



BIS Working Papers


No 1160

Inequality and the Zero Lower Bound

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

Monetary and Economic Department

January 2024



JEL classification: D31, E12, E21, E31, E43, E52, E58

Keywords: heterogeneous agents, HANK models, neural networks, non-linear dynamics

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 2024. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Inequality and the Zero Lower Bound

Jesús Fernández-Villaverde
University of Pennsylvania
NBER, CEPR

Joël Marbet
Banco de España

Galo Nuño
Banco de España
BIS, CEMFI, CEPR

Omar Rachedi
ESADE Business School
Universitat Ramon Llull*

December 3, 2023

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.

Keywords: Heterogeneous agents, HANK models, neural networks, non-linear dynamics.

JEL codes: D31, E12, E21, E31, E43, E52, E58.

*We thank Guido Ascari, Olivier Blanchard, Anmol Bhandari, Edouard Challe, Jordi Galí, Benjamin Johannsen, Lilia Maliar, 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 BIS, Banco de España or the Eurosystem.

1 Introduction

This paper examines the interplay between household inequality and the zero lower bound (ZLB) of nominal interest rates and its 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 accommodate negative shocks to the economy effectively.

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 room for the central bank to stabilize the economy away from the ZLB. Hence, any factor that changes precautionary savings, such as higher earnings volatility, will affect the real rates and 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 earnings risk in the spirit of Bewley (1980), Huggett (1993), and Aiyagari (1994). Households can smooth consumption via borrowing. However, their credit capacity is limited by a borrowing constraint. Households also experience preference shocks, which are the source of aggregate uncertainty. As in Erceg et al. (2000), wages are sticky: the labor union sets nominal wages subject to adjustment costs à la Rotemberg (1982). The model also features a firm that produces using labor and a central bank that sets the nominal interest rate according to a Taylor rule subject to the ZLB constraint. Importantly, we explicitly account for the risk of hitting the ZLB in households' expectations, as in Fernández-Villaverde et al. (2015a).

We solve our model non-linearly by introducing a novel neural network algorithm proposed by Fernández-Villaverde et al. (2023) that allows us to compute the stochastic equilibrium dynamics of a ZLB-HANK economy instead of using standard computation methods

for HANK models based on linearization (e.g., [Ahn et al., 2018](#), and [Auclert et al., 2021](#)). We demonstrate how neural networks can successfully approximate the non-linear laws of motion of the economy’s states. 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 since 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 labor income, 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. Hence, 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 it would in the absence of the ZLB, and the central bank will undershoot its inflation target. Agents respond to this deflationary bias by increasing their precautionary savings, which lowers real rates: higher demand for savings

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 (also called the risky steady state, e.g., [Coourdacier et al., 2011](#)) is a point at which all aggregate variables are constant, and the realization of the aggregate demand shock is zero. However, agents are aware that non-zero realizations can come in the future. Furthermore, 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.

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](#), and [Bianchi et al., 2021](#)), we show that heterogeneity makes it more acute.²

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 last 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. When we lower the inflation target from 4% (the average inflation between 1980 and 1999) to 1.7% (the average inflation from 2000 onward), the decline in the real interest rate in the SSS of our model is 19 basis points (bps). In comparison, the real rate falls 16 bps in the ZLB-RANK. In other words, heterogeneity reduces the real rate by three additional bps. 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 others, [Lucas, 1980](#), or [King and Watson, 1997](#)).

However, 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 it occurs 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 28 bps, or 14% 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 decline in the real rate. 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](#)), and the fall in productivity growth. In summary, our model encompasses situations where the 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

²The natural rate of interest is also endogenous in [Eggertsson and Krugman \(2012\)](#) due to deleveraging.

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 is connected 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. From this perspective, our quantitative analysis complements the analytical results in [Bilbiie \(2021a\)](#), showing that heterogeneity amplifies liquidity-trap recessions.

Our work also aligns with the literature focused on solving the full stochastic dynamics of HANK economies.³ 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. [Lin and Peruffo \(2022\)](#) combine the sequence-state Jacobian methodology in [Auclert et al. \(2021\)](#) with a tractable two-regime shock structure for aggregate uncertainty. Methodologically, our paper builds on [Fernández-Villaverde et al. \(2023\)](#), who pioneered the use of neural networks to solve heterogeneous agent models globally. Our approach is also related to [Maliar et al. \(2021\)](#), [Azinovic et al. \(2022\)](#), and [Han et al. \(2022\)](#), who also employ deep learning to globally solve 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.

³Algorithms for non-linear heterogeneous agent economies based on perfect foresight, such as [Boehl \(2023\)](#), cannot address how heterogeneity impacts households' expectations on the risk of hitting the ZLB.

2 Model

We postulate an economy populated by heterogeneous households, a labor recruiter, a labor union subject to nominal wage rigidities, a firm that produces using labor, 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 earnings 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 [Eggertsson \(2003\)](#) 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 lifetime expected discounted utility:

$$\begin{aligned} \max_{\{c_{i,t}, b_{i,t}, h_{i,t}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \xi_t \left[u(c_{i,t}) - v(h_{i,t}) \right] \\ \text{s.t. } c_{i,t} + b_{i,t} + \tau_t = w_t s_{i,t} h_{i,t} + \frac{R_{t-1}}{\pi_t} b_{i,t-1}, \end{aligned} \quad (1)$$

$$b_{i,t} \geq \underline{b}, \quad (2)$$

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 the CRRA utility functions:

$$u(c_{i,t}) = \frac{c_{i,t}^{1-\sigma}}{1-\sigma}$$

and

$$v(h_{i,t}) = \chi \frac{h_{i,t}^{1+\nu}}{1+\nu}.$$

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 and lump-sum taxes, τ_t , with labor earnings, $w_t s_{i,t} h_{i,t}$, where w_t denotes the real wage. Households can 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 $\pi_t = \frac{P_t}{P_{t-1}}$, where P_t is the price of the final good. The borrowing constraint implies that the exogenous bound \underline{b} limits households' bond positions.

Labor market: The labor market features a representative perfectly competitive labor recruiter and a representative perfectly competitive labor union. The labor union purchases the different varieties of labor services from the households, and sells them to the labor recruiter. The nominal rigidities in the model arise because the labor union faces wage adjustment costs in the spirit of [Erceg et al. \(2000\)](#).

The labor recruiter purchases labor services $s_{i,t}h_{i,t}$ at the nominal wage $W_{i,t}$ from the labor union and bundles them into an aggregate effective labor input, H_t ,

$$H_t = \left[\int_0^1 s_{i,t} h_{i,t}^{\frac{\varepsilon-1}{\varepsilon}} di \right]^{\frac{\varepsilon}{\varepsilon-1}},$$

where ε is the elasticity of substitution across labor varieties. The aggregate labor input is then sold to the firm at the nominal wage W_t :

$$W_t = \left[\int_0^1 s_{i,t} W_{i,t}^{1-\varepsilon} di \right]^{\frac{1}{1-\varepsilon}}.$$

The optimal demand of labor services of household i from the labor recruiter is:

$$h_{i,t} = \left(\frac{W_{i,t}}{W_t} \right)^{-\varepsilon} H_t.$$

The labor union sets the nominal wage $W_{i,t}$ at which it sells each effective unit of labor $s_{i,t}h_{i,t}$, but it faces wage adjustment costs à la [Rotemberg \(1982\)](#). In addition, we follow [Hagedorn et al. \(2019\)](#) by assuming that the wage adjustment costs are virtual. While the adjustment costs affect labor unions' optimal pricing decisions, they do not result in any transfer of real resources (e.g., they are effort costs), simplifying the interpretation of the results as we change other model parameters. Accordingly, the labor union decides the nominal wage of each differentiated labor service by maximizing expected profits subject to the labor recruiter's optimal demand:

$$\begin{aligned} \max_{\{W_{i,k}\}_{k=t}^{\infty}} \mathbb{E}_t \sum_{k=t}^{\infty} \beta^k \left\{ \int_0^1 \left(\frac{s_{i,k} h_{i,k} W_{i,k}}{P_k} - \frac{v(h_{i,k})}{u'(C_k)} - \frac{\theta}{2} \left[\log \left(\frac{W_{i,k}}{W_{i,k-1}} \frac{1}{\tilde{\pi}} \right)^2 h_{i,k} \right] di \right) \right\}, \quad (3) \\ \text{s.t. } h_{i,t} = \left(\frac{W_{i,t}}{W_t} \right)^{-\varepsilon} H_t, \end{aligned}$$

where $u'(C_t)$ is the marginal utility of aggregate consumption and $\tilde{\pi}$ is the inflation rate

targeted by the central bank. We follow Bayer et al. (2019) in the specification of the adjustment costs, except that we extend their setting to the case of trend inflation, possibly different from zero.

The parameter θ controls wage stickiness. 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 wage dispersion, a mechanism unrelated to the main focus of our paper.

Appendix B shows that in a symmetric equilibrium in which the labor union sets the same nominal wage, $W_{i,t} = W_t$, and hires the same amount of working hours, $h_{i,t} = H_t$, across the differentiated labor services, the wage Phillips curve is

$$\log\left(\frac{\pi_t^w}{\tilde{\pi}}\right) = \beta \mathbb{E}_t \left[\log\left(\frac{\pi_{t+1}^w}{\tilde{\pi}}\right) \frac{H_{t+1}}{H_t} \right] + \frac{\varepsilon}{\theta} \left[\frac{v'(H_t)}{u'(C_t)} - \frac{\varepsilon - 1}{\varepsilon} w_t \right], \quad (4)$$

where $\pi_t^w = \frac{W_t}{W_{t-1}}$ is gross wage inflation, $v'(H_t)$ is the marginal disutility of working, and $w_t = \frac{W_t}{P_t}$ is the real wage.

Firm: A good is produced by a representative firm and sold in a perfectly competitive final good market. The firm manufactures the good Y_t using a linear technology in labor, $Y_t = L_t$. Since the real wage is constant, price inflation coincides with wage inflation, $\pi_t = \pi_t^w$.

Central bank: 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 $R_t = 1$.

Fiscal authority: The fiscal authority raises lump-sum taxes τ_t on the households to finance a fixed amount of outstanding debt \tilde{B} , such that the government budget constraint is satisfied $\tau_t = (r_t - 1)\tilde{B}$. Given \tilde{B} and the equilibrium interest rate, the lump-sum taxes τ_t are set to clear the budget constraint.

Market clearing: Since the definition of a 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 firm: $H_t = L_t$. 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 value added equals the total output of the firm and households' total consumption: $Y_t = L_t = \int_0^1 c_{it} di$.

3 Calibration

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

Table 1: Baseline Parameterization

Parameter		Value	Target/Source
Panel A. Preferences			
β	Discount factor	0.99577	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	0.8696	Labor supply equals 1 in the DSS
Panel B. Aggregate Risk			
ρ_ξ	AR coefficient of process for ξ	0.6	Bianchi et al. (2021)
ω_ξ	Standard deviation of ξ shock	0.01225	10% ZLB frequency
Panel C. 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	33% Borrowers
\underline{b}	Borrowing limit	-0.58	Two-month average labor income
Panel D. Labor Market			
ε	Elasticity across differentiated labor services	11	10% wage markup
θ	Rotemberg wage adjustment cost	100	12-month duration of wages
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 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 of changes in inflation to $\phi_\pi = 2.5$. We set $\beta = 0.99577$ 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 FOMC’s November 2018 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 $\omega_\xi = 0.01225$.

On the labor market, we set the elasticity across differentiated labor services to $\varepsilon = 11$, such that the wage markup is 10%. The wage adjustment parameter, $\theta = 100$, implies an equivalent Calvo-type protocol with a 12-month average wage duration, the conventional value used in the literature.

Regarding the idiosyncratic risk, we first set $\underline{b} = -0.58$, which equals about two months’ average wages. Then, we calibrate $\rho_s = 0.94$ and $\omega_s = 0.075$ so that the model reproduces a share of borrowers of 33% ([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. We then convert the AR(1) process for labor earnings into a three-point Markov chain using [Rouwenhorst’s \(1995\)](#) method.

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 $\chi = 0.8696$ such that the aggregate value of the efficiency units of hours equals one in the DSS.

Regarding the fiscal authority, we follow the one-asset calibration strategy of [McKay et al. \(2016\)](#) and fix $\tilde{B} = 25\%$ of annual GDP, in line with the estimate of liquid wealth in the U.S. economy derived by [Kaplan et al. \(2018\)](#).⁴

⁴[McKay et al. \(2016\)](#) compute that aggregate liquid assets equal 140% of annual GDP, since they consider net liquid assets as the sum of all deposits in financial institutions, government bonds, corporate bonds net of revolving consumer credit, and corporate equity. Instead, our target is based on [Kaplan et al. \(2018\)](#), who categorize corporate equity as illiquid assets (corporate equity equals 102% of GDP). As discussed in [Kaplan et al. \(2018\)](#), a larger amount of liquid assets in the economy reduces households’ marginal propensity to consume, downplaying the role of households’ heterogeneity in driving aggregate dynamics.

All in all, our calibration choices make our economy consistent with both the average amount of household wealth inequality and the cyclical nature of income inequality in the U.S. economy. First, the Gini index of wealth in the model is 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. Second, the income of wealth-poor (and high-MPC) households is more pro-cyclical than that of wealth-rich (and low-MPC) households. Specifically, while the total income for the bottom 90th of the wealth distribution has a correlation of 0.94 with GDP, the analogous correlation for the total income of the top 10th is just 0.40. This implies that high-MPC households become relatively poorer in recessions, which is consistent with the empirical evidence of [Patterson \(2023\)](#).⁵

Crucial to this result is the fact that the source of nominal rigidity is wage stickiness rather than price stickiness. In this way, there are no profits to be redistributed across households, which tends to substantially curtail the counter-cyclical nature of income inequality ([Bilbiie et al., 2022](#)).⁶ As emphasized by [Bilbiie \(2008, 2020, 2021b\)](#) and [Auclert \(2019\)](#), both the level and the cyclical nature of inequality in HANK economies are key to assessing how household heterogeneity shapes macroeconomic dynamics.

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 wage rigidities, heterogeneous agents, and a ZLB.

Solving for the competitive equilibrium of a HANK model is conceptually similar to estimating a DGP in econometrics. In both situations, we look for a function that maps variables into endogenous objects of interest. In our case, the variables are the states of the model and the endogenous objects of interest are the moments of the distribution of asset holdings by the households. The key difference between a standard econometric estimation and solving our model is that, in econometrics, the variables and the object of interest are empirical observations (e.g., education and wages from a micro panel), while, for us, they are the outcomes of the decisions of the agents *within* the model. However, formally, the problem is the same. In both situations, we face a functional approximation problem (e.g., finding the expectation of wages conditional on education vs. finding inflation as a function of the states of the model).

⁵For further evidence on the counter-cyclical nature of income inequality conditional on demand shocks, see [Coibion et al. \(2017\)](#) and [Mumtaz and Theophilopoulou \(2017\)](#).

⁶[Auclert et al. \(2023\)](#) show how nominal wage stickiness allows HANK models to be simultaneously consistent with estimates of MPCs, the marginal propensity to earn, and the fiscal multiplier.

Fernández-Villaverde et al. (2023) have argued that neural networks are an effective way to tackle this functional approximation problem. Neural networks satisfy two important theoretical conditions. First, they are universal non-linear approximators: neural networks can approximate arbitrarily well any Borel-measurable finite-dimensional function, thus encompassing a large class of DGPs. Second, neural networks break the curse of dimensionality of high-dimensional functional approximations by switching from the original space of the problem to a more convenient geometrical representation. These two properties make neural networks superior to other conventional approximations in economics (such as Chebyshev polynomials or non-parametric regressions) in many applications of interest. Moreover, neural networks are easy to code thanks to the availability of state-of-the-art open-source libraries, particularly suitable to parallel processing, and require less domain knowledge (i.e., a priori understanding of the salient numerical features of the function to approximate) than other alternatives.

Neural networks are particularly useful for HANK models with the ZLB because they can capture the full non-linear dynamics of the economy without having to select an ex-ante structure for the perceived law of motion (PLM) of the distribution of agents. This selection is a challenging task when the ZLB is present because the saving decisions become much more non-linear on aggregate shocks.

More concretely, we build perceived laws of motion (PLM) that predict $\log \pi_t^w$ and $\log \left(\frac{\pi_{t+1}^w}{\pi_t^w} \right) \frac{H_{t+1}}{H_t}$, given the aggregate states (R_{t-1}, ξ_t) using neural networks and moments of the household distribution in the spirit of Krusell and Smith (1998).⁷ Our solution method is extremely accurate when households keep track of 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.

In Appendix A, we provide a detailed description of our algorithm, but some important results are highlighted here. 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).

⁷The word “perceived” captures the idea that these are the law of motions actually used by the agents.

Figure 1: Non-Linearity due to the ZLB.

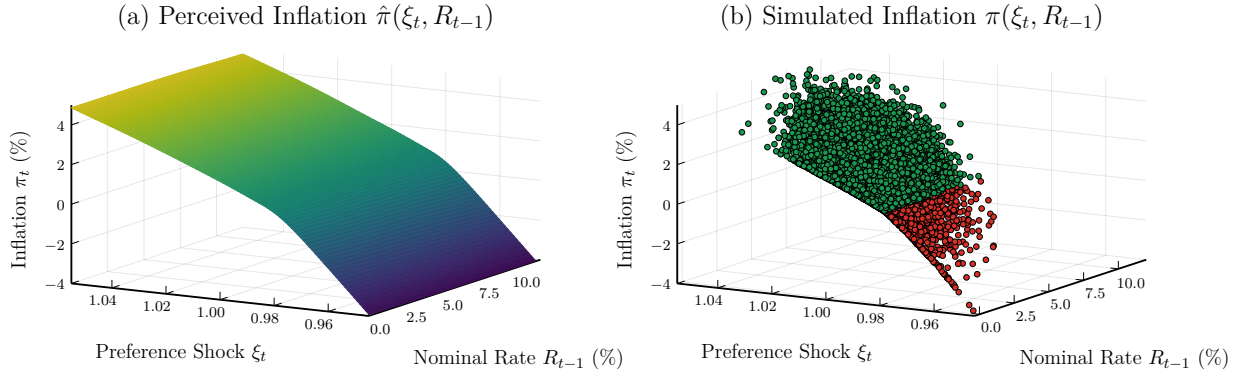
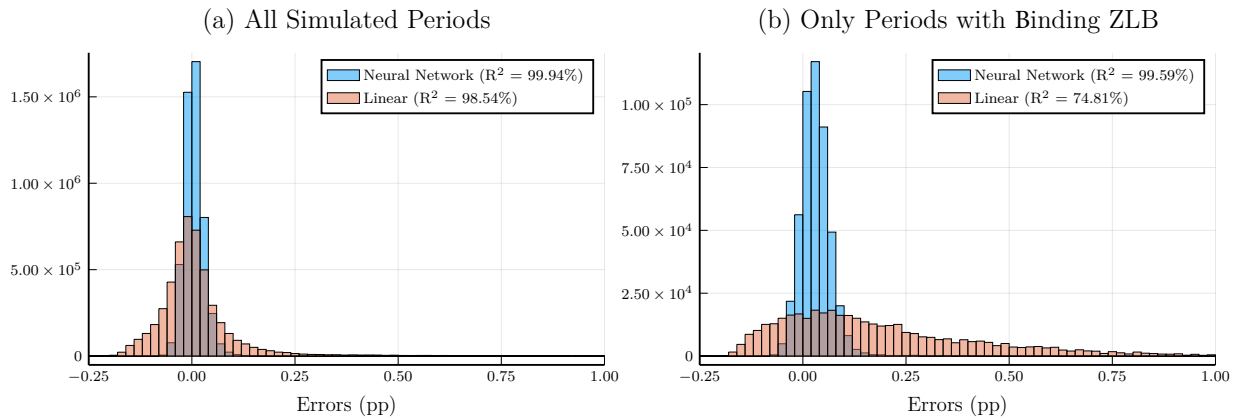


Figure 2 compares the nowcast errors on the inflation dynamics generated by our neural network approach with those implied by a naive application of the [Krusell and Smith \(1998\)](#) method, in which we predict the inflation rate as a log-linear function of the state variables (see Appendix A for the 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 98.54% to 99.94%. 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 74.81%, the neural network algorithm yields an R^2 of 99.59%, a much more satisfactory fit.

Figure 2: Nowcast Errors for Inflation.



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 parameterization as in Section 3 except that the ZLB is absent from the Taylor rule that the central bank follows.

Figure 3: Ergodic Distributions.

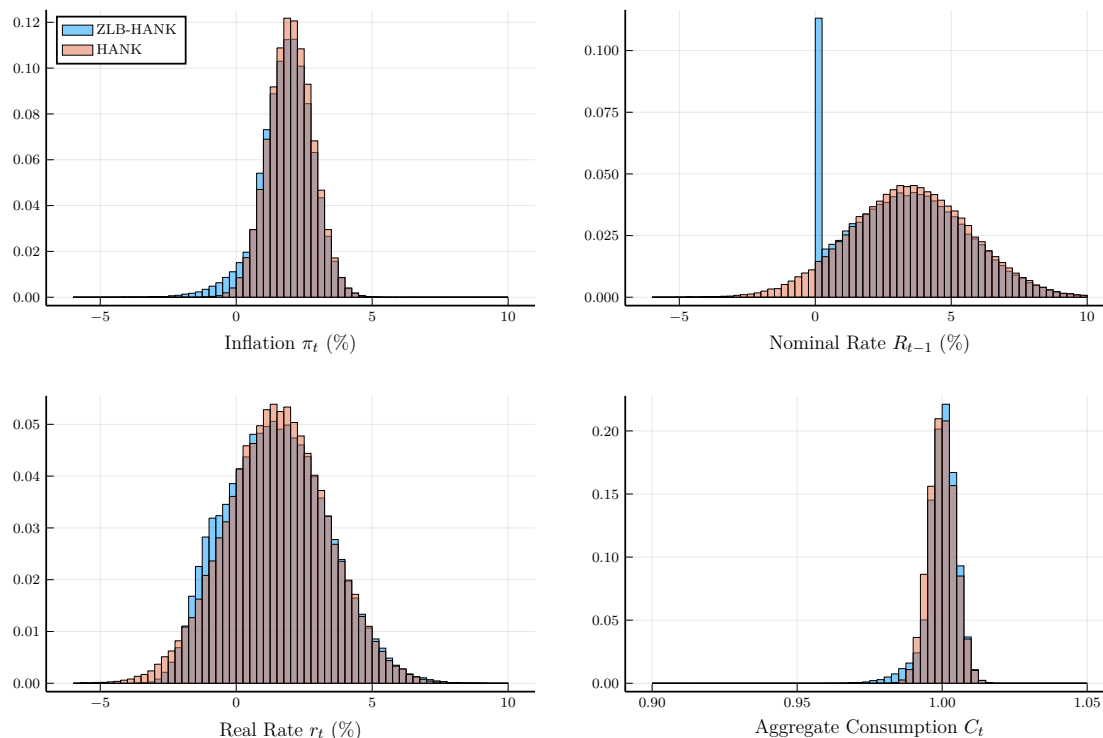


Figure 3 reports the ergodic distribution of inflation, the nominal interest rate, the real

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: Impulse Response Functions.

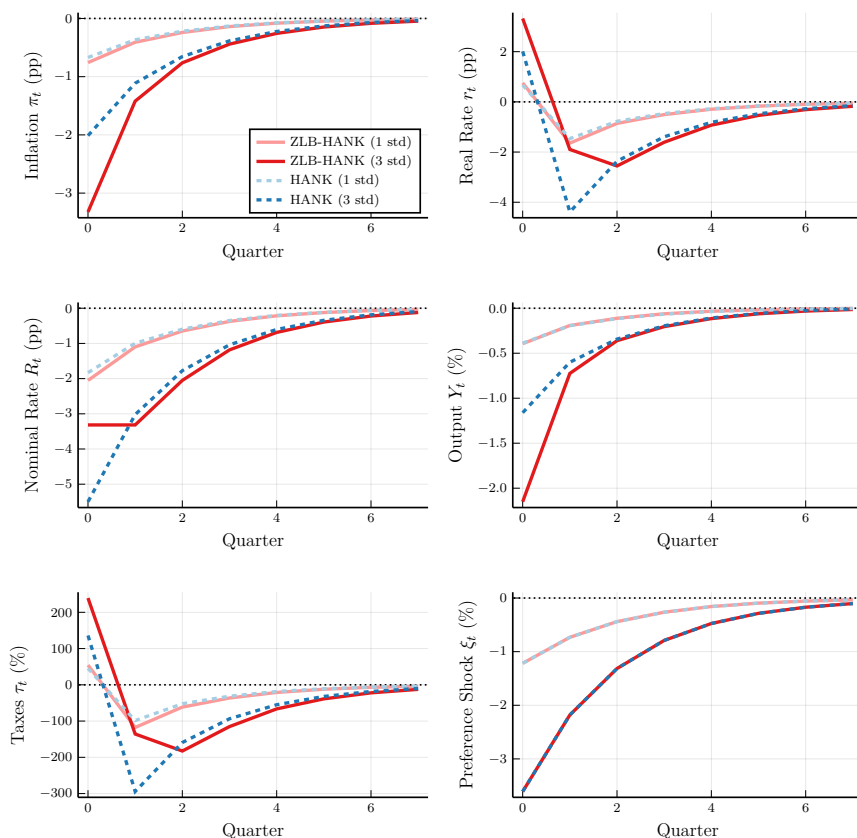


Figure 4 evaluates the effect of the non-linearity of the ZLB on the dynamics of inflation and output by plotting the impulse response functions (IRFs) at the SSS to small and large demand shocks in the ZLB-HANK and the HANK economies.⁸ Specifically, we consider a small demand shock (a one-standard-deviation 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.⁹ However, when the shock is large, the

⁸Our model is non-linear. Hence, the IRFs are state-dependent and the size and sign of the shock matter for their shape.

⁹Agents 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 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 -1.5% .¹⁰ 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. While real wages are constant by construction, the real interest rate increases upon a shock, especially when at the ZLB (as we have deflation, but the nominal rate is stuck at zero). Since interest payments on government debt are financed via taxation, lump-sum taxes also rise upon a shock, especially at the ZLB. We discuss next how these changes in real prices have important redistributive consequences.

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 interest earnings, labor income, and taxes for households in the 10th and 99th percentile of the wealth distribution, across the three realizations of the labor shock (ordered from the lowest to the highest realization). We also report the aggregate relative response of income.¹¹


Figure 5 provides two valuable insights. First, in the ZLB-HANK economy, aggregate income falls more sharply than in the HANK economy on average. This is a direct consequence of the central bank’s inability to accommodate the negative demand shock fully. Second, the ZLB constraint leads to a relative decline in total income for wealth-poor households compared to the counterfactual HANK without ZLB but to a relative increase for wealth-rich individuals. This is because working hours experience a sharp decline. Poorer households have either low or negative bond holdings (in our calibration, households at the 10th per-

ZLB. We will revisit this point below.

¹⁰The initial point is the SSS of each model. In the ZLB-HANK, it is 3.32%, and this is why the impulse response is bounded at -3.32 pp. In the HANK, it is 3.48%, and this is why a decline of more than 5 pp implies a fall below -1.5% .

¹¹Appendix C reports separately the responses of individual income in the ZLB-HANK and HANK economies.

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



Note: The graph reports the income response at impact to a three-standard-deviation negative demand shock for each of the three labor-earnings 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. For each group, the bars are ordered from the lowest realization of the labor-earnings shock up to the highest one.


centile have negative bond holdings) and rely more on labor income. 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, borrowers face relatively higher real interest payments.¹²

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. More concretely, the presence of the ZLB amplifies the drop in consumption of low-income wealth-poor individuals by about 1.4 pp, roughly twice the drop for wealthy households.

This drop is due to two different effects. For wealthy households, the relative consumption drop is driven by the intertemporal substitution effect, as households save relatively more

¹²This result implies that the income of wealth-poor high-MPC households is much more cyclical than that of wealth-rich low-MPC households in the ZLB-HANK economy relative to the standard HANK model.

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



Note: The graph reports the consumption response at impact to a three-standard-deviation negative demand shock for each of the three labor-earnings realizations at the 10th wealth percentile and the 99th wealth percentile. For each group, the bars are ordered from the lowest realization of the labor-earnings shock up to the highest one.

due to the higher path for real rates in the ZLB-HANK compared to the HANK. This explains why the drop is similar across labor income levels. For poor households, which lack the assets to smooth consumption (8% of households are at the borrowing constraint at the SSS), the income effect dominates. As low-income “hand-to-mouth” agents display much larger MPCs, this explains why the relative drop in consumption is much larger for these households compared to high-income “hand-to-mouth” ones.

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 low-income 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 sizable 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, ZLB-RANK, and RANK.

Variable	ZLB-HANK		HANK		ZLB-RANK		RANK	
	DSS	SSS	DSS	SSS	DSS	SSS	DSS	SSS
Inflation	2.00%	1.92%	2.00%	1.99%	2.00%	1.93%	2.00%	1.99%
Nominal Rate	3.50%	3.32%	3.50%	3.48%	3.69%	3.52%	3.69%	3.68%
Real Rate	1.50%	1.40%	1.50%	1.48%	1.69%	1.60%	1.69%	1.68%
(Shadow) ZLB Frequency	-	9.62%	-	(5.70%)	-	8.90%	-	(5.37%)

Table 2 reports the results of our comparison exercise, showing the DSS and SSS values for inflation, the nominal interest rate, and the real interest rate, 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 the 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 one bp, two bps, and two bps, respectively. In contrast, introducing the ZLB into the ZLB-HANK economy reduces the SSS values for inflation, the nominal rate, and the real rate by approximately 8 bps, 18 bps, and 10 bps, respectively.

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 9.62% 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 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 5.70% 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 that between ZLB-HANK and HANK. For instance, the frequency of ZLB spells goes from a (shadow) 5.37% in RANK to 8.90% in ZLB-RANK (a difference of 3.53%) while it goes from a (shadow) 5.70% in HANK to 9.62% in ZLB-HANK (a difference of 3.92%). 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	1.69%	3.69%	2.00%
ZLB-HANK SSS	1.40%	3.32%	1.92%
(i) Total	0.29pp	0.37pp	0.08pp
ZLB-RANK DSS	1.69%	3.69%	2.00%
ZLB-HANK DSS	1.50%	3.50%	2.00%
(ii) Precautionary Savings Idiosyncratic Risk	0.19pp	0.19pp	0.00pp
(i)-(ii) Deflationary Bias	0.10pp	0.18pp	0.08pp

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, 1.69%, and the ZLB-HANK SSS, 1.40%, which is 29 bps lower in the latter economy (about a fifth 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 :

$$1.69\% = \left(\frac{1}{\beta = 0.995477} \right)^4.$$

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 0.995477. 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 19 of the 29 bps of the total difference.

Then, we can return to the difference between the ZLB-HANK DSS and SSS from Table 3, which is 10 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 wage decisions of labor unions. 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. It also makes it more likely that a negative demand shock will push the economy to the ZLB and trigger a deflationary spiral. Similarly, labor unions bias their wage changes downward to avoid paying large wage 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 labor unions change their wage decisions. See the columns for ZLB-HANK and HANK in Table 2, which confirm previous findings by Adam and Billi (2007), Nakov (2008), 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 9.62%, while this frequency is 8.90% 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.

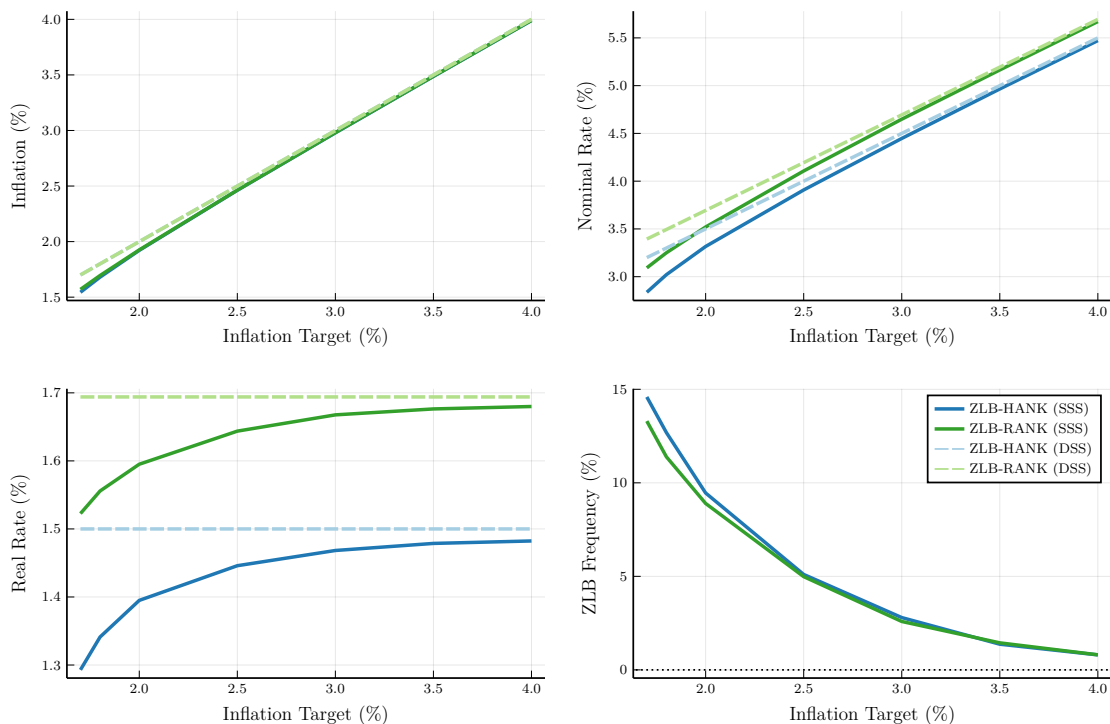
More pointedly, our analysis shows that the real rate in the ZLB-HANK economy is en-

dogenously 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 levels 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 1.54%.¹³ This 16 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 13% when the inflation target is 1.7%.

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



As we vary the inflation target, the real interest rate at the DSS does not change (dashed

¹³This 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.

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. 19 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 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 earnings

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 28 bps (in ZLB-RANK, since there is no heterogeneity, the drop is still the 16 bps we reported in the previous subsection).¹⁴

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 18% 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. Importantly, insofar as trends in population aging and inequality are key drivers of real rates of advanced economies and are expected to accelerate in the post-COVID economy, the channels emphasized in this paper in curtailing real rates and boosting the probability of hitting the ZLB could be even more relevant once the current inflationary pressures abate.

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 labor unions biasing their wage changes downward to avoid being forced to pay high wage adjustment costs at the ZLB.

¹⁴Since the bulk of our mechanism depends on the way in which the risk of hitting the ZLB alters households' desire to accumulate precautionary savings, a TANK economy is ill-equipped to quantitatively account for how changes in the level of the inflation target and wealth inequality impact real rates.

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



To disentangle both mechanisms, we increase the value of θ , the Rotemberg wage adjustment cost, from 100 (equivalent to a 12-month average wage duration) to 200 (equivalent to about a 24-month average wage duration). This way, labor unions 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 10 bps to 2 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 wage rigidities make the deflationary bias smaller by stabilizing the economy (except in the limit case of full price flexibility when the ZLB is inconsequential), thus confirming the primary role of households' precautionary savings in driving the deflationary bias.

6 Conclusion

In this paper, we have shown 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 findings, we make five 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 pp 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 28 bps. Although other factors, such as the aging of the 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., and Rognlie, M. (2023). MPCs, MPEs, and multipliers: A trilemma for New Keynesian models. *Review of Economics and Statistics*, 105(3):700–712.
- 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. *Mimeo*.
- Azinovic, M., Gaegauf, L., and Scheidegger, S. (2022). Deep equilibrium nets. *International Economic Review*, 63(4):1471–1525.
- Bayer, C., Lüttinge, 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. *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.

- Bilbiie, F. (2008). Limited asset markets participation, monetary policy and (inverted) aggregate demand logic. *Journal of Economic Theory*, 140(1):162–196.
- Bilbiie, F. (2020). The New Keynesian cross. *Journal of Monetary Economics*, 114:90–108.
- Bilbiie, F. (2021a). HALT: Heterogeneous-agent liquidity traps. *Mimeo*.
- Bilbiie, F. (2021b). Monetary policy and heterogeneity: An analytical framework. *Mimeo*.
- Bilbiie, F. O., Känzig, D. R., and Surico, P. (2022). Capital and income inequality: An aggregate-demand complementarity. *Journal of Monetary Economics*, 126:154–169.
- Boehl, G. (2023). HANK on speed: Robust nonlinear solutions using automatic differentiation. *Mimeo*.
- 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.
- 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., Kueng, L., and Silvia, J. (2017). Innocent bystanders? Monetary policy and inequality. *Journal of Monetary Economics*, 88:70–89.
- 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. (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.

- Erceg, C. J., Henderson, D. W., and Levin, A. T. (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics*, 46(2):281–313.
- 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. *Mimeo*.
- Gorodnichenko, Y., Maliar, L., Maliar, S., and Naubert, C. (2021). Household savings and monetary policy under individual and aggregate stochastic volatility. *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. *Mimeo*.
- Han, J., Yang, Y., and E, W. (2022). DeepHAM: A global solution method for heterogeneous agent models with aggregate shocks. *Mimeo*.
- Hillenbrand, S. (2021). The Fed and the secular decline in interest rates. *Mimeo*.
- 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. *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.

- 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.
- Lin, A. and Peruffo, M. (2022). Aggregate uncertainty, HANK, and the ZLB. *Mimeo*.
- 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.
- Mumtaz, H. and Theophilopoulou, A. (2017). The impact of monetary policy on inequality in the UK. An empirical analysis. *European Economic Review*, 98:410–423.
- 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.
- Patterson, C. (2023). The matching multiplier and the amplification of recessions. *American Economic Review*, 113(4):982–1012.
- 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. *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:

$$c^{-\sigma} - \mu = \beta \mathbb{E} \left[\frac{\xi'}{\xi} \frac{R}{\pi'} c'^{-\sigma} \right],$$

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

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

$$\tilde{b}' = \frac{R_{-1}}{\pi} b + wsh - \left[\mu + \beta \mathbb{E} \left(\frac{\frac{\xi' R}{\xi \pi'}}{\left(\frac{R}{\pi'} b' + w's'h' - \tau' - b'(b') \right)^\sigma} \right) \right]^{-1/\sigma}, \quad (\text{A.1})$$

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 equation (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 $h = H$ is arising from the labor unions' problem. 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 $\tilde{b}(b, R_{-1}, s, \xi)$ 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^w , H_t , and τ_t . We follow [Krusell and Smith \(1998\)](#) and [Fernández-Villaverde et al. \(2023\)](#) to build PLMs that predict $\log \pi_t^w$ and $\log \left(\frac{\pi_{t+1}^w}{\pi_t^w} \right) \frac{H_{t+1}}{H_t}$, given the aggregate states (R_{t-1}, ξ_t) using neural networks and moments of the household distribution. Given this, we can back out H_t from the wage inflation equation and all other “aggregate” variables required for solving the individual problem.¹⁵

The initial guesses for our PLMs are such that initially inflation is at its target and the inflation expectation term is such that the implied marginal costs in the labor union’s problem are always at their DSS value. 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 ($i \in \{\pi^w, \mathbb{E}\pi^w\}$) 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 \(2010\)](#).¹⁶
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 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^w 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.

¹⁵For comparison, we also found PLMs by computing linear regressions:

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

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.

¹⁶ 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 so that we can update the PLMs slowly.

5. Train the neural networks on the simulated data.
6. Evaluate the neural network on a dense grid of points. Let \widetilde{D}_i be the resulting matrices.
7. Compute the PLMs for the next iteration using the updating formula:

$$\widetilde{\widetilde{D}}_i = \lambda_{PLM} \widetilde{D}_i + (1 - \lambda_{PLM}) D_i$$

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

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

B The Wage Phillips Curve

This section provides additional details on the derivation of the wage Phillips curve in equation (4) of Section 2.

Let us start by plugging the optimal labor demand of the labor recruiter for each differentiated labor service, $h_{i,t} = \left(\frac{W_{i,t}}{W_t}\right)^{-\varepsilon} H_t$, into equation (3), which yields the following labor union's maximization problem:

$$\max_{\{W_{i,k}\}_{k=t}^{\infty}} \mathbb{E}_t \sum_{k=t}^{\infty} \beta^k \left\{ \int_0^1 \left(\frac{s_{i,k} W_{i,k}}{P_k} \left(\frac{W_{i,k}}{W_k} \right)^{-\varepsilon} H_k - \frac{v \left(\left(\frac{W_{i,k}}{W_k} \right)^{-\varepsilon} H_k \right)}{u'(C_k)} \right. \right. \\ \left. \left. - \frac{\theta}{2} \left[\log \left(\frac{W_{i,k}}{W_{i,k-1}} \frac{1}{\widetilde{\pi}} \right)^2 \left(\frac{W_{i,k}}{W_k} \right)^{-\varepsilon} H_k \right] di \right) \right\}. \quad (\text{B.1})$$

Since $\int_0^1 s_{i,t} di = 1$, in a symmetric equilibrium in which $W_{i,t} = W_t$ for all $i \in [0, 1]$, the maximization problem can be written as:

$$\max_{\{W_k\}_{k=t}^{\infty}} \mathbb{E}_t \sum_{k=t}^{\infty} \beta^k \left\{ \left(\frac{W_k}{P_k} H_k - \frac{v(H_k)}{u'(C_k)} - \frac{\theta}{2} \left[\log \left(\frac{W_k}{W_{k-1}} \frac{1}{\widetilde{\pi}} \right)^2 H_k \right] \right) \right\}. \quad (\text{B.2})$$

The first-order condition of this problem is:

$$(1 - \varepsilon) \frac{1}{P_t} \left(\frac{W_{i,t}}{W_t} \right)^{-\varepsilon} H_t + \frac{v'(H_t)}{u'(C_t)} \varepsilon \left(\frac{W_{i,t}}{W_t} \right)^{-\varepsilon-1} \frac{H_t}{W_t} \\ - \theta \log \left(\frac{W_{i,t}}{W_{i,t-1}} \frac{1}{\widetilde{\pi}} \right) \frac{H_t}{W_{i,t}} + \beta \theta \mathbb{E}_t \left[\log \left(\frac{W_{i,t+1}}{W_{i,t}} \frac{1}{\widetilde{\pi}} \right) \frac{H_{t+1}}{W_{i,t}} \right] = 0. \quad (\text{B.3})$$

In a symmetric equilibrium in which $W_{i,t} = W_t$ for all $i \in [0, 1]$, the first-order condition

reads:

$$(1 - \varepsilon) \frac{1}{P_t} H_t + \frac{v'(H_t)}{u'(C_t)} \varepsilon \frac{H_t}{W_t} - \theta \log \left(\frac{\pi_t^w}{\tilde{\pi}} \right) \frac{H_t}{W_t} + \beta \theta \mathbb{E}_t \left[\log \left(\frac{\pi_{t+1}^w}{\tilde{\pi}} \right) \frac{H_{t+1}}{W_t} \right] = 0. \quad (\text{B.4})$$

Rearranging equation (B.4) yields the wage Phillips curve in equation (4).


C More on Individual Income Responses

Figure 5 in Section 5.2 reports the response of individual income in the ZLB-HANK economy expressed in relative terms with respect to those implied by the HANK economy to a three-standard-deviation negative demand shock that hits either economy at its SSS. The figure indicates that while aggregate income falls more sharply in the ZLB-HANK economy than in the HANK model on average, there is substantial heterogeneity across the wealth distribution: the individual income of the wealth-poor drops relatively more in the ZLB-HANK economy, whereas the wealth-rich are relatively better off, since they gain from the relative surge in real rates at the ZLB, which boosts the return on savings in bonds.

Here, we report the individual income responses to a three-standard-deviation negative demand shock separately for the HANK economy and the ZLB-HANK economy. Figure C.1 displays the response of total income at impact in the HANK economy (represented by diamonds) and decomposes them into contributions from changes in interest earnings, labor income, and taxes for households in the 10th and 99th percentile of the wealth distribution, across the three realizations of the labor shock (ordered from the lowest to the highest realization), as well as the aggregate response of income. Figure C.2 displays similar moments as implied by the ZLB-HANK economy.


Consistent with the relative responses of Figure 5, we find that aggregate income drops by roughly 3 pp in the ZLB-HANK economy, which almost doubles the reduction in aggregate income experienced in the HANK economy. This result masks substantial heterogeneity across the wealth distribution. While the income of wealth-poor households drops relatively more in the ZLB-HANK economy due to the relatively larger reduction in labor income, the income of wealth-poor households surges relatively more in the ZLB-HANK economy as the hikes in interest income more than offset the drop in labor income.

Figure C.1: Individual Income Responses - HANK.



Note: The graph reports the income response in the HANK economy for each of the three labor-earnings 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. For each group, the bars are ordered from the lowest realization of the labor-earnings shock up to the highest one.

Figure C.2: Individual Income Responses - ZLB-HANK.



Note: The graph reports the income response in the ZLB-HANK economy for each of the three labor-earnings 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. For each group, the bars are ordered from the lowest realization of the labor-earnings shock up to the highest one.

Previous volumes in this series

1159 January	The 'plucking' model of the unemployment rate floor: Cross-country estimates and empirics	Jing Lian Suah
1158 January	Financial development and the effectiveness of macroprudential and capital flow management measures	Yusuf Soner Başkaya, Ilhyock Shim, Philip Turner
1157 December	Fintech vs bank credit: How do they react to monetary policy?	Giulio Cornelli, Fiorella De Fiore, Leonardo Gambacorta and Cristina Manea
1156 December	Monetary policy frameworks away from the ELB	Fiorella De Fiore, Benoit Mojon, Daniel Rees, Damiano Sandri
1155 December	Monetary Tightening, Inflation Drivers and Financial Stress	Frederic Boissay, Fabrice Collard, Cristina Manea, Adam Shapiro
1154 December	Foreign investor feedback trading in an emerging financial market	Ingomar Krohn, Vladyslav Sushko, Witit Syntayakul
1153 December	Foreign institutional investors, monetary policy, and reaching for yield	Ahmed Ahmed, Boris Hofmann, Martin Schmitz
1152 December	The Heterogeneous Impact of Inflation on Households' Balance Sheets	Clodomiro Ferreira, José Miguel Leiva, Galo Nuño, Álvaro Ortiz, Tomasa Rodrigo, and Sirenia Vazquez
1151 November	The financial origins of regional inequality	Anne Beck and Sebastian Doerr
1150 November	Markups and the asymmetric pass-through of cost push shocks	Enisse Kharroubi, Renée Spigt, Deniz Igan, Koji Takahashi, Egon Zakrajšek
1149 November	Housing affordability: a new data set	Nina Biljanovska, Chenxu Fu, Deniz Igan
1148 November	Firm heterogeneity, capital misallocation and optimal monetary policy	Beatriz González, Galo Nuño, Dominik Thaler, Silvia Albrizio
1147 November	Central Bank Digital Currency and Privacy: A Randomized Survey Experiment	Syngjoo Choi, Bongseob Kim, Young-Sik Kim, Ohik Kwon
1146 November	On par: A Money View of stablecoins	Iñaki Aldasoro, Perry Mehrling, Daniel H. Neilson

All volumes are available on our website www.bis.org.