

Inequality and the Zero Lower Bound*

Jesús Fernández-Villaverde

Joël Marbet

University of Pennsylvania, NBER, CEPR

CEMFI

Galo Nuño

Omar Rachedi

Banco de España

ESADE Business School

March 26, 2021

Abstract

This paper argues that the effects of the zero lower bound (ZLB) on aggregate dynamics crucially depend on household inequality. We establish this result within a heterogeneous agent New Keynesian (HANK) model that features an occasionally-binding ZLB. Importantly, we solve numerically for the fully non-linear stochastic equilibrium using a novel neural-network algorithm. We first highlight how the presence of the ZLB in a HANK economy alters the response of households' decisions and macroeconomic aggregates to demand shocks. We then show that the interaction of the central bank's inflation target and the amount of wealth inequality is a key driver of the level of real interest rates and the frequency of ZLB events. This stands in contrast of standard macroeconomic models, in which the level of real rates is pinned down as an exogenous parameter. In our setting, a drop in the inflation target reduces the level of the real interest rate because households increase their precautionary savings against the higher risk of ZLB events. As a result, ZLB events become even more likely. This channel is further amplified at higher levels of wealth inequality.

Keywords: Zero Lower Bound, Heterogeneous Agents, HANK, Non-linear Dynamics, Real Interest Rates, Inequality.

JEL Classification Codes: D31, E12, E21, E31, E43, E52, E58.

*The views expressed in this paper are those of the authors and do not necessarily represent the views of the Banco de España or the Eurosystem.

1 Introduction

The level of real interest rates in advanced economies is at a historic low. While three decades ago the natural interest rate was around 3%, nowadays the level is estimated to be well below 1% (Del Negro et al., 2017; Fiorentini et al., 2018). This fact has major implications for the conduct of monetary policy (Barsky, Justiniano and Melosi, 2014), as it implies a much higher likelihood of experiencing situations in which policy rates are constrained by the zero lower bound (ZLB). This explains why in the aftermath of the Great Recession virtually all central banks have been constrained by the lower bound on policy rates for several years.¹ This situation has been further exacerbated by the coronavirus crisis, such that nowadays the policy rate is expected to be stuck around zero for about three years (Lilley and Rogoff, 2020).

Although interest rates play a major role in shaping the stance of monetary policy, standard macroeconomic models tend to feature a long-run neutral monetary policy, with the implication that the level of interest rates is pinned down by structural parameters unrelated to (monetary and fiscal) policy. In this paper we argue that the level of real interest rates crucially depends on the interaction between the stance of monetary policy – i.e., the central bank’s inflation target – and the amount of wealth inequality in the economy.

To establish this result, we extend the workhorse representative-agent New Keynesian (RANK) model in two ways: (*i*) by allowing for household heterogeneity due to uninsurable idiosyncratic labor income risk, in the spirit of Bewley (1980), Huggett (1993), and Aiyagari (1994), and (*ii*) by explicitly incorporating the risk of hitting the ZLB in households’ expectations, as in Fernández-Villaverde et al. (2015). This framework generalizes the heterogeneous agent New Keynesian (HANK) model of Gornemann, Kuester and Nakajima (2016), McKay, Nakamura and Steinsson (2016), and Kaplan, Moll and Violante (2018) by incorporating in the dynamics of the model the non-linearity implied by the existence of an occasionally-binding ZLB. To do so, we propose a novel neural-network algorithm that allows for the computation

¹In few jurisdictions the policy rates have actually been set to negative values. However, policy rates are constrained by the effective lower bound (ELB), that is, a level of the interest rates under which the central bank cannot implement further rate reductions. Throughout this paper, we consider the ZLB as the ELB, but our results can be generalized simply by taking a stand of how negative policy rates can be.

of the non-linear stochastic equilibrium of a HANK economy. Importantly, we also highlight how the presence of the ZLB substantially alters the dynamics of both households' decision and macroeconomic aggregates with respect to an otherwise identical HANK model that does not explicitly incorporate the non-linearity in the policy rate.

Our model features a continuum of ex-ante homogeneous households, which are ex-post heterogeneous due to an uninsurable and idiosyncratic labor earning risk. Households can smooth the effect of labor earning risk on their consumption via borrowing. However, households' credit capacity is limited by an ad-hoc borrowing limit. On the production side, there is a continuum of intermediate-good firms that produce using labor, and set prices subject to Rotemberg (1982) adjustment costs. Finally, the model features a monetary authority that sets the nominal interest rate with a Taylor rule, subject to the explicit constraint of the ZLB.

The model features preferences shocks as a source of aggregate uncertainty, as in Christiano, Eichenbaum and Rebelo (2011). This shock is meant to be a reduced-form for any variation in uncertainty, fiscal policy, or credit market tightness that could alter households' risk appetite. We use this exogenous variation in preferences following the extensive literature that has shown how changes in the degree of households' patience can be a main driver of a ZLB event. Together, the aggregate preference shock, the idiosyncratic labor earning risk, and the presence of a ZLB generate a non-linear HANK economy. We discipline the model following the a one-asset calibration of McKay, Nakamura and Steinsson (2016) and Kaplan, Moll and Violante (2018). In this way, our economy is consistent with the heterogeneity in labor earnings and liquid wealth across households, as well as generating a realistic value – in the lower end of the ballpark of the available estimates – for the average marginal propensity to consume. The stochastic steady state (SSS) of our baseline economy with a 2% inflation target features a 10% ZLB frequency and a real rate just below 1%, in line with the data (Coibion et al., 2016). When we abstract from households' heterogeneity and consider a RANK model with the same parametrization, the ZLB frequency drops to 8% and the real rate rises by 26 basis points (bps).

What drives these differences between the RANK and the HANK economy? The reason is twofold. First, the presence of the idiosyncratic risk reduces the real interest rate in the

deterministic steady state (DSS) of the HANK economy. Although this result traces back to the seminal work of Aiyagari (1994), our setting grants it a novel perspective, as the precautionary savings reduce the room of manoeuvre for the central bank’s policy rate. In this way, the precautionary savings in the DSS alter the ZLB frequency, and consequently affect also the behavior of the macroeconomic variables in the SSS. Importantly, 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 of the lack of the ZLB.

Second, the interaction of aggregate uncertainty with wealth heterogeneity also increases precautionary savings at the SSS, and thus exert a downward force in the real interest rate vis-à-vis the RANK economy. As ZLB events affect disproportionately more wealth-poor agents, the presence of aggregate uncertainty and the possibility of the occurrence of large recessions in which the policy rate is constrained leads households – and especially those at the bottom of the wealth distribution – to increase their buffer of precautionary savings (or equivalently, reduce their borrowing positions).

These two channels then suggest that, as long as any variation in central bank’s inflation target may alter the frequency of ZLB events, then it would also alter households’ decisions about their optimal buffer of precautionary savings, and ultimately affect the level of the real interest rate in the SSS. In addition, the effect of changes in the inflation target on the level of real rates would be much larger in economies with more wealth inequality, as in this case households at the bottom of the distribution have a stronger motive for precautionary savings. This is the key mechanism that we put forward in this paper: the long-run level of the real interest rate crucially depends on the interaction between the stance of monetary policy and households’ inequality. Thus, our model generates a *long-run Fisher equation* in which the level of the real rate depends positively on the inflation target. That is, if we denote the steady-state nominal interest rate, real rate, inflation rate, and inflation target by i , r , π , and $\bar{\pi}$, respectively, then $i(\bar{\pi}) = r(\bar{\pi}) + \pi(\bar{\pi})$, such that $dr/d\bar{\pi} > 0$. Crucially, the magnitude of this derivative increases with the amount of wealth inequality.

We then leverage our model to evaluate the implications of the interaction of different

inflation targets and different amounts of wealth inequality on the level of the real interest rate. We show that reducing the inflation target from 4% – which corresponds to the average inflation rate between 1980 and 1999 – to 1.7% – which corresponds to the average inflation rates observed from the year 2000 on – reduces the real rate in the SSS by 21 bps. If in addition we consider a rise in wealth inequality – which in the model is achieved by increasing the volatility of the idiosyncratic labor earning process to replicate the same change in wealth inequality observed in the U.S. over the last three decades – then the drop in the level of the real rate at the SSS is amplified by an additional 29 bps. If we put these changes in the context of the decline in the level of the real rates observed over the last decades, the interaction between the drop in the inflation target and the rise in wealth inequality explains roughly a fourth of it. From this perspective, we propose a novel channel that is complementary to the factors proposed by the literature in accounting for the drop in the real rate, such as population aging (Carvalho, Ferrero and Nechio, 2016; Aksoy et al., 2019) and the differences in households’ marginal propensity to save out of permanent income (Mian, Straub and Sufi, 2021).

The non-neutrality of the inflation target on the level of real rates due to the presence of the ZLB has already been studied in the context of RANK models, in which changes in the inflation target alter real rates through a deflationary spiral channel: when the inflation target is sufficiently low and the probability of hitting the ZLB is relatively high, agents correctly expect that the ZLB may be binding in the future and internalize that the central bank may not be able to fully stabilize inflation. As a result, nominal interest rates decrease, reducing real rates, and eventually increasing the likelihood of hitting the ZLB even more.² However, we show that the quantitative relevance of the deflationary spiral channel hinges on the presence of the precautionary savings channel generated by market incompleteness: if we reduce the inflation target from 4% to 1.7% in a RANK model under exactly the same parametrization of our HANK economy, the real rate drops just by 12 bps. Thus, our mechanism generates

²While this mechanism has been already uncovered by Adam and Billi (2007), Nakov (2008), Hills, Nakata and Schmidt (2019), and Bianchi, Melosi and Rottner (2020), these papers mainly highlight this mechanism to argue for variation in the systematic monetary policy stance. For instance, Bianchi, Melosi and Rottner (2020) show that the presence of deflationary spirals calls for an asymmetric Taylor rule that allows the central bank to overshoot more likely its inflation target.

a fourfold increase in the relevance of deflationary spiral channel, from 12 to 50 bps. Thus, accounting for households' inequality is key in generating a quantitatively relevant long-run Fisher equation.

Our novel mechanism on the interaction between the inflation target and the amount of wealth inequality on the level of real rates hinges on the fact that the model explicitly incorporates the risk of hitting the ZLB. While there is a strand of the literature that studies the stochastic dynamics at the ZLB in RANK economies (e.g., Christiano, Eichenbaum and Rebelo, 2011; Coibion, Gorodnichenko and Wieland, 2012; Andrade et al., 2019), the analysis of the ZLB in economies with heterogeneous agents is usually carried out by imposing perfect foresight, and thus discarding the anticipation mechanism which is at the core of our paper (McKay, Nakamura and Steinsson, 2016; Guerrieri and Lorenzoni, 2017). Instead, we follow Fernández-Villaverde et al. (2015) by solving a fully non-linear model, in which households know ex-ante the probability of hitting an occasionally-binding ZLB. To do so, we move away from standard computation methods used in the heterogeneous-agent literature, such as Krusell and Smith (1998), which heavily rely on linear properties. We rather follow the work of Fernández-Villaverde, Hurtado and Nuño (2020), which shows how to solve non-linear heterogeneous-agent model via the use of neural networks. As long as neural networks can approximate the non-linear laws of motion of the states of the economy generated by the presence of the ZLB, we can derive the fully non-linear solution of the model. In this way, the model allows us to study the transmission of changes in the inflation target under aggregate uncertainty and in the presence of the ZLB constraint. From this point of view, this work is the first in introducing aggregate uncertainty and non-linear dynamics within a HANK economy.

The closest paper to ours is Auclert and Rognlie (2020), who analyze the link between inequality and aggregate demand. They find that an increase in income inequality in the deterministic steady state can lead to a substantial decline of consumption, output, and the real rate, especially when the monetary authority is always constrained by the ZLB.³ Our

³In a similar spirit, McKay and Reis (2016) show that the presence of the ZLB matters for the overall effects of automatic stabilizers.

paper looks at how inequality alters the *ex-ante* incentives to accumulate precautionary savings when households face an occasionally-binding ZLB, a mechanism that affects the levels of consumption, output, and the real rate in the stochastic steady state.

2 Model

2.1 Environment

The economy is populated by a unit measure of ex-ante homogeneous households, which are ex-post heterogeneous due to the realizations of a persistent uninsurable labor earning risk. The idiosyncratic shock follows a Markov chain process. Households' utility function over consumption and leisure is subject to an aggregate preference shock, which follows an AR(1) process. Households finance consumption expenditures with labor income as well as firms' profits, and are allowed to borrow up to an exogenous limit.

The production side of the economy consists of two levels. There is a continuum of intermediate good producers, which produce different varieties of the intermediate good using labor, subject to a Rotemberg (1982) price-setting cost. Then, the different varieties of the intermediate goods are bundled into the final consumption good by the final good producer.

Finally, the economy features a monetary authority that sets the nominal interest rate following a standard Taylor rule, subject to the ZLB constraint, as well as a fiscal authority, which levies progressive labor income taxes on the households to finance a fixed amount of outstanding public debt.

2.2 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, and follows a Markov chain with transition matrix Ω . We normalize the process such that the average realization of the idiosyncratic shock is $\int s_{it} di = 1$. The aggregate shock ξ_t is a preference shifter, and follows an

AR(1) process in logs

$$\log \xi_t = \rho_\xi \log \xi_{t-1} + \zeta_t, \quad (1)$$

where $\rho_\xi \in (0, 1)$ and ζ_t is normally distributed with mean 0 and variance σ_ξ . The use of the demand shocks follow the work of Krugman, Dominquez and Rogoff (1998), Eggertsson et al. (2003), and Eggertsson and Krugman (2012), which show that shifts in households' preferences – a reduced form for any change that affect households' leveraging capacity – is a powerful driver of ZLB events:

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

$$\max_{\{c_{i,t}, b_{i,t}, h_{i,t}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t \xi_t u(c_{i,t}, h_{i,t}) \quad (2)$$

$$\text{s.t. } c_{i,t} + b_{i,t} = w_t s_{i,t} h_{i,t} - \tau_t (w_t s_{i,t} h_{i,t})^{1-\gamma} + \Pi_t s_{i,t} + \frac{R_{t-1}}{\pi_t} b_{i,t-1}, \quad (3)$$

$$b_{i,t} \geq \underline{b}. \quad (4)$$

where β denotes the time-discount parameter. Households' optimization problem is subject to the budget constraint in Equation (3) and the borrowing constraint in Equation (4). 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 progressive taxation, where the parameter τ_t controls the level of the taxations and γ captures the degree of progressivity. When $\gamma = 0$, the tax rate is flat and does not depend on the level of households' labor earnings.⁴ Households also trade one-period non-contingent bonds, which yield the gross nominal return R_{t-1} . We then refer to the gross real return on bonds as $r_t = \frac{R_{t-1}}{\pi_t}$, which divides the gross nominal return by gross inflation. The borrowing constraint posits that households bond position is limited by the exogenous limit \underline{b} .

⁴Correia et al. (2013), McKay and Reis (2016), Wolf (2021) discuss the role of changes in the fiscal system in complementing monetary policy at the ZLB. In our framework, we abstract from this debate and consider a constant degree of progressivity over time.

Households have GHH preferences as in Greenwood, Hercowitz and Huffman (1988), such that

$$u(c_{i,t}, h_{i,t}) = \frac{1}{1-\sigma} \left(c_{i,t} - \chi \frac{h_{i,t}^{1+\nu}}{1+\nu} \right)^{1-\sigma}, \quad (5)$$

where σ denotes the risk-aversion parameter, ν is the inverse of the Frisch elasticity, and χ is a leisure shifter.

2.3 Firms

2.3.1 Final Good Firm

There is a final good that is produced by a representative final good firm and sold in a perfectly competitive final good market. The final good firm produces 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

$$Y_t = \left(\int_0^1 y_{j,t}^{\frac{\varepsilon-1}{\varepsilon}} dj \right)^{\frac{\varepsilon}{\varepsilon-1}}, \quad (6)$$

where ε is the elasticity of substitution between the different intermediate-input varieties. This CES aggregator yields the following optimal iso-elastic demand for intermediate good j

$$y_{j,t} = \left(\frac{p_{j,t}}{P_t} \right)^{-\varepsilon} Y_t, \quad (7)$$

where $p_{j,t}$ is the price of variety j , and P_t is the aggregate price level. These two price indexes are related together through the following equation,

$$P_t = \left(\int_0^1 p_{j,t}^{1-\varepsilon} dj \right)^{\frac{1}{1-\varepsilon}}. \quad (8)$$

2.3.2 Intermediate Good Firms

Each intermediate good firm j produces its variety using labor services $l_{j,t}$ according to the production function

$$y_{j,t} = l_{j,t}^\alpha, \quad (9)$$

where α captures the degree of diminishing returns to scale of the production function.

The cost minimization problem of intermediate good firms imply the marginal costs $m_{j,t}$ are defined as

$$m_{j,t} = \frac{w_t}{\alpha l_{j,t}^{\alpha-1}}. \quad (10)$$

Since all firms operate with the same labor-to-output ratio, both marginal costs and demanded labor services are equated across variety producers, that is, $m_{j,t} = m_t$ and $l_{j,t} = l_t$, for all $j \in [0, 1]$.

Intermediate good firms cannot freely adjust their prices, but rather face price adjustment costs as in Rotemberg (1982). This feature introduces nominal rigidities in the model. We follow Bayer et al. (2019) in the specification of the adjustment costs, with the only difference that we extend their settings to the case of a trend inflation possibly different from zero. In this case, the adjustment costs, $\Theta_{j,t}$, take the following functional form

$$\Theta_{j,t} \equiv \Theta \left(\frac{p_{j,t}}{p_{j,t-1}} \right) = \frac{\theta}{2} \left[\log \left(\frac{p_{j,t}}{p_{j,t-1} \times \tilde{\pi}} \right) \right]^2 Y_t, \quad (11)$$

where $\tilde{\pi}$ denotes the trend inflation which is targeted by the central bank, and θ crucially captures the degree of price stickiness. When $\theta = 0$, the economy does not feature nominal rigidities, and therefore collapses to a standard neoclassical heterogeneous-agent model with imperfect competition on the supply side.

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,

$$\mathbb{E}_t \sum_{k=t}^{\infty} \beta^k \left[\Pi_k(p_{j,k}) - \Theta \left(\frac{p_{j,k}}{p_{j,k-1}} \right) \right], \quad (12)$$

where the profits of the intermediate good firm, $\Pi_k(p_{j,k})$, equal

$$\Pi_k(p_{j,k}) = \left(\frac{p_{j,k}}{P_k} - m_k \right) \left(\frac{p_{j,k}}{P_k} \right)^{-\varepsilon} Y_k. \quad (13)$$

In the pricing protocol, we follow Hagedorn, Manovskii and Mitman (2019) by assuming that the price adjustment costs are virtual. Basically, the adjustment costs do affect firms' decision of setting optimal prices, but eventually do not result in any transfer of real resources. Accordingly, the profits of each firm, $\Pi_t(p_{j,t})$, describe the entire resources that are rebated back to the households. In this procedure, we further assume that profits are rebated to each household according to its idiosyncratic productivity level, such that we are consistent with the empirical observation that earnings-rich households tend to receive a disproportionately larger share of firm profits.

Finally, the overall solution of the problem of the intermediate good firms yields the following New Keynesian Philips curve

$$\log \left(\frac{\pi_t}{\tilde{\pi}} \right) = \beta E_t \left[\log \left(\frac{\pi_{t+1}}{\tilde{\pi}} \right) \frac{Y_{t+1}}{Y_t} \right] + \frac{\varepsilon}{\theta} \left(m_t - \frac{\varepsilon - 1}{\varepsilon} \right), \quad (14)$$

where $\pi_t = \frac{P_t}{P_{t-1}}$ is gross inflation rate.

2.4 Government

The government consists of a monetary authority and a fiscal authority. The monetary authority consists of a central bank which sets the gross nominal interest rate R_t according to a standard Taylor rule

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

\tilde{R} is the gross nominal rate in the deterministic steady state, and \tilde{Y} denotes the level of output in the deterministic steady state.

This specification of the Taylor rule implies two important considerations. First, the mone-

tary authority faces a ZLB constraint, such that the gross nominal interest rate cannot go below one. Second, when the central bank is not constrained by the ZLB, then it sets the nominal interest rate by reacting to changes in the inflation rate from its target – the parameter ϕ_π determines the strength with which this happens – and by reacting to changes in output from its steady state level – the parameter ϕ_y pins down the strength of this second channel.

On the other hand, the fiscal authority raises progressive labor earning taxes on the households to finance a fixed amount of outstanding debt \tilde{B} , such that the government budget constraint equals

$$\int_0^1 \tau_t (w_t s_{i,t} h_{i,t})^{1-\gamma} di = (r_t - 1) \tilde{B}. \quad (15)$$

We consider a fixed amount of outstanding debt so that we can calibrate our economy to be consistent with the amount of liquid assets of the U.S. economy, in line with the calibration strategy of McKay, Nakamura and Steinsson (2016) and the one-asset economy of Kaplan, Moll and Violante (2018).

2.5 Market Clearing Conditions

A competitive equilibrium for this economy implies the following market clearing conditions. First, the labor market clears, so that the overall efficiency units of hours provided by households equal the labor services demanded by the intermediate good firms, that is,

$$\int_0^1 l_{j,t} dj = \int_0^1 s_{i,t} h_{i,t} di. \quad (16)$$

Second, the bond market clears, so that the overall bond positions of households equal the outstanding government bonds issued by the fiscal authority, that is,

$$\tilde{B} = \int_0^1 b_{i,t} di. \quad (17)$$

Finally, the aggregate resource constraint posits that total output equals the sums of the value added of the intermediate good firms, and also equals the overall consumption of the households, that is,

$$Y_t = \int_0^1 l_{j,t}^\alpha dj = \int_0^1 c_{it} di. \quad (18)$$

3 Calibration

We calibrate the model to the U.S. economy by setting one time period of the model to correspond to a quarter. Table 1 shows the current parametrization of the model and the chosen targets. First, we set the gross inflation target of the monetary authority to $\tilde{\pi} = \exp(0.02/4)$, so that the annual inflation target is 2% in the baseline economy. As far as the Taylor rule parameters are concerned, we set the sensitivity of the nominal rate to output deviations from the steady state to $\phi_y = 0.1$, and that to changes in inflation to $\phi_\pi = 2.5$. Finally, to discipline the level of the real rate in the deterministic steady state, we set the time discount factor to $\beta = 0.997$. This choice implies that the real interest rate in the deterministic steady state equals 1%. Our choice binds the value of the real interest in the deterministic steady state to a value which is in the ballpark of the estimates provided by the literature (Del Negro et al., 2017; Fiorentini et al., 2018), and also consistent with the November 2018 FOMC’s Summary of Economic Projections. However, in the quantitative results we will show that different values of the inflation target imply a large variation in the level of real interest rates in the stochastic steady state. This happens even if all the economies share the same value of real interest rates in the deterministic steady state.

Next, we calibrate the demand shock process, which is the only source of aggregate risk in our model. Since we want to understand how household heterogeneity and changes in the monetary policy stance interact alter households’ decision by modifying the risk of hitting the ZLB, we require our model to be consistent with the frequency of ZLB episodes observed in the U.S. economy. To do so, we first set the persistence of the autoregressive process to $\rho_\xi = 0.6$, in line with the parametrization of Bianchi, Melosi and Rottner (2020), and then set the standard

deviation to $\omega_\xi = 0.0105$, so that – under a 2% inflation target – the model reproduces a 10% ZLB frequency, in line with the frequency observed in the U.S. in the post-war period (Coibion et al., 2016).

Table 1: Baseline Parametrization

| Parameter | | Value | Target/Source |
|-----------------------------|--|-------------|---|
| Panel A. Aggregate Risk | | | |
| ρ_ξ | AR coefficient of process for ξ | 0.6 | Bianchi, Melosi and Rottner (2020) |
| ω_ξ | Standard deviation of ξ shock | 0.0105 | 10% ZLB frequency |
| Panel B. Idiosyncratic Risk | | | |
| ρ_s | AR coefficient of process for s_t | 0.8 | 10% Average MPC |
| ω_s | Standard deviation of s_t shock | 0.05 | 30% Borrowers |
| \underline{b} | Borrowing limit | -0.29 | Monthly average labor income |
| Panel C. Preferences | | | |
| β | Discount factor | 0.997 | 1% 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 | 0.71 | Labor supply equals 1 in the DSS |
| Panel D. Production | | | |
| ε | Demand elasticity | 7.67 | 15% price markup |
| α | Labor share | 1 | Constant returns to scale |
| θ | Rotemberg price adjustment cost | 79.41 | Equivalent to 0.75 Calvo parameter |
| Panel E. Monetary Authority | | | |
| $\tilde{\pi}$ | 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 deviations from steady state output | 0.1 | Standard value |
| Panel F. Fiscal Authority | | | |
| B | Government outstanding bonds | 0.25 | Total liquid assets = 25% annual GDP |
| γ | Degree of tax progressivity | 0.18 | Heathcote, Storesletten and Violante (2017) |

As far as the calibration of the idiosyncratic risk is concerned, we want our economy to be consistent with the share of borrowers and savers observed in the data, and with the empirical estimates of the marginal propensity to consume (MPC). To do so, we first set the borrowing limit to $\underline{b} = -0.29$, so that it equals one monthly average wage. Given this tightness of the

borrowing constraint, we calibrate the labor earning process so that the model reproduces a share of borrowers of 30% (Kaplan, Violante and Weidner, 2014), and an average MPC of around 10%, which is at the lower end of the estimates provided by the literature (Johnson, Parker and Souleles, 2006; Parker et al., 2013; Broda and Parker, 2014). These values are obtained by setting the parameters of the AR(1) process governing the changes in households' efficiency unit of hours to $\rho_s = 0.8$ and $\omega_s = 0.05$. We then convert the AR(1) process into a three point Markov Chain using the Rouwenhorst (1982) method. This parametrization leads to a Gini index of wealth of 0.93, which is just slightly higher than the value of 0.87 that Kuhn and Ríos-Rull (2016) find when looking at the Gini index of liquid wealth of the 2013 U.S. economy.

After having disciplined the amount of aggregate and idiosyncratic uncertainty, we calibrate the parameters of the households as follows. We start by fixing the risk aversion to the standard value of $\sigma = 1$, and the Frisch elasticity of labor supply 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 $\chi = 0.71$ such that the aggregate value of the efficiency units of hours equals one in the deterministic steady state.

On the production side, we set the demand elasticity across different intermediate-input varieties to $\epsilon = 7.67$, such that the price markup is 15%, in line with the values commonly used in the New Keynesian literature (e.g., Christiano, Eichenbaum and Evans, 2005). We then fix the share of labor in value added to $\alpha = 1$, so that the production function features constant return to scale. The price adjustment parameter is set to $\theta = 79.41$ following Bianchi, Melosi and Rottner (2020), which use the equivalence result – at a first-order approximation – between the frequency in price adjustments derived by the Rotemberg protocol and the Calvo protocol to set the adjustment cost such that the prices are reset every four quarter, as it is commonly assumed in the literature.

Regarding the fiscal authority, we assume that the government has to finance a fixed amount

of outstanding bonds which amount to 25% of annual GDP, in line with the estimate of liquid wealth in the U.S. economy derived by Kaplan, Moll and Violante (2018). Then, we set the degree of tax progressivity to $\gamma = 0.18$, in line with the estimates of Heathcote, Storesletten and Violante (2017) using PSID data. Finally, we set the tax shifter τ_t by clearing the government budget constraint, which implies a value of 0.975.

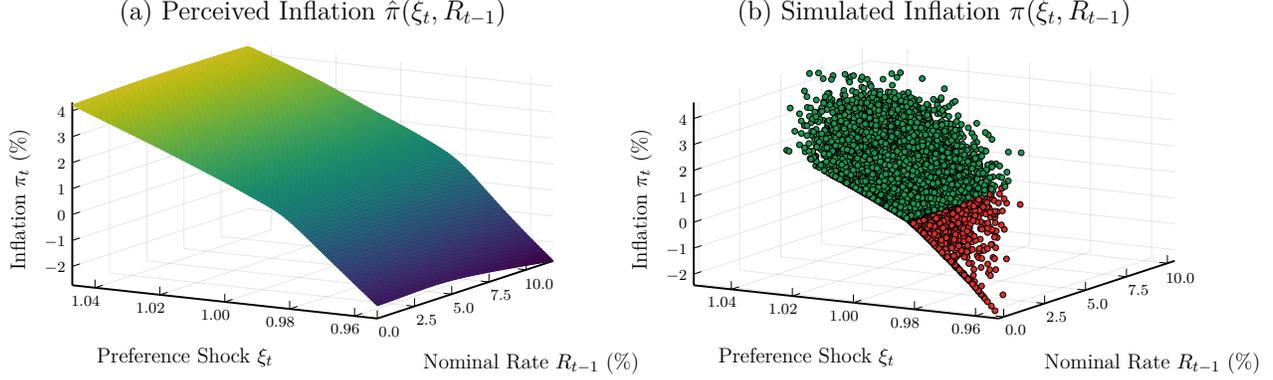
4 Algorithm

The computation of heterogeneous agent is complicated by the fact that an equilibrium requires the households to keep track of how the distribution of bonds evolves over time, so that it allows them to be able to properly form expectations. However, this feature makes the solution computationally intractable as one needs to keep track of an infinite-dimensional object. Thus, this class of model is often solved by assuming bounded rationality, as in Krusell and Smith (1998). Basically, the solution algorithm imposes that households need to keep track of just a moment of the bond distribution, and this moment tends to be the mean. Then, the Krusell and Smith (1998) hinges on linear law of motion approach, in which the individual policy function first solves for the updated individual policy function, and then updates the perceived law of motion derived within a stochastic simulation. The algorithm then repeats these steps until the convergence of the perceived law of motion.

However, our economy is inherently non-linear, and therefore the linear approach of Krusell and Smith (1998) would not capture the full dynamics of the economy. An important innovation of this paper is to introduce the solution of a non-linear economy within the literature of HANK models. This allows us to evaluate the stochastic steady state in an economy that features price rigidities, heterogeneous agents, and a zero lower bound. In doing so, we follow the approach of Fernández-Villaverde, Hurtado and Nuño (2020), showing that neural networks can fully approximate in a non-parametric way the non-linearity of the law of motion of a heterogeneous agent economy. We report a thorough description of the algorithm in Appendix A.

Figure 1 shows how our algorithm uncovers a marked non-linearity in both in both the

Figure 1: Non-Linearity due to the ZLB.

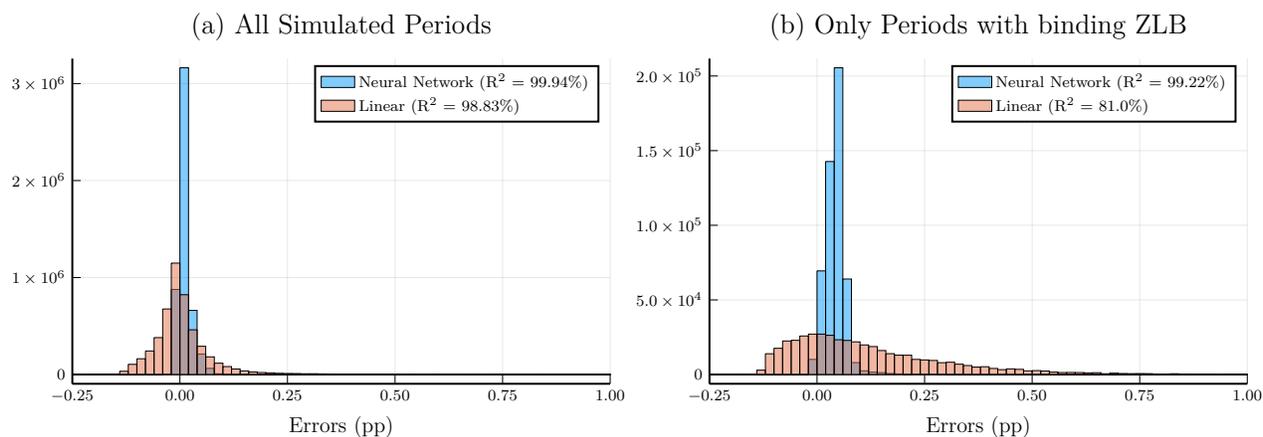


perceived and simulated inflation rate, a non-linearity which is generated by the presence of the zero lower bound, which holds for sufficiently low levels of the realizations of the demand shock process ξ_t .

To further corroborate the relevance of our approach, Figure 2 compares the nowcast errors on the dynamics of inflation generated by our neural network approach, with those implied by the 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. Panel (a) reports the errors for all simulated periods, and shows that the neural network approach increases the fit of households' expectations by increasing the R^2 from 98.8% to 99.9%. This improvement, is concentrated in the left-hand part of the distribution of the errors, that is, cases in which inflation tend to be lower than expected due to the deflationary spirals generated amidst a ZLB event. Indeed, our approach is able to capture properly the non-linear dynamics at the ZLB.

Panel (b) shows that the linear Krusell and Smith (1998) provides a very poor fit for inflation dynamics when the ZLB constraint binds, with an R^2 of 81.0%, whereas the neural network algorithm generates an R^2 of 99.2%. Hence, neural networks allow us to properly characterize the non-linear dynamics of the stochastic steady state at the ZLB, a feature that would not be present if we were to rely on standard computational algorithms.

Figure 2: Nowcast Errors for Inflation.



5 Results

We use the calibrated model to evaluate to what extent the presence of the ZLB constraint on the nominal rates set by the monetary authority alters the behavior of household-level and aggregate-level variables in a HANK economy, and uncover how the differences in the dynamics between our ZLB-HANK and the standard HANK economy depend crucially on both the level of the inflation target and the amount of the wealth inequality.

We start by comparing the responses of aggregate variables and households' optimal decision to the realizations of the aggregate demand shock in a HANK economy and in our ZLB-HANK model. We then illustrate how the presence of the ZLB constraint alters the stochastic steady state of the economy. Finally, we consider a set of exercises in which we either vary solely the inflation target of the monetary authority, or we change both the inflation target and the amount of wealth inequality.

5.1 The Macro of the ZLB: Aggregate Impulse-Response Functions

Our paper is the first to study the stochastic dynamics of a heterogeneous agent model by explicitly taking into account the non-linearity generated by the ZLB constraint. A natural question is that to understand the qualitative and quantitative implication of the non-linearity on the dynamics at the stochastic steady state.

The relevance of the non-linearity generated by the presence of the ZLB constraint can be appreciated by comparing the ergodic distribution of the key variables of our ZLB-HANK model with the implications of a counterfactual version of the baseline economy that abstracts from the the ZLB, that is, the standard HANK model.

Figure 3: Ergodic Distributions.

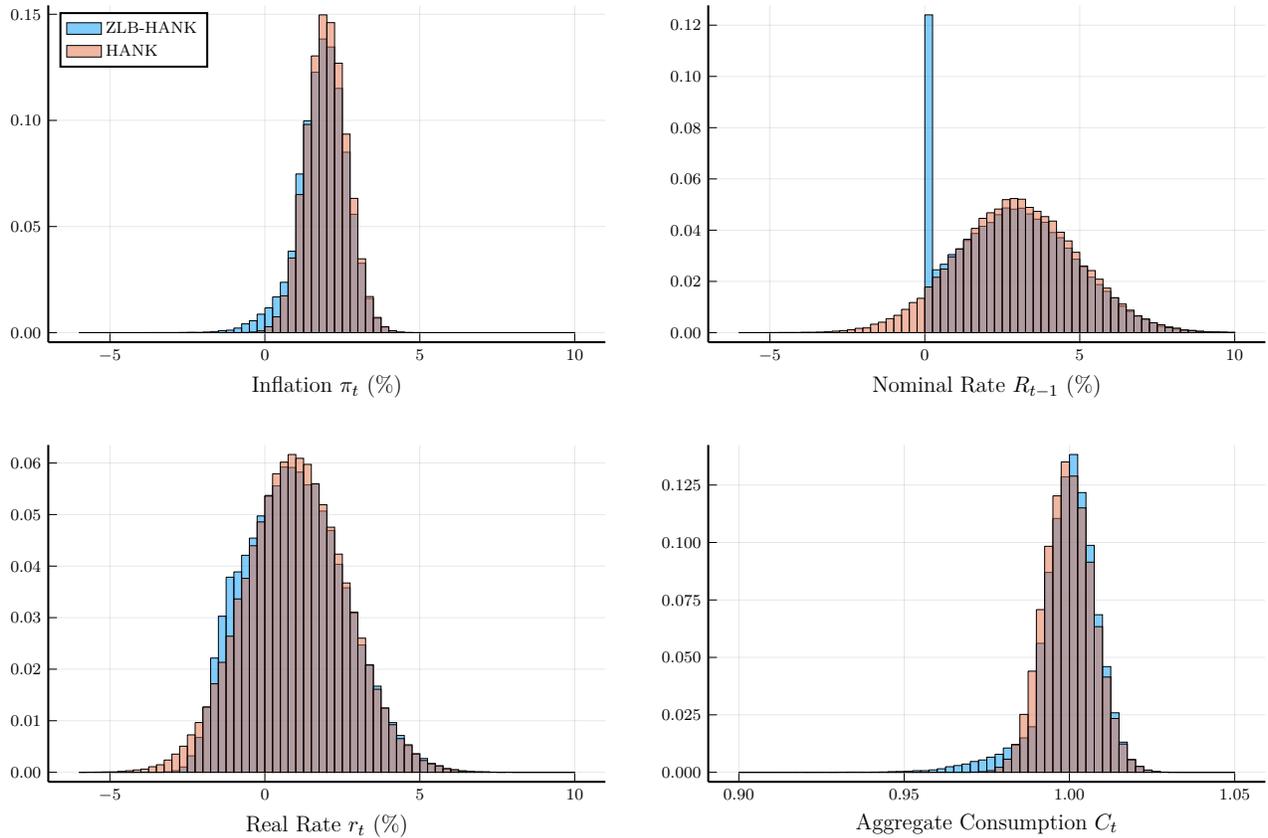
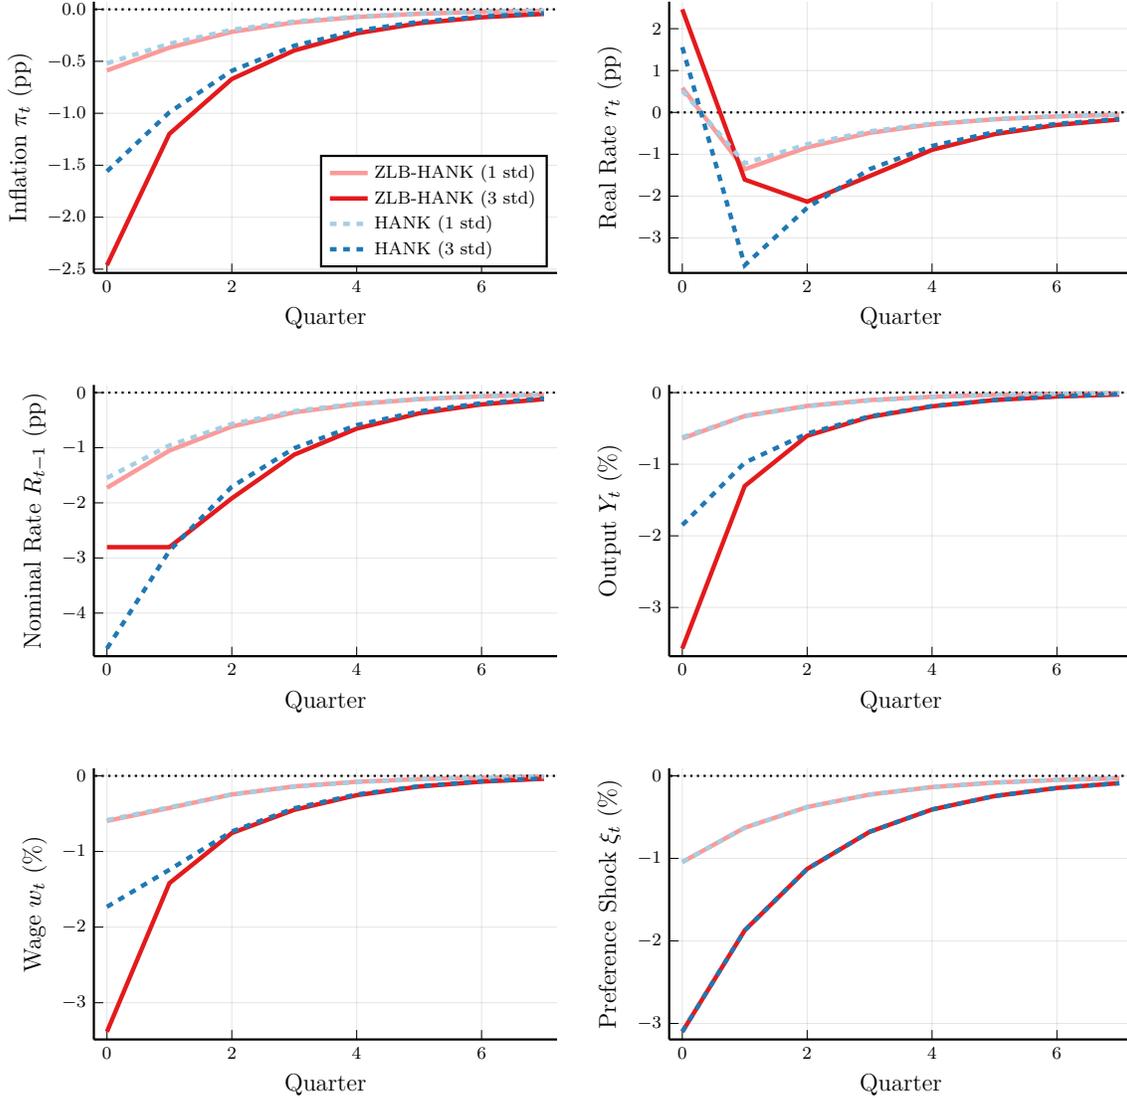


Figure 3 reports the ergodic distribution of inflation, the nominal interest rate, the real interest rate, and aggregate consumption in both the baseline ZLB-HANK and the HANK economy. The graph shows how the presence of the ZLB skews the dynamics of the model to the left: it corresponds to cases in which the nominal interest rate is constrained by the ZLB, and the economy experiences a sharp drop in aggregate consumption amidst a deflationary spiral. All these dynamics are absent in the standard HANK model.

The effect of the non-linearity of the ZLB constraint on the dynamics of inflation and output can be also evaluated by looking at the impulse-response function to small and large demand

shocks (i.e., a one standard deviation shock and a three standard deviation shock, respectively) in both the ZLB-HANK and the HANK economy, which we report in Figure 4.

Figure 4: Impulse Response Functions.



The plot shows that the impact response of the model to small demand shocks coincide both in the model with and without ZLB, as the shock is too small to generate a sufficiently large drop in the nominal interest rate that makes the ZLB binding. Instead, the responses to large demand shocks differ starkly depending on the presence or not of the ZLB. In the baseline model with the ZLB, a large shock brings the nominal interest rate down to zero, and it cannot go down further. This situation is accompanied by a much larger drop in both inflation and

output, as the monetary authority cannot reduce the interest rate to be as accommodative as it should be to minimize the negative effects of such a large shock. Figure 4 then highlights the important asymmetries generated by the ZLB: output and inflation drop much following big negative shocks when the constraint binds. Accordingly, ZLB events are characterized by deflation and substantial consumption losses for all households.

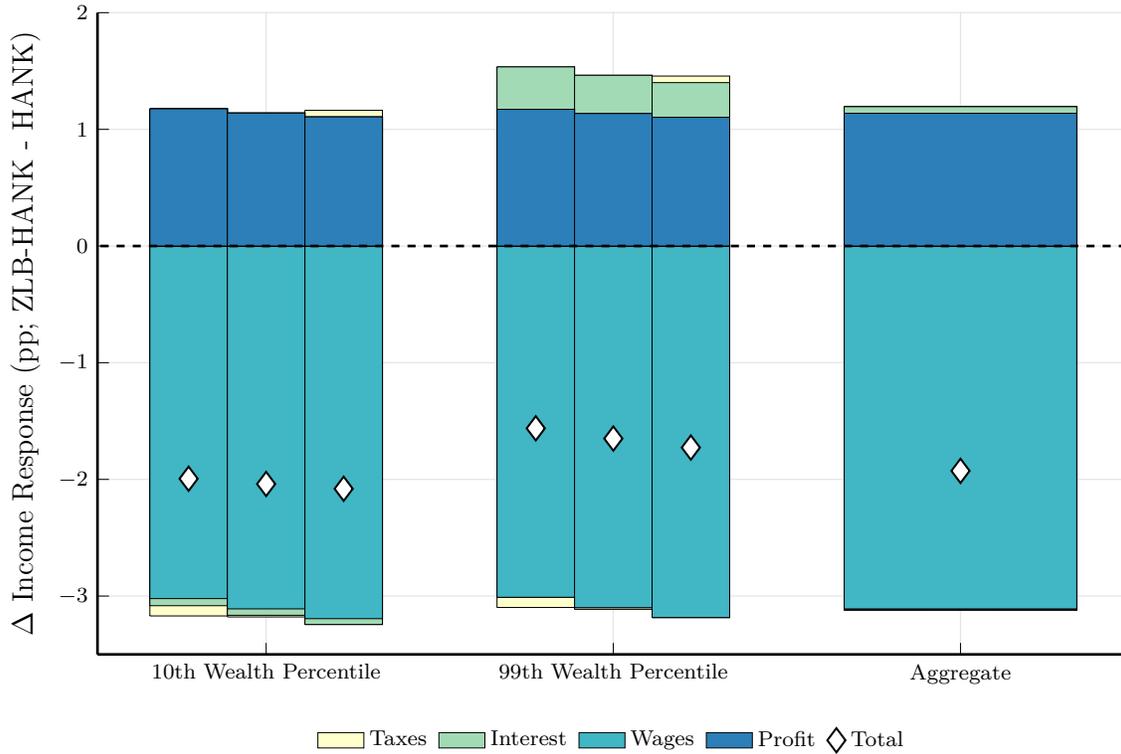
5.2 The Micro of the ZLB: Households' Impulse-Response Functions

The previous sector has highlighted how the presence of the ZLB constraint alters substantially the macroeconomic dynamics in response to (large realizations of) the demand shocks. However, the ZLB also alters the distributional consequences over business cycle fluctuations. To establish this result, we report in Figure 5 the differential impact response of total income to a negative demand shock by splitting total income across all its different components. Specifically, we compute the difference in the response of total income in the ZLB-HANK and the HANK economy (that is, the standard heterogeneous agent NK model that does not feature the ZLB), and disentangle the overall change in that of taxes, interest earnings, labor earnings, and profit earnings. We then report these values across different levels of labor earnings – by considering the relative response of income across the three different realizations of the individual labor earnings process – and across different levels of wealth – by considering the relative response of individuals in the first decile of the wealth distribution with that of wealth-rich households, that is, households in the top 1% of the wealth distribution. We also report the aggregate relative response of income.

Figure 5 indicates several important results on the distributional effects of a recessionary shocks once accounting for the ZLB constraint. First, overall the response of income is negative, for any realization of labor earnings, any position in the wealth distribution, as well as when looking at the behavior of aggregate labor income. This is related to the results of the previous section, which suggest that for sufficiently large and negative realizations of the demand shock, the drop in the nominal interest rate is so large to hit the ZLB constraint. In such a case, the

monetary authority cannot reduce further its stance to ease the effects of the recession, and thus the drop in output, income, and consumption are exacerbated relative to the very same HANK economy that does not feature the ZLB constraint.

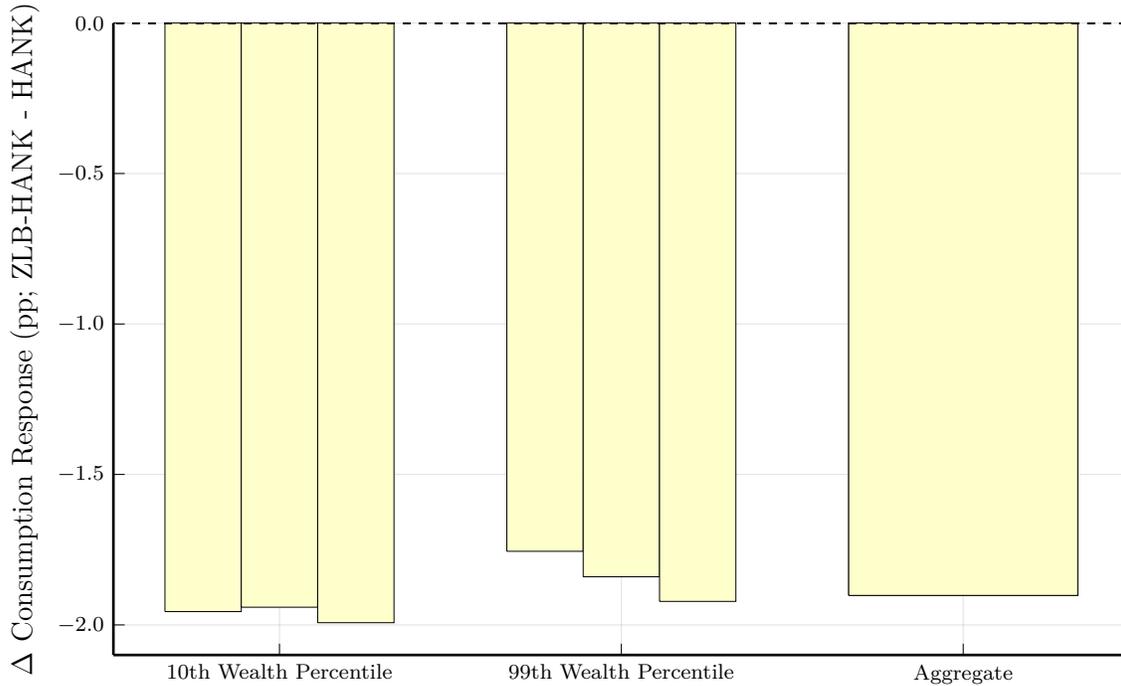
Figure 5: Individual Income Responses - ZLB-HANK vs. HANK.



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 in the changes in taxes, interest revenues, wages, and profits.

Second, the presence of the ZLB constraint amplifies the drop in the total income of wealth-poor households relative to that of wealth-rich individuals. Although both relative responses are negative, that associated to the households in the first decile of the wealth distribution is relatively larger. Again, this difference is due to the lack of the monetary authority of being able to stabilize the economy when the nominal interest rate hits zero. Indeed, we find what leads the drop in total income is the combination of the larger relative reduction in households' wages, and the relatively larger hike in interest payments. Thus, since the ZLB amplifies the drop in wages as a response to a recessionary demand shock – while raising the response of interest payments – it generates a differential impact along the wealth distribution such that

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.

the overall negative effects are concentrated on those individuals whose main income source are the labor earnings, that is, the wealth-poor households.

The presence of the ZLB alters the distributional consequences of business cycle fluctuations also when looking at the response of individual consumption. Figure 6 reports the differences in the impact response of consumption in the ZLB-HANK and HANK economy to a negative demand shock, and it does so for the three level of the labor earnings process, as well as for the households in first decile of wealth distribution, the individuals at the top 1% of the wealth distribution, and the differential response of aggregate consumption. Once again, we find that the response of consumption is relatively more negative when accounting for the ZLB constraint, and once again this is amplified at lower level of the wealth distribution. Specifically, the presence of the ZLB amplifies the drop in consumption of wealth-poor individuals by up to 0.2 percentage points, which accounts for 10% of the total additional negative response of consumption for the households in the first decile.

Overall, this evidence suggests that accounting explicitly for the presence of the ZLB con-

straint alters the effects of business cycle fluctuations of a HANK economy both at the aggregate level – with recessions becoming much more severe upon large negative realizations of the demand shock process – and the individual level – with recessions whose burden becomes even larger for households at the lower end of the wealth distribution, thus increasing consumption, income, and wealth inequality during a severe downturn.

5.3 The Deterministic and Stochastic Steady States

Standard models tend to predict that the real interest rate in the deterministic steady state is pinned down by structural parameters, and so there is no role for monetary policy to influence it. What we will show next is that although our model features a similar result in the deterministic steady state, it generates a relationship between changes in the inflation target and changes in real interest rates in the stochastic steady state.

What is the difference between these two steady state concepts? The deterministic steady state is an equilibrium in which all variables are constant and households ignore any source of aggregate risk. Basically, in our case it represents the scenario in which the standard deviation of the demand shock goes to zero, that is, $\omega_\xi = 0$. Instead, the stochastic steady state (also known as the risky steady state, e.g. Coeurdacier, Rey and Winant, 2011)) describes the equilibrium in which all variables are constant and households take their optimal decisions by taking into account the aggregate risk and by explicitly considering that the standard deviation of the demand shock is positive, that is, $\omega_\xi > 0$, although no shock arrives along the equilibrium path.

The difference between these two steady states is then that in the deterministic steady state households do not anticipate the effect of future aggregate shocks, and this case is often referred to as the perfect foresight equilibrium. Instead, in the stochastic steady state households are well aware of the existence of future aggregate shocks that may hit the economy. In either case, households internalize the risk due to the idiosyncratic labor earning process.

Importantly, the stochastic steady state is a more accurate description of the long run be-

havior of non-linear models as ours. For instance, in the deterministic steady state of our model the level of the inflation rate is 2%, which coincides with the inflation target of the central bank. Instead, at the stochastic steady state the level of inflation is relatively lower, and equals 1.9%. This happens because households internalize the possible occurrence of large negative demand shocks that may bring the economy to the ZLB constraint, triggering deflationary episodes: the deflationary spiral channel. This examples shows the crucial relevance of considering the dynamics of the stochastic steady state of a model that features a non-linearity such as the ZLB, and in the next sections we will highlight how our model implies that in the stochastic steady state changes in the inflation target do affect the level of real interest rates.

We then leverage the differences in the value of key macroeconomic variables at the DSS and SSS across different model economies. Specifically, we compare first our ZLB-HANK model to an otherwise identical heterogeneous agent NK economy which does not feature the non-negativity constraint on the policy rate, that is, the standard HANK economy, and then we compare the ZLB-HANK model to an otherwise identical economy which abstracts from individual heterogeneity and thus features a representative households, the ZLB-RANK.

Table 2 reports the results of this comparison exercise for the ZLB-HANK, the HANK with no ZLB constraint, and the ZLB-RANK model, showing the DSS and SSS values for inflation, the nominal interest rate, the real interest rate, output, and the frequency – as well as the duration in quarters – of ZLB events. In the case of the HANK economy without the bound on policy rates, we report the shadow frequency and duration of ZLB events, defined as all those circumstances in which the nominal interest rate of the HANK economy equals to or is less than zero.

The table shows that the two economies are characterized by exactly the same values for all these key macroeconomic variables at the DSS. However, the values at the SSS differ across the two model versions. Specifically, the HANK economy that does not feature the ZLB constraint displays SSS values for the inflation, nominal rate, and real rate which are virtually equal to those at the DSS. Indeed, the difference amounts to 1 basis point (bp), 3.5 bp, and 2.5 bp, respectively. Instead, the introduction of the ZLB reduces the SSS values of inflation, the

nominal rate, and the real rate by around 8 bp, 16bp, and 8bp, respectively.

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

| Variable | ZLB-HANK | | HANK | | ZLB-RANK | |
|--------------------------------------|----------|--------|------|---------|----------|-------|
| | DSS | SSS | DSS | SSS | DSS | SSS |
| Inflation | 2.0% | 1.91% | 2.0% | 1.99% | 2.0% | 1.93% |
| Nominal Rate | 3.0% | 2.80% | 3.0% | 2.96% | 3.22% | 3.08% |
| Real Rate | 1.0% | 0.89% | 1.0% | 0.97% | 1.22% | 1.15% |
| (Shadow) ZLB Frequency | - | 10.17% | - | (6.09%) | - | 8.35% |
| (Shadow) ZLB Duration in Quarters | - | 1.65 | - | (1.50) | - | 1.60 |

What drives this difference between the two model economies at the SSS? The answer is the presence of the ZLB constraint, which alters the dynamics of inflation, the nominal and real rates. This can be observed also by comparing the frequency and duration of ZLB events across model economies: while the frequency and duration of ZLB occurrences equal 10.165% and 1.6528 quarters in our ZLB-HANK, these (shadow) moments are lower in the standard HANK economy, and equal to 6.094% and 1.4951. As the ZLB does not allow the monetary authority to properly stabilize the economy when the nominal rate hits zero, this leads to more frequent and longer ZLB spells. Since the previous section shows that the ZLB affects relatively more the income and consumption of wealth-poor individuals, the ZLB-HANK economy is then characterized by a higher amount of households' precautionary savings, as individuals try to increase their buffer of savings (or equivalently, try to reduce their borrowing amounts) so to insure their consumption stream in those recessions in which the nominal rate hits zero. Higher precautionary savings exert a downward pressure on the level of the real interest rate, which then reduces the room of manoeuvre for the nominal rates of the central bank, and eventually

makes the ex-post realization of the ZLB events even more likely.

When comparing the DSS and SSS values of the ZLB-HANK and the ZLB-RANK models, we find that although both economies share the same value for the inflation target in the DSS, the real interest rate in the DSS of the ZLB-HANK is 22 bp lower than that of the ZLB-RANK. In the RANK economy the real interest rate is uniquely pinned down by time discount factor β . The heterogeneous-agent counterpart features the very same parametrization – including that of the discount factor – but it is characterized by a lower DSS real rate because of households’ precautionary savings against the idiosyncratic labor-earning risk.⁵ Although this result traces back to the seminal work of Aiyagari (1994), our setting grants it a novel perspective, as it generates another dimension through which households’ precautionary savings reduce the room of manoeuvre for the central bank’s policy rate. In this way, the precautionary savings in the DSS alter the ZLB frequency, and consequently affect also the behavior of the macroeconomic variables in the SSS. Indeed, the ZLB-HANK economy has a frequency of ZLB events which is around 10%, well above the 8% frequency generated by the RANK counterpart. Importantly, 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 of the lack of the ZLB.

Overall, the analysis of the DSS and the SSS shows that the real rate in the ZLB-HANK economy is endogenously determined by the interaction of the presence of the ZLB constraint and households’ heterogeneity. To properly measure the different channels that lead to a lower real rate in our baseline economy, we propose in Table 3 a decomposition exercise. First, by comparing the DSS and SSS values of the ZLB-RANK economy, we can back out the relevance of the *deflationary bias* channel highlighted by Adam and Billi (2007), Nakov (2008), Hills, Nakata and Schmidt (2019), and Bianchi, Melosi and Rottner (2020): when the inflation target is sufficiently low and the probability of hitting the ZLB is relatively high, agents correctly expect that the ZLB may be binding in the future and internalize that the central bank may

⁵At the DSS households do not require to increase their buffer of savings against the realization of aggregate risk, as this equilibrium ignore any realization of the aggregate demand shock. However, households still face the additional idiosyncratic uncertainty due to the uninsurable labor earning process.

Table 3: Decomposition Exercise.

| | Real Rate | Nominal Rate | Inflation |
|--|-----------|--------------|-----------|
| ZLB-RANK DSS | 1.22% | 3.22% | 2.0% |
| ZLB-RANK SSS | 1.15% | 3.08% | 1.93% |
| (i) Deflationary Bias | 0.08pp | 0.14pp | 0.07pp |
| ZLB-RANK DSS | 1.22% | 3.22% | 2.0% |
| ZLB-HANK DSS | 1.0% | 3.0% | 2.0% |
| (ii) Precautionary Savings Idiosyncratic Risk | 0.22pp | 0.22pp | 0.0pp |
| ZLB-RANK DSS | 1.22% | 3.22% | 2.0% |
| ZLB-HANK SSS | 0.89% | 2.8% | 1.91% |
| (iii) Total | 0.33pp | 0.42pp | 0.09pp |
| (iii)-(i)-(ii) Precautionary Savings Aggregate Risk | 0.03pp | 0.05pp | 0.02pp |

not be able to fully stabilize inflation. As a result, nominal interest rates decrease, reducing real rates, and eventually increasing the likelihood of hitting the ZLB even more. We find that the deflationary bias reduces the level of real rates, nominal rates, and inflation by 8 bp, 14 bp, and 7 bp, respectively.

Second, we compare the DSS values of the ZLB-RANK economy with those of our baseline heterogenous-agent model. As discussed above, this case allows us to identify the effect of households' *precautionary savings against their idiosyncratic labor-earning risk*. This channel leads to a 22 bp drop in the real rate, whereas it does not affect the inflation rate, as inflation target in both the HANK and RANK economies equals 2%. As a result, we also find a 22 bp bias for the nominal interest rate.

Finally, by computing the difference in the real rate, nominal rate, and inflation between the ZLB-RANK at the DSS and the ZLB-HANK at the SSS, and subtracting the relevance of the previous two mechanisms – the deflationary bias and the precautionary savings against the idiosyncratic risk – we can then measure the relevance of the *precautionary savings against the realization of aggregate uncertainty*. This channel is the one which is quantitatively less relevant, as it generates a reduction in the real rate, nominal rate, and inflation rate of 3 bp, 5 bp, and 2bp, respectively.

In the next two sections, we show that the downward bias in the level of the real interest rate in our ZLB-HANK economy depends crucially on the interaction between the inflation target set by the monetary authority and the amount of wealth heterogeneity, such that the bias gets substantially amplified in economies with relatively lower inflation targets and relatively larger inequality.

5.4 Inflation Target and Real Interest Rates

Instead, in our environment the precautionary savings amplify the deflationary spiral channel, through which the equilibrium level for the real rate depends on households' expectations about the costs of experiencing a ZLB event. Since the monetary authority can alter the frequency of

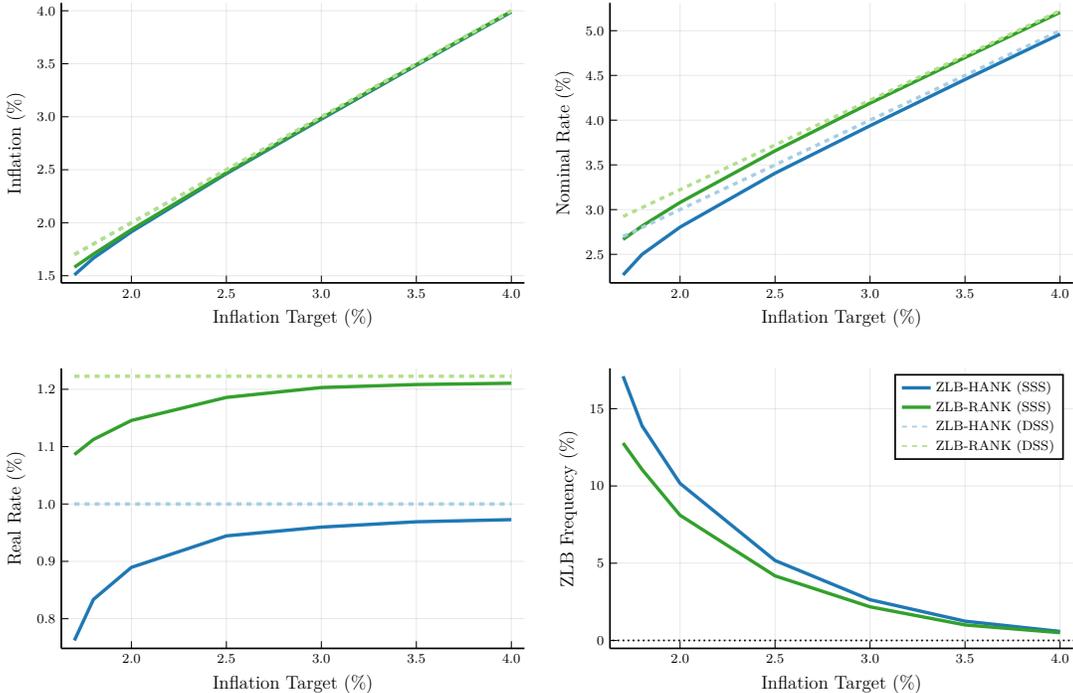
the ZLB spells by modifying its inflation target, it ends up affecting households' expectations and eventually the level of real interest rates. In other words, monetary policy is not neutral. More specifically, the model features a long-run Fisher equation, $i(\tilde{\pi}) = r(\tilde{\pi}) + \pi(\tilde{\pi})$, in which the real rate in the stochastic steady state depends on the central banks' inflation target $\tilde{\pi}$. In this setting, a higher inflation target raises the SSS level of the real rate, that is, $dr/d\tilde{\pi} > 0$.

To establish this result, we compare the level of the real interest rate in different model economies, which uniquely differ in the parameter $\bar{\pi}$, that is, the level of the inflation target of the monetary authority. Figure 7 reports the DSS and SSS levels of the real rate, the nominal rate, and the frequency of the ZLB for inflation targets between 1.7% and 4% in both the ZLB-HANK economy, as well as the ZLB-RANK model. First, the figure shows that – in both the RANK and HANK economy – when the inflation target is between 3% and 4%, then the central bank manages to achieve its inflation target in the stochastic steady state. Indeed, the inflation rate in the stochastic steady state almost coincides with that at the deterministic steady state. Moreover, the probability of experiencing a ZLB event is so low that the non-linearity in the setting of the nominal interest rate is not a quantitatively relevant force in shaping households' expectations and optimal decisions. Accordingly, the economy behaves as if there is no ZLB constraint, and changes in the inflation target do not alter the level of real interest rates.

However, when the inflation target of the central bank goes below 3%, then the non-linearity due to the ZLB kicks in, and the implications of the stochastic steady state diverge substantially from those of the deterministic steady state. First of all, the central bank start undershooting its inflation target: when the target is 1.7%, the inflation rate in the stochastic steady state is around 1.5%. This 0.2 percentage points undershooting comes together with a very high likelihood that the economy experiences a ZLB event. Indeed, the probability of hitting the floor in the nominal interest rate raises from below 4% when the inflation target is 4% up to around 18% when the inflation target is 1.7%. The ZLB frequency is relatively less sensitive to changes in the inflation target in the RANK model, as it raises from below 4% when the inflation target is 4% up to 13% when the inflation target is 1.7%.

In this exercise, we find that although the DSS level of the real interest rate is lower in the

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

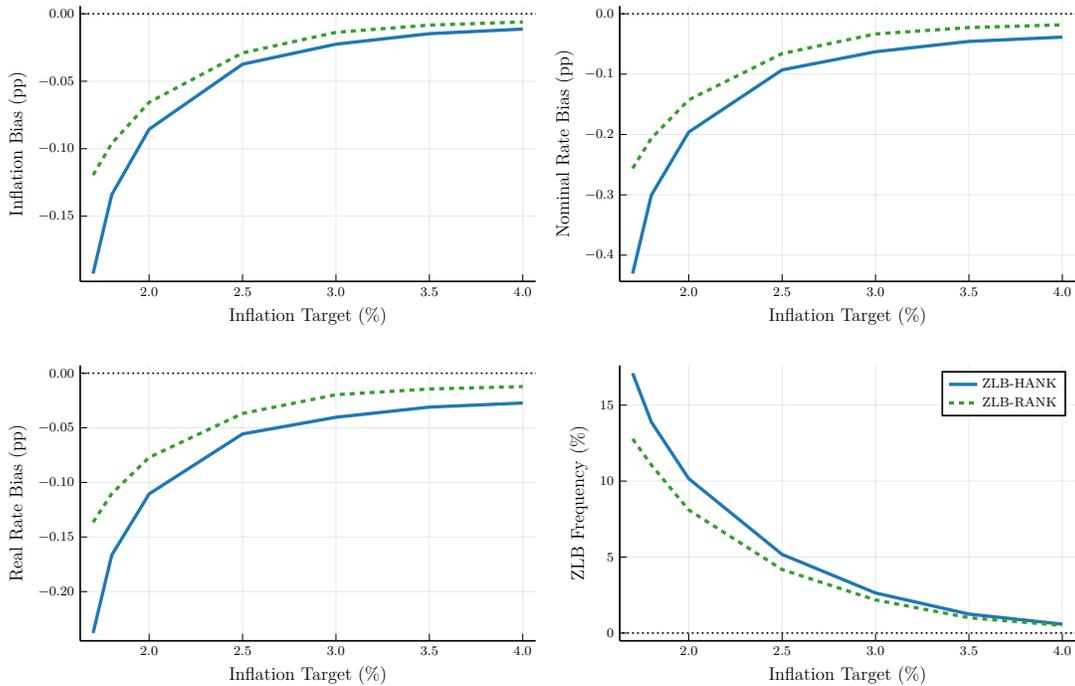


HANK economy than in the RANK – due to the precautionary savings against the idiosyncratic labor earning risk – in each economy the DSS level of the real rate is constant and independent of the monetary policy stance. Instead, at the SSS there is a relevant quantitative relationship between the inflation target and real interest rates when the target is below 3%. Indeed, we find that although the level of real rates is very close to the 2% of the deterministic steady state when the target is above 3%, the real rates drop substantially as the target shrinks. When the inflation target reaches 1.7%, then the level of real rates fall to 0.75%, amounting to roughly a 25 basis point drop from the level of real rates associated with the 4% inflation target. This is a novel prediction of the model which crucially hinges on incorporating explicitly the non-linearity of the ZLB constraint in the dynamics at the stochastic steady state. If we were to abstract from the ZLB or consider the deterministic steady state dynamics, then the model would predict the standard result that the monetary policy stance does not affect the level of the real interest rate.

The relationship between the level of the real rate at the SSS and the level of the inflation target is slightly less powerful in the RANK model, that is, the magnitude of the derivative

$dr/d\tilde{\pi}$ in RANK is lower than the same derivative $dr/d\tilde{\pi}$ in HANK. All these dynamics can be better appreciated in Figure 8, which reports the differences between the SSS and DSS values of the RANK and HANK economies at each value of the inflation target. Specifically, the drop in the real rate in the ZLB-RANK going from an inflation target of 4% to 1.7% is 10 bps lower than in the ZLB-HANK economy.

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



What drives the relationship between the inflation target and the real interest rates? To inspect this relationship, let us refer again to the Fisher equation. In the steady state, standard models pin down the level of the real rate as the inverse of the households' time discount parameter, whereas the level of inflation is a policy parameter. Thus, given these two structural parameters, standard models recover the level of the nominal interest rate.

Although our economy features the same property in its deterministic steady state, the dynamics change in the stochastic steady state. In this case, the values of both the inflation rate and the real interest rate are endogenously determined, so that the model features an equilibrium long-run Fisher equation. Namely, the link between the inflation target and the real interest rate is twofold. On the one hand, the relationship works through changes in households'

precautionary savings. As the probability of ZLB events is almost as high as 20% with a 1.6% inflation target, and those events affect disproportionately more wealth-poor agents, a decline in the inflation target which raises the probability that policy rates become constrained has the implication of raising households' precautionary savings. As a result, the interest rate declines. We will provide further evidence on this mechanism in the next section.

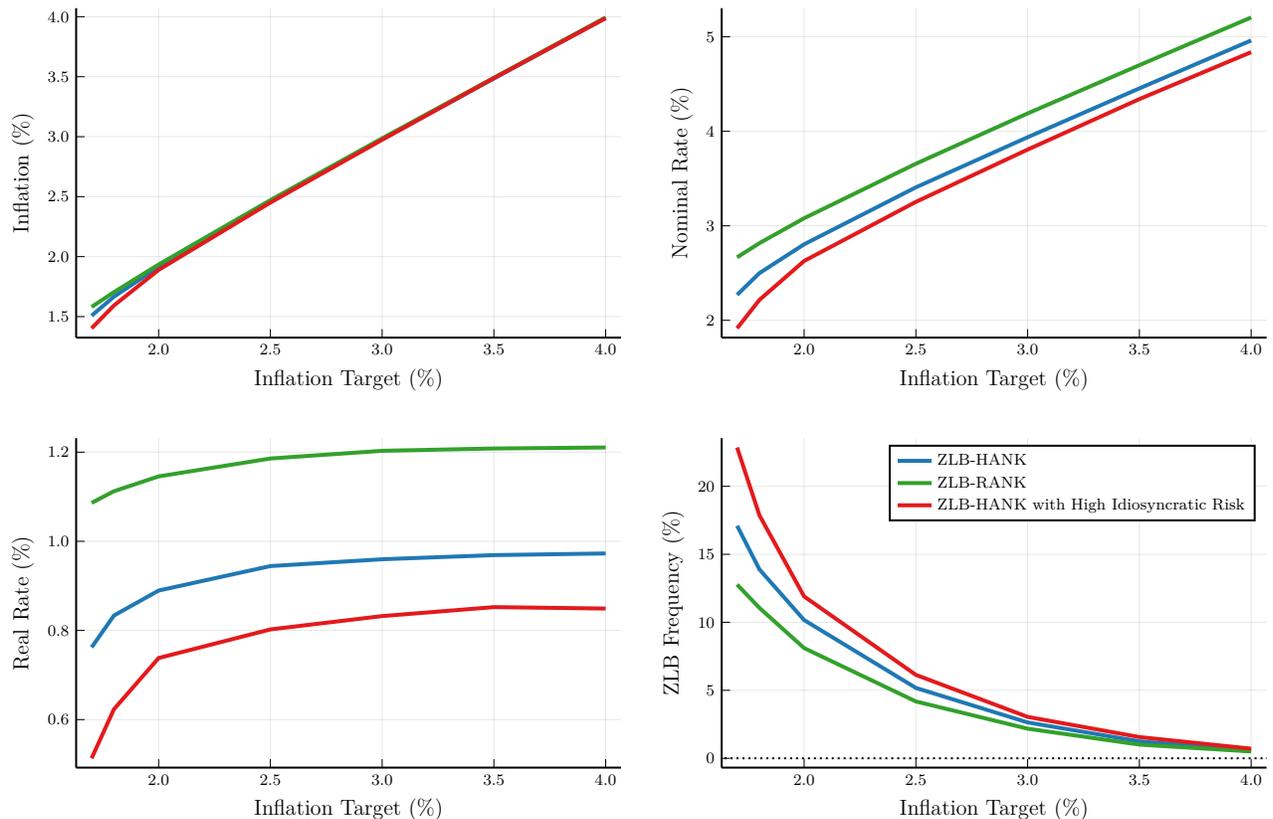
Importantly, the non-neutrality of the inflation target arises already in the RANK model, as indicated in Figures 7 and 8. In this economy, the inflation target alters real rates through a deflationary spiral channel already highlighted by Adam and Billi (2007), Nakov (2008), Hills, Nakata and Schmidt (2019), and Bianchi, Melosi and Rottner (2020): when the inflation target is sufficiently low and the probability of hitting the ZLB is relatively high, agents correctly expect that the ZLB may be binding in the future and internalize that the central bank may not be able to fully stabilize inflation. As a result, nominal interest rates decrease, reducing real rates, and eventually increasing the likelihood of hitting the ZLB even more. Thus, this mechanism self reinforces itself, with the ultimate implication that reductions in the inflation target may substantially raise the probability of hitting the ZLB and reduce the level of real rates. While the non-neutrality of the inflation target in the RANK economy is quite limited, in the next section we show that changes in the inflation target can lead to quantitatively relevant variations in the real rate at high level of wealth inequality.

5.5 The Role of Wealth Inequality

What is the role of household heterogeneity in generating the relationship between changes in the inflation target and the variation in the real interest rate at the stochastic steady state? The non-neutrality of monetary policy crucially depends on the interaction of the inflation target with the amount of wealth inequality in the economy. To establish this result, we compare the effects of changes in the inflation target of our baseline ZLB-HANK economy and its ZLB-RANK counterpart with one additional ZLB-HANK economy which differs from the benchmark case for the amount of wealth inequality. More specifically, we increase the standard deviation

of the labor earning risk process, ω_s , such that the Gini index goes from the value of 0.93 our baseline economy to 0.96. This increase in the Gini index of wealth corresponds to the change observed in the U.S. economy between 2007 and 2013 (Kuhn and Ríos-Rull, 2016). Figure 9 reports the SSS values for the inflation target, the nominal interest rate, the real rate, and the ZLB frequency for the baseline HANK model, the RANK counterpart, and the additional economy with higher idiosyncratic risk.

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



We find that the non-neutrality of monetary policy does depend substantially on the level of wealth inequality: when we increase the amount of wealth inequality with the higher level of idiosyncratic risk, then the drop in the real rate between the 4% and 1.7% inflation targets amounts to 34 bp, a change which almost halves the level of the real rate. Thus, our mechanism can generate a variation in the real rate which is three times as large as that of the deflationary spiral channel, and thus suggest that accounting for households' inequality is key in generating a quantitatively relevant long-run Fisher equation, in which the monetary policy stance does

affect the long-run values of the key macroeconomic aggregates.

Why does a higher wealth inequality lead to a larger drop in the real rate for lower levels of the inflation target? Again, the answer lies in the role of the precautionary savings. Households accumulate an additional buffer of savings (or alternatively, reduce their borrowing exposure) to shield their consumption stream during a ZLB recession. This is especially important for low-wealth individuals, who bear the burden of these severe downturns. In the low idiosyncratic risk economy, as the sensitivity of the ZLB frequency to the inflation target is quite reduced, ranging from 9% to 14%, then households do not boost their precautionary savings at lower levels of the inflation target. However, in the high idiosyncratic risk economy the frequency of the ZLB gets as high as 23% at the 1.7% inflation target. Since the ZLB is visited once every four quarters on average and the higher level of wealth inequality means that there are many more individuals at the lower end of the distribution who bear the costs of the ZLB spells, households decide to substantially increase their saving buffers, leading to a drop in the level of the real rate.

Overall, these results are consistent with the observed dynamics of the U.S. economy over the recent decades. As the level of 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 – then the economy has started experiencing ZLB events, the inflation rate has been constantly below the 2% target, and the level of real interest rates has dropped substantially. Thus, our results provide a novel rationale that could comprehensively explain all these events jointly.

6 Concluding Remarks

This paper shows how to explicitly incorporate the presence of the ZLB constraint into a heterogeneous agent New Keynesian economy. To do so, we propose a novel neural network algorithm that allows to solve for the fully non-linear stochastic equilibrium. In this way, we can show that the presence of the ZLB constraint substantially alters the dynamics and the properties of the HANK economy.

First, the presence of the ZLB alters the response of aggregate output, inflation, and the real rate to large shocks. In the cases that the nominal rate hits the ZLB, the monetary authority loses its stabilizational power. As a result, the drop in output is twice as large as in a model without the constraint on the policy rates, and it generates a relatively more muted response on the real interest rate, accompanied by a larger decrease in the wage. Thus, the presence of the ZLB further depresses labor income vis-à-vis capital income during a severe recession.

Second, the presence of the ZLB amplifies the negative recessionary effects on income and consumption for wealth-poor households, as the income of these individuals is disproportionately tilted towards labor earnings. In addition, wealth-poor households do not have a sufficiently large buffer to insure their consumption stream during a downturn. Consequently, the ZLB leads to a larger increase in income and consumption inequality during severe recessions.

Third, households react to the additional risk posed by the presence of the ZLB by raising their precautionary savings, even above the level that characterizes any heterogeneous agent economy at its deterministic steady state. Overall, the precautionary savings exert a downward pressure on interest rates, thus reducing the room of manoeuvre for the central banks policy stance. This behavior leads to an even larger ex-post equilibrium frequency and duration of the ZLB events.

Fourth, we uncover a novel relationship between the stance of monetary policy – and namely, the inflation target set by the central bank in its Taylor rule – and the level of real interest rates. Usually, the level of real interest rates in standard New Keynesian models is pinned down by structural parameters, and so although the level of real rates may shape the optimal conduct of monetary policy, there is no feedback effect coming out from between monetary policy decisions. Instead, in our environment the precautionary savings lead to an equilibrium level for the real rate which depends on households' expectations about the costs of experiencing a ZLB event. Since the monetary authority can alter the frequency of the ZLB spells by modifying its inflation target, it ends up affecting households' precautionary savings and eventually the level of real interest rates. In other words, monetary policy is not neutral. This would not be the case if we were to derive the model in a perfect foresight equilibrium, in which the non-linearity of

the ZLB are plugged in the model but without properly affecting households' expectations and optimal decisions.

Finally, the non-neutrality of monetary policy is largely amplified at higher values of wealth inequality. The model predicts that the interaction between a drop in the inflation target from 4% to 1.8% – which corresponds to the variation in the average inflation rate over the last three decades – and an increase in the Gini index of wealth of three percentage points – which corresponds to the observed rise in wealth inequality between 2007 and 2013 – reduces the level of the real interest rate by around 50 basis points. This decline in the level of the real interest rate explains around 25% of the 200 basis point reduction over the last decades that has been recently estimated by Del Negro et al. (2017) and Fiorentini et al. (2018). Hence, we provide a novel explanation for the drop in real interest rates, such that this decline may not be only due to demographic factors, reduction in the productivity of the economy, or variation in the convenience yields demanded by investors, but also may be a byproduct of the more aggressive stance of the central bank in reducing the level of the inflation target.

References

- Adam, Klaus, and Roberto M Billi.** 2007. “Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates.” *Journal of Monetary Economics*, 54(3): 728–752.
- Aiyagari, S Rao.** 1994. “Uninsured Idiosyncratic Risk and Aggregate Saving.” *Quarterly Journal of Economics*, 109(3): 659–684.
- Aksoy, Yunus, Henrique S Basso, Ron Smith, and Tobias Grasl.** 2019. “Demographic Structure and Macroeconomic Trends.” *American Economic Journal: Macroeconomics*, 11(1): 193–222.
- Andrade, Philippe, Jordi Galí, Herve Le Bihan, and Julien Matheron.** 2019. “The Optimal Inflation Target and the Natural Rate of Interest.” *Brookings Papers on Economic Activity*, Fall 2019 Edition.

- Auclert, Adrien, and Matthew Rognlie.** 2020. “Inequality and Aggregate Demand.” Mimeo.
- Barsky, Robert, Alejandro Justiniano, and Leonardo Melosi.** 2014. “The Natural Rate of Interest and its Usefulness for Monetary Policy.” *American Economic Review*, 104(5): 37–43.
- Bayer, Christian, Ralph Lütticke, Lien Pham-Dao, and Volker Tjaden.** 2019. “Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk.” *Econometrica*, 87(1): 255–290.
- Bewley, Truman.** 1980. “Interest Bearing Money and the Equilibrium Stock of Capital.” Mimeo.
- Bianchi, Francesco, Leonardo Melosi, and Matthias Rottner.** 2020. “Hitting the Elusive Inflation Target.” Mimeo.
- Broda, Christian, and Jonathan A Parker.** 2014. “The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption.” *Journal of Monetary Economics*, 68: S20–S36.
- Carvalho, Carlos, Andrea Ferrero, and Fernanda Nechio.** 2016. “Demographics and Real Interest Rates: Inspecting the Mechanism.” *European Economic Review*, 88: 208–226.
- Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber.** 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, Lawrence J, Martin Eichenbaum, and Charles L Evans.** 2005. “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy.” *Journal of Political Economy*, 113(1): 1–45.
- Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo.** 2011. “When is the Government Spending Multiplier Large?” *Journal of Political Economy*, 119(1): 78–121.

- Coeurdacier, Nicolas, Helene Rey, and Pablo Winant.** 2011. “The Risky Steady State.” *American Economic Review*, 101(3): 398–401.
- Coibion, Olivier, Marc Dordal-i Carreras, Yuriy Gorodnichenko, and Johannes Wieland.** 2016. “Infrequent but Long-Lived Zero Lower Bound Episodes and the Optimal Rate of Inflation.” *Annual Reviews*, 497–520.
- Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland.** 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.
- Correia, Isabel, Emmanuel Farhi, Juan Pablo Nicolini, and Pedro Teles.** 2013. “Unconventional Fiscal Policy at the Zero Bound.” *American Economic Review*, 103(4): 1172–1211.
- Del Negro, Marco, Domenico Giannone, Marc P Giannoni, and Andrea Tambalotti.** 2017. “Safety, Liquidity, and the Natural Rate of Interest.” *Brookings Papers on Economic Activity*, 2017(1): 235–316.
- Eggertsson, Gauti B, and Paul Krugman.** 2012. “Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach.” *The Quarterly Journal of Economics*, 127(3): 1469–1513.
- Eggertsson, Gauti B, et al.** 2003. “Zero Bound on Interest Rates and Optimal Monetary Policy.” *Brookings Papers on Economic Activity*, 2003(1): 139–233.
- Fernández-Villaverde, Jesús, Grey Gordon, Pablo Guerrón-Quintana, and Juan F Rubio-Ramirez.** 2015. “Nonlinear Adventures at the Zero Lower Bound.” *Journal of Economic Dynamics and Control*, 57: 182–204.
- Fernández-Villaverde, Jesús, Samuel Hurtado, and Galo Nuño.** 2020. “Financial Frictions and the Wealth Distribution.” Mimeo.

- Fiorentini, Gabriele, Alessandro Galesi, Gabriel Pérez-Quirós, and Enrique Sentana.** 2018. “The Rise and Fall of the Natural Interest Rate.” Mimeo.
- Gornemann, Nils, Keith Kuester, and Makoto Nakajima.** 2016. “Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy.” Mimeo.
- Greenwood, Jeremy, Zvi Hercowitz, and Gregory W Huffman.** 1988. “Investment, Capacity Utilization, and the Real Business Cycle.” *The American Economic Review*, 402–417.
- Guerrieri, Veronica, and Guido Lorenzoni.** 2017. “Credit Crises, Precautionary Savings, and the Liquidity Trap.” *The Quarterly Journal of Economics*, 132(3): 1427–1467.
- Hagedorn, Marcus, Iourii Manovskii, and Kurt Mitman.** 2019. “The Fiscal Multiplier.” Mimeo.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante.** 2017. “Optimal Tax Progressivity: An Analytical Framework.” *The Quarterly Journal of Economics*, 132(4): 1693–1754.
- Hills, Timothy S, Taisuke Nakata, and Sebastian Schmidt.** 2019. “Effective Lower Bound Risk.” *European Economic Review*, 120: 103321.
- Huggett, Mark.** 1993. “The Risk-free Rate in Heterogeneous-agent Incomplete-insurance Economies.” *Journal of economic Dynamics and Control*, 17(5-6): 953–969.
- Johnson, David S, Jonathan A Parker, and Nicholas S Souleles.** 2006. “Household Expenditure and the Income Tax Rebates of 2001.” *American Economic Review*, 96(5): 1589–1610.
- Kaplan, Greg, Benjamin Moll, and Giovanni L Violante.** 2018. “Monetary Policy According to HANK.” *American Economic Review*, 108(3): 697–743.

- Kaplan, Greg, Giovanni L Violante, and Justin Weidner.** 2014. “The Wealthy Hand-to-Mouth.” *Brookings Papers on Economic Activity*, 1: 77–153.
- Krugman, Paul R, Kathryn M Dominquez, and Kenneth Rogoff.** 1998. “It’s Baaack: Japan’s Slump and the Return of the Liquidity Trap.” *Brookings Papers on Economic Activity*, 1998(2): 137–205.
- Krusell, Per, and Anthony A Smith, Jr.** 1998. “Income and Wealth Heterogeneity in the Macroeconomy.” *Journal of Political Economy*, 106(5): 867–896.
- Kuhn, Moritz, and José-Víctor Ríos-Rull.** 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.
- Lilley, Andrew, and Kenneth Rogoff.** 2020. “The Case for Implementing Effective Negative Interest Rate Policy.” *Strategies for Monetary Policy*.
- Maliar, Lilia, Serguei Maliar, and Fernando Valli.** 2010. “Solving the Incomplete Markets Model with Aggregate Uncertainty using the Krusell–Smith Algorithm.” *Journal of Economic Dynamics and Control*, 34(1): 42–49.
- McKay, Alisdair, and Ricardo Reis.** 2016. “The Role of Automatic Stabilizers in the US Business Cycle.” *Econometrica*, 84(1): 141–194.
- McKay, Alisdair, Emi Nakamura, and Jón Steinsson.** 2016. “The Power of Forward Guidance Revisited.” *American Economic Review*, 106(10): 3133–58.
- Mian, Atif, Ludwig Straub, and Amir Sufi.** 2021. “Indebted Demand.” *The Quarterly Journal of Economics*, forthcoming.
- Nakov, Anton.** 2008. “Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate.” *International Journal of Central Banking*, 4: 73–127.

- Parker, Jonathan A, Nicholas S Souleles, David S Johnson, and Robert McClelland.** 2013. “Consumer Spending and the Economic Stimulus Payments of 2008.” *American Economic Review*, 103(6): 2530–53.
- Rotemberg, Julio J.** 1982. “Sticky prices in the United States.” *Journal of Political Economy*, 90(6): 1187–1211.
- Rouwenhorst, Geert K.** 1982. “Asset Pricing Implications of Equilibrium Business Cycle Models.” *Frontiers of Business Cycle Research*.
- Wolf, Christian.** 2021. “Interest Rate Cuts vs. Stimulus Payments: A Macro Equivalence Result.” Mimeo.
- Young, Eric 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 computational algorithm for solving a fully-nonlinear HANK economy around its stochastic steady state is based on the stochastic simulation algorithm as in Maliar, Maliar and Valli (2010), whereas the bond distribution is represented as a histogram following the procedure of Young (2010). We start by reporting the Euler equation of the individual household problem, that is⁶

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

where μ is the Lagrange multiplier associated with the borrowing constraint of the households. Using the budget constraint we can rewrite this as

$$\begin{aligned} \tilde{b}' &= \frac{R_{-1}}{\pi} b + w s h + \Pi s - T \\ &- \left[\mu + \beta E \left(\frac{\frac{\xi'}{\xi} \frac{R}{\pi'}}{\left(\frac{R}{\pi'} b' + w' s' h' - \tau (w' s' h')^{1-\gamma} + \Pi' s' - b'(b') - \chi \frac{h'^{1+\nu}}{1+\nu}\right)^\sigma} \right) \right]^{-1/\sigma} - \chi \frac{h^{1+\nu}}{1+\nu} \end{aligned} \quad (\text{A.2})$$

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 points⁷ for $b \in [b_{min}, b_{max}]$ and $R \in [R_{min}, R_{max}]$ ⁸, in addition to the grids for the idiosyncratic state $s \in \{s_m\}_{m=1}^M$ and the aggregate state $\xi \in \{\xi_j\}_{j=1}^J$. The individual problem is then solved on this grid according to the following algorithm.

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 deterministic steady state bond policy function for all aggregate states (R_{-1}, ξ) .

⁶Throughout this section, with an abuse of notation we refer to individual firm-level and household-level variables without reporting the j and i indexes, respectively.

⁷Note for b we use the polynomial rule from Maliar, Maliar and Valli (2010) to place more grid points near the borrowing constraint. The grid for R is evenly spaced, however.

⁸In economies with a zero lower bound (ZLB) we can set $R_{min} = 1$

2. For each grid point (b, R_{-1}, s, z) , substitute the assumed bond policy function $b'(b, R_{-1}, s, z)$ in the right hand side of (A.2), set the Lagrange multiplier equal to zero and compute the new bond policy function in the left-hand side of equation (A.2). The required labor policy h can be computed from $h = \left(\frac{w}{\chi}\right)^{1/\nu}$. For each point that does not belong to $[b_{min}, b_{max}]$, set $\tilde{b}'(b, R_{-1}, s, z)$ equal to the corresponding boundary value.
3. Compute the bond function for the next iteration $\tilde{\tilde{b}}'(b, R_{-1}, s, \xi)$ using the following updating formula

$$\tilde{\tilde{b}}'(b, R_{-1}, s, \xi) = \lambda_b \tilde{b}'(b, R_{-1}, s, \xi) + (1 - \lambda_b) b'(b, R_{-1}, s, \xi) \quad (\text{A.3})$$

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

Iterate on steps 2 and 3 until the maximum 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, Maliar and Valli (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 inflation π_t , the real wage w_t and profits Π_t . We use a Krusell and Smith (1998) type of algorithm to predict $\log \pi_t$ and $\log \left(\frac{\pi_{t+1}}{\tilde{\pi}}\right) \frac{Y_{t+1}}{Y_t}$, given the aggregate states (R_{t-1}, ξ_t) . Given this, we can back out marginal cost m_t from the NK Philips curve and all other ‘‘aggregate’’ variables which are required for solving the individual problem. Following Fernández-Villaverde, Hurtado and Nuño (2020), we use neural networks to predict the laws of motion.⁹

The initial guesses for the PLMs are made such that initially inflation is always at the steady

⁹For comparison, we also compute linear regressions of the following form

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

This is 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.

state value (i.e. the inflation target) and the inflation expectation term is such that marginal costs are always at their steady state value, i.e. $m_t = \frac{\varepsilon-1}{\varepsilon}$. Given this, the whole model can be solved with the following approach.

Algorithm 2 (Stochastic Simulation)

1. *Generate and fix a time series of length T for the aggregate shocks.*
2. *Set initial matrices D_i ($i \in \{\pi, E\pi\}$) for the PLMs¹⁰ and an initial distribution of bonds. We use the distribution of bonds in the deterministic steady state and represent this distribution as a histogram as in Young (2010).*
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 (i.e., the aggregate amount of bonds). 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 PLM for the next iteration using the following updating formula*

$$\tilde{D}_i = \lambda_{PLM} \tilde{D}_i + (1 - \lambda_{PLM}) D_i \tag{A.4}$$

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

¹⁰ D_π and $D_{E\pi}$ 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 don't use the prediction of the neural network directly, to be able to slowly update the PLMs.

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