7 plots the DSS and SSS levels of inflation, the real rate, the nominal rate, and the frequency of the ZLB for $\tilde{\pi}$ between 1.7% and 4% in both the ZLB-RANK and ZLB-HANK economies. First, the graph shows that, in both economies, the central bank successfully achieves its inflation target at the SSS when $\tilde{\pi}$ falls within the range of 3% to 4%. For those high inflation targets, the probability of experiencing a ZLB event is so low that the non-linearity of the model is not quantitatively relevant. The economy behaves in practice as if there is no ZLB constraint, and changes in the inflation target between 3% and 4% do not alter the level of real interest rates.

However, when the inflation target goes below 3%, the non-linearity caused by the ZLB kicks in, and the SSS and DSS diverge substantially. First, the central bank undershoots its inflation target. When the target is 1.7%, the inflation rate in the SSS is 1.51%.⁴ This 19 bps undershooting comes together with the economy spending 15% of quarters at the ZLB, compared with less than 1% of quarters when the inflation target is 4%. Importantly, the ZLB-HANK economy is more sensitive with respect to the probability of hitting the ZLB as a function of the inflation target than the ZLB-RANK economy. In the latter, the proportion of quarters spent at the ZLB goes from less than 1% when the inflation target is 4% to 11.6% when the inflation target is 1.7%.

As we vary the inflation target, the real interest rate at the DSS does not change (dashed lines in the bottom left panel of Figure 7) since idiosyncratic labor income risk is either independent of the inflation target in the ZLB-HANK economy or non-existent in the ZLB-RANK economy. In comparison, the real rate in the SSS increases with the inflation target because we reduce the probability of hitting the ZLB (solid lines in the bottom left panel of Figure 7). Notably, the sensitivity of the real rate in the SSS to changes in the inflation target is greater in the ZLB-HANK economy compared to the ZLB-RANK economy (i.e., $dr/d\tilde{\pi}$ in ZLB-HANK is larger than $dr/d\tilde{\pi}$ in ZLB-RANK). In other words, if we were to look at the DSS, we would conclude erroneously that the monetary policy stance does not affect the real interest rate level in the economy’s ergodic distribution. Moreover, even if we were to look at the SSS of an economy without heterogeneity, we would miss the strength of the relationship between the inflation target and the real rate.

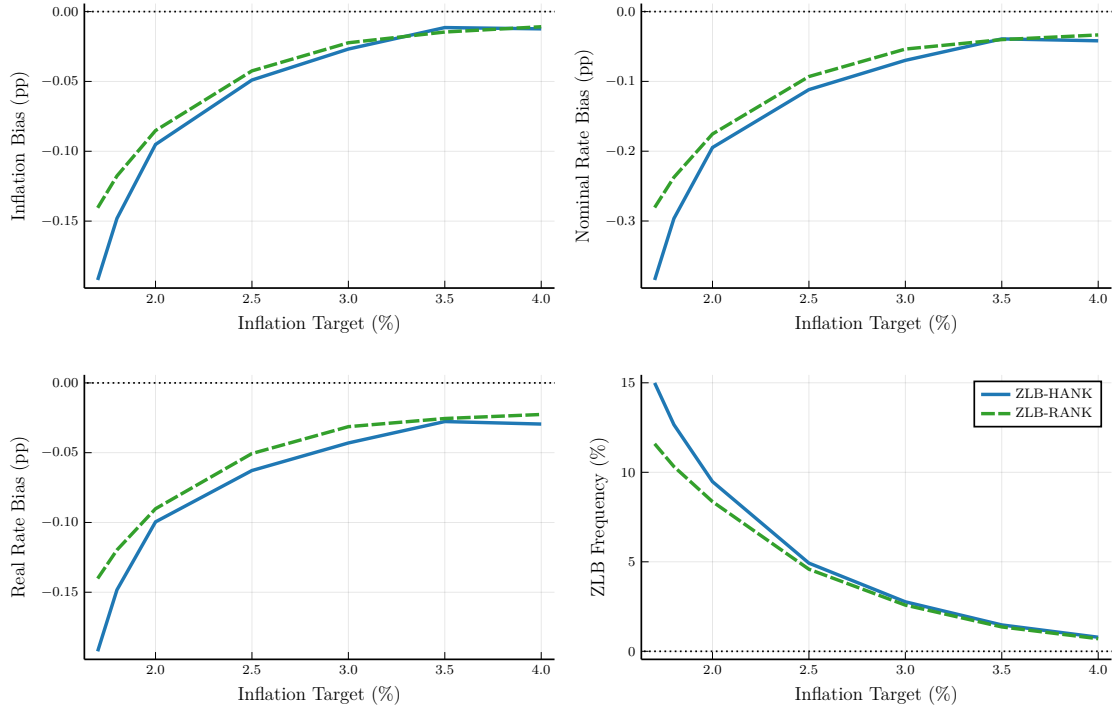
These dynamics can be better appreciated in Figure 8, which reports the differences between the SSS and DSS values of the ZLB-RANK and ZLB-HANK economies at each value of the inflation target. For instance, the drop in the real rate in the ZLB-RANK going from an inflation target of 4% to 1.7% is 16 bps vs. 12 bps in the ZLB-HANK economy.

What drives the relationship between the inflation target and the real and nominal interest rates? To inspect this relationship, let us refer again to the Fisher equation. In the DSS of standard models, the real rate is the inverse of the households’ time discount parameter, whereas the level of inflation is a policy parameter. Given these two structural parameters, standard models determine the level of the nominal interest rate.

In the ZLB-HANK economy, in comparison, we have the deflationary bias we described above triggered by the possibility of hitting the ZLB, which lowers both the real rate and inflation below its target, with both forces lowering the nominal rate. While this deflationary

⁴Notice that we are reporting that this undershooting is happening at the SSS, where the economy is not at the ZLB. In other words, the undershooting is caused by the agents reacting to the possibility of being at the ZLB in the future, not because of the deflationary spiral triggered when we hit the ZLB.

Figure 8: Differences between the SSS and DSS as a Function of the Inflation Target.



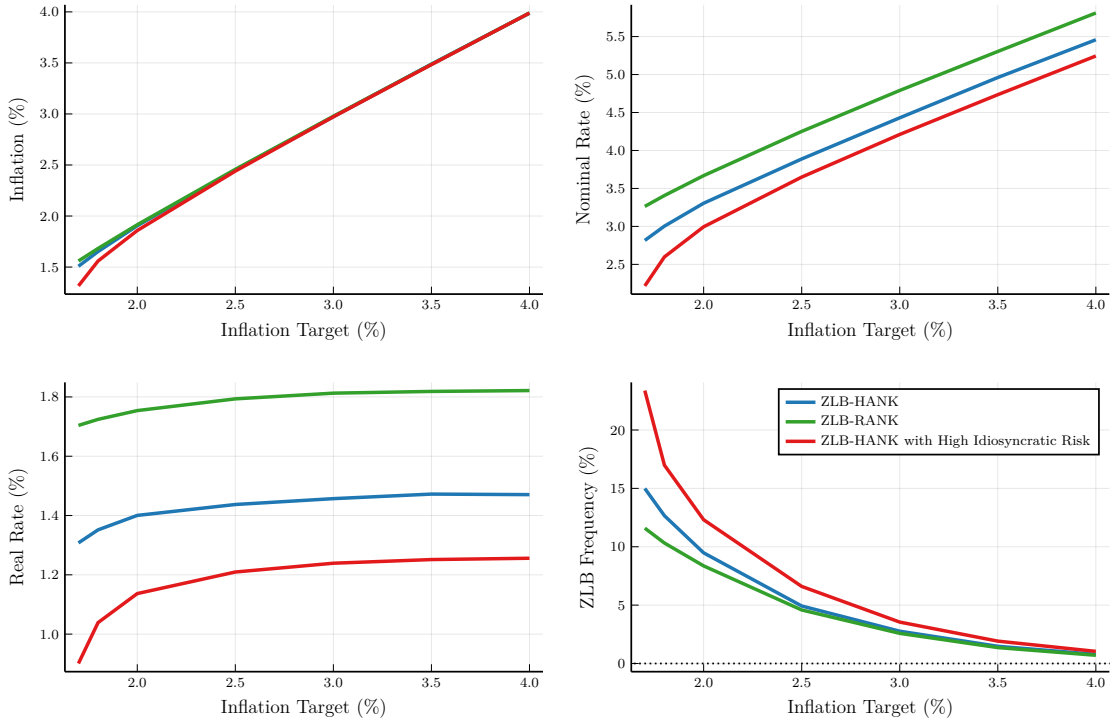
bias also exists in the ZLB-RANK economy, its importance is larger in the ZLB-HANK case because idiosyncratic uncertainty also lowers the real rate and makes hitting the ZLB more likely. While the non-neutrality of the inflation target in the ZLB-RANK economy is limited, the following subsection shows that changes in $\tilde{\pi}$ can lead to quantitatively relevant variations in the real rate at a high level of wealth inequality in the ZLB-HANK economy.

5.5 The Role of Wealth Inequality

In our previous discussions, we established that household heterogeneity matters for the real interest rate and, with it, for the long-run non-neutrality of monetary policy. Now we show that this importance grows with the level of wealth inequality. To prove this point, Figure 9 reports the SSS values for the inflation target, the nominal interest rate, the real rate, and the ZLB frequency for ZLB-HANK, ZLB-RANK, and ZLB-HANK with high idiosyncratic labor income risk. More specifically, we increase the standard deviation of the labor earning risk process, ω_s , so the Gini index grows from 0.81 in our baseline calibration in Section 3 to 0.84. This increase in the Gini index corresponds to the change observed in the U.S. economy in the early 2000s (Kuhn and Ríos-Rull, 2016).

We find that the non-neutrality of monetary policy depends substantially on the level of wealth inequality. When we increase idiosyncratic labor-income risk, the drop in the real rate when the inflation target goes from 4% to 1.7% is 36 bps (in ZLB-RANK, since there

Figure 9: The Interaction Between the Inflation Target and Wealth Inequality.



is no heterogeneity, the drop is still the 12 bps we reported in the previous subsection).

Why does higher wealth inequality result in a more pronounced decline in the real rate at lower levels of the inflation target? Once again, the answer lies in the role of precautionary savings. In the baseline economy with low idiosyncratic risk, the relatively infrequent occurrence of the ZLB implies that households do not significantly increase their precautionary savings with respect to the case without the ZLB even when the inflation target is low. However, in the high idiosyncratic risk economy, the frequency of encountering the ZLB rises dramatically, reaching as high as 23% when the inflation target is set at 1.7%. The high probability of a ZLB spell, coupled with a higher risk of encountering a ZLB spell when the labor income shock is low, significantly increases households' savings demand. This shift in savings demand leads to a substantial decrease in the level of the real rate.

Our results are consistent with the observed dynamics of the U.S. economy over recent decades. As the average inflation rate has shrunk from around 4% in the 1980s and 1990s to below 2% in the 2000s –and these changes have happened contemporaneously with a secular rise in the amount of wealth inequality, the economy has started experiencing ZLB events. Consequently, until the recent burst of inflation, the inflation rate has constantly been below the 2% target, and the level of real interest rates has dropped substantially. Our results provide a novel rationale that jointly accounts (at least partially) for all these events.

5.6 The Role of Nominal Rigidities

We explained above that the deflationary bias is caused by households' desire to self-insure against low wages at the ZLB by saving more, and firms biasing their price changes downward to avoid being forced to pay high price adjustment costs at the ZLB.

To disentangle both mechanisms, we increase the value of θ , the Rotemberg price adjustment cost, from 117 (equivalent to a 12-month average price duration) to 200 (equivalent to an 18-month average price duration). This way, firms become even more concerned about paying high adjustment costs when a large negative demand shock hits the economy. Interestingly, higher adjustment costs reduce the deflationary bias in the ZLB-HANK model from 9 bps to 5 bps. A similar result appears in the ZLB-RANK economy.

The result is not surprising. It has been known since [Long and Summers \(1986\)](#) that higher nominal rigidities ameliorate deflationary spirals. Our model replicates this finding. Thus, we can interpret the deflationary bias as arising from households' precautionary savings against negative demand shocks. Nominal price rigidities make the deflationary bias smaller by stabilizing the economy (except in the limit case of full price flexibility when the ZLB is inconsequential), an effect that more than counterbalances the downward pricing bias of firms.

6 Conclusion

This paper presents a methodological contribution and a set of substantive findings. Methodologically, we show how to solve a ZLB-HANK model non-linearly using neural networks and document how this approach significantly affects the model's quantitative implications.

In terms of substantive findings, we make several key observations. First, the ZLB exacerbates the effects of large negative demand shocks in a HANK economy by preventing the central bank from fully accommodating them. Second, we show that the ZLB disproportionately impacts wealth-poor households, which experience larger drops in consumption due to the recession's impact on wages.

Third, as in [Aiyagari \(1994\)](#), in our model, household heterogeneity is linked to the real interest rate, even without aggregate shocks. Thus, we demonstrate that heterogeneity makes the economy more vulnerable to the ZLB than an economy with a representative household, as smaller negative shocks can push it into ZLB territory. This effect is driven by precautionary saving, which paradoxically makes the ZLB more likely.

Fourth, we prove that monetary policy is not neutral in the long run, as the inflation target chosen by the central bank affects the level of precautionary savings and, consequently, the real interest rate. This mechanism is absent from standard HANK models or models

that assume the ZLB is only an initial state but not a recurrent one.

Fifth, the non-neutrality of monetary policy is amplified when wealth inequality is high. Our calibrated ZLB-HANK model predicts that a drop in the inflation target from 4% to 1.7% combined with a 3 percentage point increase in the Gini index of wealth inequality (which corresponds to the observed rise in wealth inequality in the U.S. in the early 2000s) reduces the level of the real interest rate by around 36 bps. Although other factors like the aging of population are also likely to play a role in the fall of the real rate, we prove the role of monetary policy and wealth inequality in this phenomenon.

References

- Adam, K. and Billi, R. M. (2007). Discretionary monetary policy and the zero lower bound on nominal interest rates. *Journal of Monetary Economics*, 54(3):728–752.
- Ahn, S., Kaplan, G., Moll, B., Winberry, T., and Wolf, C. (2018). When inequality matters for macro and macro matters for inequality. *NBER Macroeconomics Annual*, 32(1):1–75.
- Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109(3):659–684.
- Aksoy, Y., Basso, H. S., Smith, R., and Grasl, T. (2019). Demographic structure and macroeconomic trends. *American Economic Journal: Macroeconomics*, 11(1):193–222.
- Andrade, P., Galí, J., Le Bihan, H., and Matheron, J. (2019). The optimal inflation target and the natural rate of interest. *Brookings Papers on Economic Activity*, Fall 2019 Edition.
- Ascari, G. (2004). Staggered prices and trend inflation: Some nuisances. *Review of Economic Dynamics*, 7(3):642–667.
- Ascari, G. and Sbordone, A. M. (2014). The macroeconomics of trend inflation. *Journal of Economic Literature*, 52(3):679–739.
- Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–67.
- Auclert, A., Bardóczy, B., Rognlie, M., and Straub, L. (2021). Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408.
- Auclert, A. and Rognlie, M. (2020). Inequality and aggregate demand. Technical report, Mimeo.
- Azinovic, M., Gaegauf, L., and Scheidegger, S. (2022). Deep equilibrium nets. *International Economic Review*, 63(4):1471–1525.
- Bayer, C., Lüticke, R., Pham-Dao, L., and Tjaden, V. (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica*, 87(1):255–290.
- Bewley, T. (1980). Interest bearing money and the equilibrium stock of capital. Technical report, Mimeo.

- Bianchi, F., Lettau, M., and Ludvigson, S. C. (2022). Monetary policy and asset valuation. *Journal of Finance*, 77(2):967–1017.
- Bianchi, F., Melosi, L., and Rottner, M. (2021). Hitting the elusive inflation target. *Journal of Monetary Economics*, 124:107–122.
- Broda, C. and Parker, J. A. (2014). The economic stimulus payments of 2008 and the aggregate demand for consumption. *Journal of Monetary Economics*, 68:S20–S36.
- Carvalho, C., Ferrero, A., and Nechio, F. (2016). Demographics and real interest rates: Inspecting the mechanism. *European Economic Review*, 88:208–226.
- Chetty, R., Guren, A., Manoli, D., and Weber, A. (2013). Does indivisible labor explain the difference between micro and macro elasticities? A meta-analysis of extensive margin elasticities. *NBER Macroeconomics Annual*, 27(1):1–56.
- Christiano, L., Eichenbaum, M., and Rebelo, S. (2011). When is the government spending multiplier large? *Journal of Political Economy*, 119(1):78–121.
- Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45.
- Coeurdacier, N., Rey, H., and Winant, P. (2011). The risky steady state. *American Economic Review*, 101(3):398–401.
- Coibion, O., Dordal-i Carreras, M., Gorodnichenko, Y., and Wieland, J. (2016). Infrequent but long-lived zero lower bound episodes and the optimal rate of inflation. *Annual Review of Economics*, 8(1):497–520.
- Coibion, O., Gorodnichenko, Y., and Wieland, J. (2012). The optimal inflation rate in New Keynesian models: Should central banks raise their inflation targets in light of the zero lower bound? *Review of Economic Studies*, 79(4):1371–1406.
- Del Negro, M., Giannone, D., Giannoni, M. P., and Tambalotti, A. (2017). Safety, liquidity, and the natural rate of interest. *Brookings Papers on Economic Activity*, 2017(1):235–316.
- Eggertsson, G. B. et al. (2003). Zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity*, 2003(1):139–233.
- Eggertsson, G. B. and Krugman, P. (2012). Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach. *Quarterly Journal of Economics*, 127(3):1469–1513.

- Fernández-Villaverde, J., Gordon, G., Guerrón-Quintana, P., and Rubio-Ramírez, J. F. (2015a). Nonlinear adventures at the zero lower bound. *Journal of Economic Dynamics and Control*, 57:182–204.
- Fernández-Villaverde, J., Guerrón-Quintana, P., Kuester, K., and Rubio-Ramírez, J. (2015b). Fiscal volatility shocks and economic activity. *American Economic Review*, 105(11):3352–84.
- Fernández-Villaverde, J., Hurtado, S., and Nuño, G. (2023). Financial frictions and the wealth distribution. *Econometrica*, 91(3):869–901.
- Fiorentini, G., Galesi, A., Pérez-Quirós, G., and Sentana, E. (2018). The rise and fall of the natural interest rate. Technical report, Mimeo.
- Gorodnichenko, Y., Maliar, L., Maliar, S., and Naubert, C. (2021). Household savings and monetary policy under individual and aggregate stochastic volatility. Technical report, Mimeo.
- Guerrieri, V. and Lorenzoni, G. (2017). Credit crises, precautionary savings, and the liquidity trap. *Quarterly Journal of Economics*, 132(3):1427–1467.
- Hagedorn, M., Manovskii, I., and Mitman, K. (2019). The fiscal multiplier. Technical report, Mimeo.
- Hillenbrand, S. (2021). The Fed and the secular decline in interest rates. Technical report, Mimeo.
- Hills, T. S., Nakata, T., and Schmidt, S. (2019). Effective lower bound risk. *European Economic Review*, 120:103321.
- Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5-6):953–969.
- Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5):1589–1610.
- Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3):697–743.
- Kaplan, G., Violante, G. L., and Weidner, J. (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*, 1:77–153.

- Kase, H., Melosi, L., and Rottner, M. (2022). Estimating nonlinear heterogeneous agents models with neural networks. Technical report, Mimeo.
- King, R. G. and Watson, M. W. (1997). Testing long-run neutrality. *Federal Reserve Bank of Richmond Economic Quarterly*, 83(3):69–101.
- Krishnamurthy, A. and Vissing-Jorgensen, A. (2012). The aggregate demand for Treasury debt. *Journal of Political Economy*, 120(2):233–267.
- Krugman, P. R. (1998). It’s baaack: Japan’s slump and the return of the liquidity trap. *Brookings Papers on Economic Activity*, 1998(2):137–205.
- Krusell, P. and Smith, Jr, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896.
- Kuhn, M. and Ríos-Rull, J.-V. (2016). 2013 update on the US earnings, income, and wealth distributional facts: A view from macroeconomics. *Federal Reserve Bank of Minneapolis Quarterly Review*, 37(1):2–73.
- Long, J. B. D. and Summers, L. H. (1986). Is increased price flexibility stabilizing? *American Economic Review*, 76(5):1031–1044.
- Lucas, R. E. (1980). Two illustrations of the quantity theory of money. *American Economic Review*, 70(5):1005–1014.
- Maliar, L., Maliar, S., and Valli, F. (2010). Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm. *Journal of Economic Dynamics and Control*, 34(1):42–49.
- Maliar, L., Maliar, S., and Winant, P. (2021). Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122:76–101.
- McKay, A., Nakamura, E., and Steinsson, J. (2016). The power of forward guidance revisited. *American Economic Review*, 106(10):3133–58.
- McKay, A. and Reis, R. (2016). The role of automatic stabilizers in the US business cycle. *Econometrica*, 84(1):141–194.
- Nakov, A. (2008). Optimal and simple monetary policy rules with zero floor on the nominal interest rate. *International Journal of Central Banking*, 4:73–127.
- Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, 103(6):2530–53.

- Rotemberg, J. J. (1982). Sticky prices in the United States. *Journal of Political Economy*, 90(6):1187–1211.
- Rouwenhorst, G. K. (1995). Asset pricing implications of equilibrium business cycle models. In Cooley, T. F., editor, *Frontiers of Business Cycle Research*. Princeton University Press.
- Schaab, A. (2020). Micro and macro uncertainty. Technical report, Mimeo.
- Young, E. R. (2010). Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations. *Journal of Economic Dynamics and Control*, 34(1):36–41.

A Computational Algorithm

Our novel algorithm for solving a fully nonlinear HANK economy is based on the stochastic simulation algorithm in [Maliar et al. \(2010\)](#), except that the bond distribution is represented as a histogram following [Young \(2010\)](#). Switching to a recursive notation, we start with the Euler equation of the individual household problem:

$$\left(c - \chi \frac{h^{1+\nu}}{1+\nu}\right)^{-\sigma} - \mu = \beta \mathbb{E} \left(\frac{\xi'}{\xi} \frac{R}{\pi'} \left(c' - \chi \frac{h'^{1+\nu}}{1+\nu}\right)^{-\sigma} \right)$$

where μ is the Lagrange multiplier associated with the borrowing constraint of the households.

Using the budget constraint, we can rewrite this condition as:

$$\begin{aligned} \tilde{b}' &= \frac{R_{-1}}{\pi} b + wsh + \Pi s \\ &- \left[\mu + \beta \mathbb{E} \left(\frac{\frac{\xi' R}{\xi \pi'}}{\left(\frac{R}{\pi'} b' + w' s' h' - \tau (w' s' h') + \Pi' s' - b'(b') - \chi \frac{h'^{1+\nu}}{1+\nu}\right)^\sigma} \right) \right]^{-1/\sigma} \\ &- \chi \frac{h^{1+\nu}}{1+\nu} - T, \end{aligned} \tag{A.1}$$

where $\mu \equiv \mu(b, R_{-1}, s, \xi)$, $b' \equiv b'(b, R_{-1}, s, \xi)$, and $b'(b') \equiv b'(b'(b, R_{-1}, s, \xi))$.

We construct grids of points for $b \in [b_{min}, b_{max}]$ and $R \in [R_{min}, R_{max}]$, in addition to the grids for the idiosyncratic state $s \in \{s_m\}_{m=1}^M$ and the aggregate state $\xi \in \{\xi_j\}_{j=1}^J$. We use the polynomial rule from [Maliar et al. \(2010\)](#) for b to place more grid points near the borrowing constraint. The grid for R is evenly spaced, however. In economies with a ZLB, we can set $R_{min} = 1$.

The individual problem is then solved on this grid according to:

Algorithm 1 (Individual Problem)

1. Make a guess for the bond policy function $b'(b, R_{-1}, s, \xi)$ on the grid. We set the initial bond policy function to the DSS bond policy function for all aggregate states (R_{-1}, ξ) .
2. For each grid point (b, R_{-1}, s, ξ) , plug the assumed bond policy function $b'(b, R_{-1}, s, \xi)$ in the right-hand side of (A.1), set the Lagrange multiplier equal to zero, and compute the new bond policy function in the left-hand side of equation (A.1). The required labor policy is $h = \left(\frac{\tau}{\chi}(ws)\right)^{\frac{1}{\nu}}$. For each point that does not belong to $[b_{min}, b_{max}]$, set $\tilde{b}'(b, R_{-1}, s, \xi)$ equal to the corresponding boundary value.

3. Compute the bond function for the next iteration $\tilde{\tilde{b}}(b, R_{-1}, s, \xi)$ using:

$$\tilde{\tilde{b}}(b, R_{-1}, s, \xi) = \lambda_b \tilde{b}'(b, R_{-1}, s, \xi) + (1 - \lambda_b) b'(b, R_{-1}, s, \xi) \quad (\text{A.2})$$

where $\lambda_b \in (0, 1]$ is an updating parameter.

Iterate on steps 2 and 3 until the sup difference between $\tilde{\tilde{b}}(b, R_{-1}, s, \xi)$ and $\tilde{b}'(b, R_{-1}, s, \xi)$ is less than a given degree of precision.

As discussed in [Maliar et al. \(2010\)](#), the algorithm satisfies the Euler equation, the budget constraint, and the complementary slackness condition by construction.

When solving the individual problem, one also needs several ‘‘aggregate’’ variables, such as π_t , τ_t , w_t , and Π_t . We follow [Krusell and Smith \(1998\)](#) and [Fernández-Villaverde et al. \(2023\)](#) to build PLMs that predict $\log \pi_t$ and $\log \left(\frac{\pi_{t+1}}{\pi} \right) \frac{Y_{t+1}}{Y_t}$, given the aggregate states (R_{t-1}, ξ_t) using neural networks and moments of the household distribution. Given this, we can back out m_t from the Phillips curve and all other ‘‘aggregate’’ variables required for solving the individual problem.⁵

The initial guesses for our PLMs are such that initially inflation is at its target and the inflation expectation term is such that marginal costs are always at their DSS value, i.e., $m_t = \frac{\varepsilon - 1}{\varepsilon}$. Given this, the whole model can be solved as follows:

Algorithm 2 (Stochastic Simulation)

1. Generate and fix a time series of length T for the aggregate shocks.
2. Set initial matrices D_i ($i \in \{\pi, \mathbb{E}\pi\}$) for the PLMs and initial distribution of bonds. We use the distribution of bonds in the DSS and represent this distribution as a histogram as in [Young \(2010\)](#).⁶
3. Given the PLMs as represented by D_i , compute a solution to the individual problem as described in [Algorithm 1](#). Off-grid values of D_i are linearly interpolated.
4. Use the individual policy functions, the PLMs, and the aggregate shocks from step 1. to

⁵For comparison, we also found PLMs by computing linear regressions:

$$\begin{aligned} \log \pi_t &= \beta_{(0,\pi)} + \beta_{(1,\pi)} \log R_{t-1} + \beta_{(2,\pi)} \log \xi_t + \beta_{(3,\pi)} \log R_{t-1} \log \xi_t \\ \log \left(\frac{\pi_{t+1}}{\pi} \right) \frac{Y_{t+1}}{Y_t} &= \beta_{(0,\mathbb{E}\pi)} + \beta_{(1,\mathbb{E}\pi)} \log R_{t-1} + \beta_{(2,\mathbb{E}\pi)} \log \xi_t + \beta_{(3,\mathbb{E}\pi)} \log \xi_{t+1} \\ &\quad + \beta_{(4,\mathbb{E}\pi)} \log R_{t-1} \log \xi_t + \beta_{(5,\mathbb{E}\pi)} \log R_{t-1} \log \xi_{t+1} + \beta_{(6,\mathbb{E}\pi)} \log \xi_t \log \xi_{t+1}. \end{aligned}$$

These are similar to the specification in [Bayer et al. \(2019\)](#) but adapted to the fact that the aggregate state ξ_t follows an AR(1) process and not a Markov chain.

⁶ D_π and $D_{\mathbb{E}\pi}$ are matrices that represent the predictions of the neural network (evaluated on dense grids for R_{t-1} , ξ_t , and ξ_{t+1}). We do not use the prediction of the neural network directly to be able to update the PLMs slowly.

simulate the economy forward as in [Young \(2010\)](#) and calculate average bond holdings B_t . If $B_t \neq \tilde{B}$, use a nonlinear solver to find π_t that implies $B_t = \tilde{B}$. This requires updating individual policies where we take the PLM for inflation expectations and all $t + 1$ policies as given.

5. Train the neural networks on the simulated data (use the series of π_t and R_{t-1} that imply $B_t = \tilde{B}$ at each simulation step).
6. Evaluate the neural network on a dense grid of points. Let \tilde{D}_i be the resulting matrices.
7. Compute the PLMs for the next iteration using the updating formula:

$$\tilde{\tilde{D}}_i = \lambda_{PLM} \tilde{D}_i + (1 - \lambda_{PLM}) D_i \tag{A.3}$$

where $\lambda_{PLM} \in (0, 1]$ is an updating parameter.

Iterate on steps 3.-7. until the average squared difference between $\tilde{\tilde{D}}_i$ and D_i is less than a given degree of precision.