Inflation, Inequality, and the Business Cycle

Clara Lindemann[⊛] Sofia Semik[⋈] Mathias Trabandt[‡]

July 17, 2025

Abstract

We introduce a nonlinear, state-dependent Phillips curve into a standard Heterogeneous Agent New Keynesian (HANK) model. We show that this nonlinearity is crucial for *jointly* matching the empirical properties of inflation and inequality. In our model, inflation and income inequality respond asymmetrically to business cycle fluctuations, increasing more sharply than declining. As a result, the model accounts for the observed positively skewed distributions of U.S. inflation rates and income inequality. In contrast, a version with a constant Phillips curve slope fails to replicate these empirical patterns, underscoring the importance of the nonlinear Phillips curve.

JEL Classification: E31, E32, E52

Keywords: HANK, inequality, nonlinear Phillips Curve, inflation, skewness, business cycle

⊛ Goethe University Frankfurt, Theodor-W.-Adorno-Platz 3, 60323 Frankfurt am Main, Germany, c.lindemann@econ.uni-frankfurt.de

⋈ Goethe University Frankfurt, Theodor-W.-Adorno-Platz 3, 60323 Frankfurt am Main, Germany, semik@em.uni-frankfurt.de

‡ Goethe University Frankfurt, Theodor-W.-Adorno-Platz 3, 60323 Frankfurt am Main, Germany, Halle Institute for Economic Research (IWH), and CEPR, mathias.trabandt@gmail.com

1 Introduction

The inflation surge following the Covid-19 crisis caught many economists and central banks by surprise. Inflation in the U.S. and many other economies rose to levels not observed in decades. A similar surprise occurred after the Global Financial Crisis and Great Recession, when inflation rates fell less than predicted at the time. Recent research has shown that standard linearized New Keynesian models have difficulty explaining inflation developments during deep crises, mirroring the surprises for economists and central banks mentioned above. Harding, Lindé, and Trabandt (2022, 2023) have shown that a nonlinear New Keynesian model with a nonlinear Phillips curve, in which the slope is state-dependent, better accounts for inflation dynamics in deep crises than the linearized model. We contribute to this literature by introducing household heterogeneity into a model with a nonlinear, state-dependent Phillips curve, allowing us to study how these observed inflation dynamics interact with the income distribution.

We include a state-dependent slope of the Phillips curve similar to e.g. Erceg, Jakab, and Lindé (2021) into an otherwise standard nonlinear HANK model similar to e.g. Auclert, Rognlie, and Straub (2024b). This model framework allows us to analyze the implications of a state-dependent Phillips curve slope on the propagation of demand and supply shocks in a heterogeneous agent environment. Importantly, it enables us to study the two-way interaction between inflation and inequality, both during inflation surges and in periods of persistently low inflation. Furthermore, we introduce countercyclical labor and profit income inequality and risk into our model to ensure that our model accounts for the observed volatility in labor and aggregate income inequality in the U.S.

Our results suggest that – due to the state-dependency of the Phillips curve slope – inflationary pressures are amplified, while deflationary pressures are dampened, allowing the model to replicate the volatility and skewness of post-war U.S. inflation, GDP growth, and income inequality. The HANK environment allows us to study how supply and demand shocks propagate into inflation and affect the distribution of income. In particular, we find that inflationary cost-push shocks and contractionary demand shocks raise income inequality by reducing output. Due to countercyclical labor and profit income inequality and income risk, households face a greater risk of receiving less labor income and/or less profit income during recessions. The resulting increase in inequality is amplified by the state-dependent slope of the Phillips curve. Conversely, reductions in inequality in response to cost-pull or expansionary demand shocks are more muted, as the positive output effects of these shocks are dampened relative to a model with a constant Phillips curve slope. Over the business cycle, these results imply that inequality is positively skewed.

Using U.S. data on the standard deviation of log household income – a standard measure of inequality used in the literature – we show that this positive skewness in inequality in the model aligns well with the data. This implies that inequality increases more strongly in recessions than it falls in booms. A HANK model without a state-dependent Phillips curve slope and with Gaussian shocks is not able to replicate these features observed in the data.

A growing strand of literature studies the relationship between inequality, inflation, and monetary policy.[1] Auclert (2019), Bilbiie (2018), and Kaplan, Moll, and Violante (2018) show that inequality affects the transmission of monetary policy in a HANK framework. Auclert et al. (2023) analyze the effects of an inflationary energy price shock and the implications for monetary and fiscal policies in a HANK model. Another strand of literature analyzes the nexus between inequality and optimal monetary policy, see e.g. Acharya, Challe, and Dogra (2023), Bhandari et al. (2021), and McKay and Wolf (2022). Bayer, Born, and Luetticke (2024) estimate a HANK model and show that business cycle fluctuations have implications for inequality. Their findings highlight how business cycle shocks propagate through heterogeneous household portfolios, amplifying inequality dynamics.

We contribute to the above cited literature by explicitly allowing for a nonlinear, state-dependent Phillips curve in a HANK model. Importantly, we focus on the ability of our model to account for the skewness observed in both inflation and income inequality in U.S. data.

Our results are also related to recent empirical evidence on the impact of inflation and monetary policy on inequality. Coibion et al. (2017) find that income and consumption inequality in the U.S. increase in response to contractionary monetary policy shocks. Furceri, Loungani, and Zdzienicka (2018) find empirical evidence that monetary easing and tightening have asymmetric effects on income inequality – similar to our model-implied results. Empirical evidence by Pallotti et al. (2023) suggests that the recent inflation surge has affected households heterogeneously. Evidence from Del Canto et al. (2023) suggests that inflationary cost-push shocks widen the welfare distribution, while expansionary monetary policy shocks tighten it.

The remainder of the paper is organized as follows. In section 2 we introduce our model and the nonlinear, state-dependent Phillips curve. Section 3 describes our model calibration. In section 4 we present our results. Section 5 concludes.

[1] More broadly our paper is related to a by now huge literature on Heterogenous Agent New Keynesian (HANK) models. It beyond the scope of this section to survey or list all contributions in this literature. For a partial body of work, see the following papers and references therein: Acharya and Dogra (2020), Alves et al. (2020), Auclert and Rognlie (2018), Auclert, Rognlie, and Straub (2020), Auclert et al. (2021), Auclert, Rognlie, and Straub (2024a), Bayer et al. (2019), Bilbiie (2019), Bilbiie (2020), Bilbiie, Känzig, and Surico (2022), Bilbiie, Primiceri, and Tambalotti (2023), Bilbiie (2024), Bilbiie, Monacelli, and Perotti (2024), Broer et al. (2020), Debortoli and Galí (2024), Fernández-Villaverde et al. (2023), Hagedorn et al. (2019), Kaplan and Violante (2018), McKay, Nakamura, and Steinsson (2016), McKay and Reis (2016), Moll (2014), and Oh and Reis (2012), among many others.

2 Model

This section develops our nonlinear HANK model. Our model setup is based on the nonlinear version of the canonical HANK model presented by Auclert et al. (2021) and Auclert, Rognlie, and Straub (2024b). It features a nonlinear Phillips curve with a state-dependent slope, Rotemberg nominal rigidities in wage setting, and a central bank that sets the nominal policy rate.

2.1 Households

Households derive utility from consumption $c_{i,t}$ and disutility from supplying labor $n_{i,t}$. They earn income from three sources: labor income $w_t e_{i,t} n_{i,t}$, profit/dividend income $d_{i,t}$, and asset income from holding risk-free assets $r_t a_{i,t-1}$. Household i solves the following maximization problem:

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{c_{i,t}^{1-\sigma}}{1-\sigma} - \varphi \frac{n_{i,t}^{1+\nu}}{1+\nu} \right) \tag{1}$$

$$\text{s.t.} \quad c_{i,t} + a_{i,t} \leq (1-\tau_t)(w_t e_{i,t} n_{i,t} + d_{i,t}) + (1+r_t) a_{i,t-1} \tag{2}$$

$$a_{i,t} \geq \underline{a}. \tag{3}$$

Here, β denotes the discount factor, ν is the inverse Frisch elasticity, σ is the inverse intertemporal elasticity of substitution, and φ scales the disutility from labor. r_t denotes the ex-post net real interest rate. w_t is the real wage, $n_{i,t}$ denotes hours worked and $e_{i,t}$ is idiosyncratic, type-specific productivity. There exist n_e idiosyncratic productivity states. τ_t denotes the income tax rate. Households can buy and sell assets subject to a borrowing constraint given by equation 3. The asset stock is idiosyncratic and type-specific.

Following Auclert and Rognlie (2018), we allow for countercyclical inequality and income risk by introducing the following labor allocation rule:

$$n_{i,t} = n_t \frac{e_{i,t}^{\zeta_n ln(n_t/n)}}{\mathbb{E}\left[e_i^{1+\zeta_n ln(n_t/n)}\right]} \tag{4}$$

where $\zeta_n < 0$. Essentially, this rule generates countercyclical inequality since in a boom, low-productivity households work more than in the steady state and high-productivity households work less than in the steady state - thus implying that inequality declines in a boom. Also, with $\zeta_n < 0$, labor income risk is countercyclical, meaning in a boom, labor income risk falls. Note that if $\zeta_n = 0$, following Auclert, Rognlie, and Straub (2024b), we assume that all households are employed by a union, working the same number of hours, which implies $n_{i,t} = n_t$. Note, too, that $\zeta_n = 0$ implies acyclical inequality and income risk.

Following Debortoli and Galí (2024), we assume that post-tax profit income is distributed among households in proportion to each household's productivity $e_{i,t}$. We also allow for countercyclicality in profit income inequality and risk as above by assuming that aggregate dividends are allocated to households as follows:

$$d_{i,t} = d_t \frac{e_{i,t}^{1+\zeta_d ln(d_t/d)}}{\mathbb{E}\left[e_i^{1+\zeta_d ln(d_t/d)}\right]} \tag{5}$$

where $\zeta_d < 0$ and d_t denotes aggregate firm profits.

We define capital income as the sum of asset and profit/dividend income, as we interpret profit income as dividends paid to households – similar to the definition in Debortoli and Galí (2024).

2.2 Phillips curve

Following Auclert, Rognlie, and Straub (2024b) we assume sticky wages but flexible prices in our model. In models with heterogeneous agents and nominal rigidities, this assumption avoids countercyclical profits and thus large undesirable redistribution effects. Labor unions select the wage rate to maximize household utility and face quadratic nominal wage adjustment costs à la Rotemberg, governed by a adjustment cost parameter ϕ. In Appendix A, we provide a step-by-step derivation of the following wage Phillips curve:

$$\pi_t^w(1+\pi_t^w) = \kappa_t \left(\varphi n_t^{\nu} - \frac{1}{\mu_w}(1-\tau_t) w_t c_t^{-\sigma}\right) n_t + \beta \mathbb{E}_t\left[\pi_{t+1}^w(1+\pi_{t+1}^w)\right] + \epsilon_t, \tag{6}$$

$\mu_w = \frac{\varepsilon}{\varepsilon - 1}$ is the steady state wage markup, and κ_t the state-dependent slope parameter of the Phillips curve. ϵ_t is a cost-push or cost-pull shock. In this paper we use the terms cost-push/cost-pull and supply shocks interchangeably.

Recent work suggests that especially in high and low inflation episodes a nonlinear, 'banana-shaped', Phillips curve is helpful to explain the inflation dynamics in the Great Recession and post-Covid inflation surge. Harding, Lindé, and Trabandt (2022, 2023) use a Kimball aggregator in a nonlinear New Keynesian representative agent (RANK) model to explain the missing deflation puzzle after the financial crisis and the post-Covid inflation surge. Rather than using a fully-fledged nonlinear Kimball aggregator setup – which introduces several more endogenous state variables that render computations in the nonlinear HANK model and interpretation of results more challenging – we follow Erceg, Jakab, and Lindé (2021) and introduce state-dependency of the Phillips curve slope by assuming the slope parameter κ_t to take the following functional form:

$$\kappa_t = \kappa e^{\chi(y_t - y)} \tag{7}$$

with $\chi \geq 0$ determining the endogenous curvature of the Phillips curve and $\kappa = \frac{\varepsilon}{\phi}$. This simple functional form captures key features from richer nonlinear, state-dependent Phillips curve setups while keeping our HANK model as transparent and computationally tractable as possible.

The functional form can be interpreted as follows: when output is at its steady state value y, the slope parameter κ_t becomes state-invariant. When output rises above steady state output, κ_t increases, therefore the Phillips Curve becomes steeper, accounting for the fact that wages are adjusted more strongly when the output gap is positive. When output falls below the steady-state value, κ_t decreases, and the Phillips curve becomes flatter.

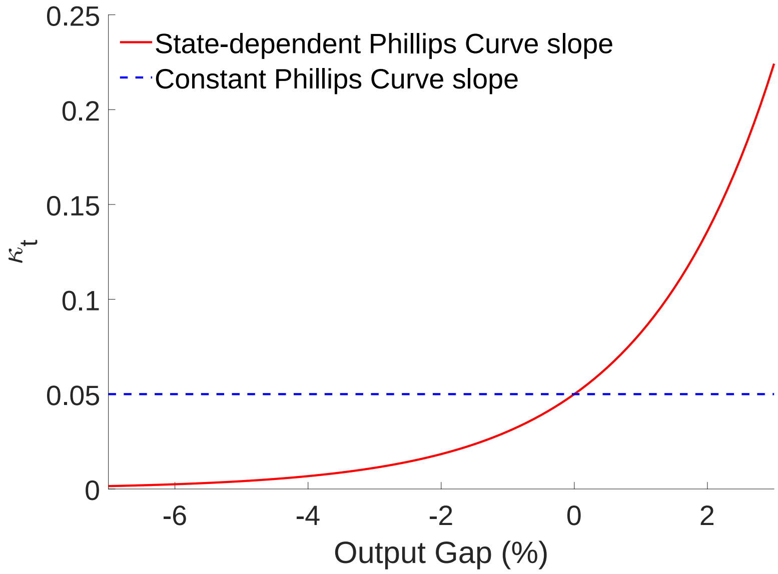

Figure 1: Relationship between output gap and the Phillips curve slope parameter κ_t.

The introduction of a state-dependent κ_t allows us to analyze the impact of high and low inflation on inequality taking into account the results by Harding, Lindé, and Trabandt (2022, 2023).

Figure 1 illustrates the relationship between κ_t and the output gap in our model. We use the following parameters $\kappa = 0.05$ and $\chi = 50$ – see the calibration section below for details how these parameters are obtained. For Figure 1 we set the range for the output gap when calculating κ_t to the minimum and maximum of U.S. CBO output gap data in our sample (1967 - 2019). The minimum output gap in the data is about -7 percent and the maximum output gap in the data is about 3 percent. According to Figure 1, a positive output gap increases the slope of the Phillips curve. Conversely, a negative output gap reduces the slope of the Phillips curve. Thus, the model features a state-dependent slope of the Phillips curve. More importantly, the figure also illustrates that the slope of the Phillips curve is very flat in deep recessions (such as e.g. the Great Recession). Conversely, the slope steepens considerably when actual output is above potential output. Overall, our specification connects to empirical evidence in e.g. Hazell et al. (2022) and Cerrato and Gitti (2025) that there is considerable variation in estimated slopes of the Phillips curve across time and business cycle states. The slope of our state-dependent Phillips curve also aligns with recent cross-country empirical evidence presented by Benigno and Eggertsson (2024).

2.3 Firms

The representative firm produces a continuum of intermediate goods $y_{j,t}$ using labor $n_{j,t}$ according to the following linear production function:

$$y_{j,t} = n_{j,t}. \tag{8}$$

The firm solves the following optimal flexible price setting problem under monopolistic competition:

$$\begin{aligned} \max_{P_{j,t}} D_{j,t} &= P_{j,t} y_{j,t} - W_t n_{j,t} && (9) \\ \text{s.t.} \quad y_{j,t} &= n_{j,t} && (10) \\ y_{j,t} &= \left(\frac{P_{j,t}}{P_t} \right)^{-\theta} y_t, && (11) \end{aligned}$$

where $D_{j,t}$ are nominal profits, W_t is the nominal wage rate, $P_{j,t}$ is the price of good j and $\theta > 1$ is the substitution elasticity between intermediate goods. Equation (11) is the demand for good j. In equilibrium, all firms charge the same price, i.e. $P_{j,t} = P_t$, so that the optimal price setting equation for firms can be expressed as:

$$P_t = \frac{\theta}{\theta - 1} W_t. \tag{12}$$

Thus, firms charge a constant mark-up $\mu_p = \frac{\theta}{\theta - 1}$ over their marginal cost (nominal wages). The real wage is therefore given by $w_t = \frac{1}{\mu_p}$. Real profits are given by

$$d_t = y_t - w_t y_t = \left(1 - \frac{1}{\mu_p} \right) y_t. \tag{13}$$

Note that real profits are procyclical in our model.

Finally, goods inflation π_t and wage inflation π_t^w are related as follows:

$$1 + \pi_t = (1 + \pi_t^w) \left(\frac{w_{t-1}}{w_t} \right). \tag{14}$$

2.4 Monetary and fiscal policy

We assume that the central bank follows a standard Taylor rule to set the nominal interest rate i_t:

$$\frac{1+i_t}{1+i} = \left(\frac{1+\pi_t}{1+\pi}\right)^{\phi_\pi} \left(\frac{y_t}{\tilde{y}_t}\right)^{\phi_y} e^{\gamma_t}, \tag{15}$$

where $\tilde{y}_t$ is potential output, which takes the value of steady-state output y in our model. i and π are the steady-state nominal interest rate and inflation rate, respectively. ϕ_π and ϕ_y denote the Taylor rule parameters on inflation and the output gap. γ_t denotes an exogenous monetary policy shock which – following a large body of work in the HANK literature – we take as a stand-in for a demand shock affecting the economy.

The government budget constraint is given by:

$$\tau_t y_t + b_t = (1+r_t)b_{t-1} + g, \tag{16}$$

where b_t is government debt and g is government consumption spending. We assume that government consumption spending is constant. We assume that the government adjusts the tax rate τ_t period-by-period to balance its budget.

2.5 Aggregation

The aggregate resource constraint and asset market clearing condition are given by, respectively:

$$y_t = c_t + g \tag{17}$$

$$b_t = \int_0^1 a_{i,t} di. \tag{18}$$

2.6 Shocks

For the monetary policy shock γ_t and the cost-push shock ε_t, we specify the following AR(1) processes:

$$\gamma_t = \rho_\gamma \gamma_{t-1} + \eta_t^\gamma, \tag{19}$$

$$\epsilon_t = \rho_\epsilon \epsilon_{t-1} + \eta_t^\epsilon, \tag{20}$$

where $\eta_t^\gamma \sim \mathcal{N}(0, \sigma_\gamma^2)$ and $\eta_t^\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$ are exogenous disturbances.

3 Calibration

We calibrate our HANK model to the following thirteen moments of key macroeconomic U.S. time series: the standard deviation, skewness, and autocorrelation of inflation and GDP growth, and the correlation between inflation and GDP growth, and the mean, standard deviation and skewness of labor income inequality and labor income plus capital income inequality. We set model parameters so that the model-implied moments are close to the mean data moments and within 95% data confidence intervals.[2]

3.1 Data

To match our model to U.S. macro data, we use U.S. data on annualized PCE inflation excluding food and energy and annualized GDP growth from 1967Q1 to 2019Q4. Figure 2 illustrates the evolution of quarterly inflation and GDP growth over the period considered. The inflation dynamics show that inflation surges are larger than inflation declines, resulting in positive skewness.

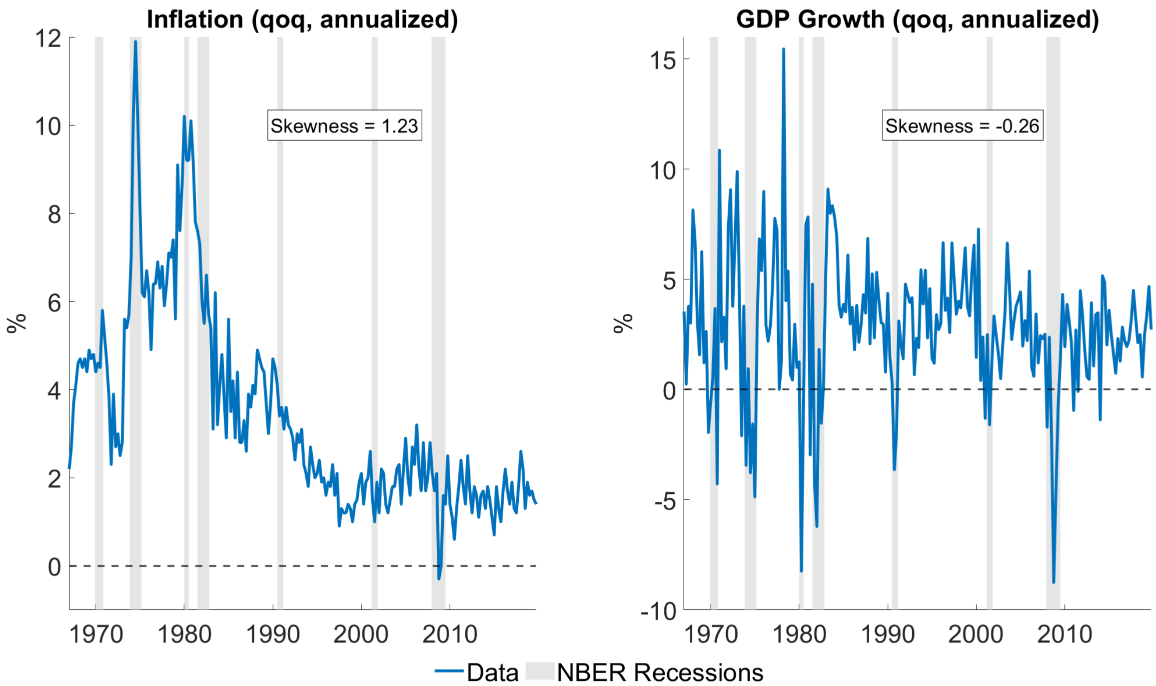

Figure 2: Quarterly U.S. inflation and GDP growth time series 1967-2019.

To measure income inequality in the U.S., we construct an inequality measure using data on household-level labor and capital income from the Current Population Survey

[2]We adopt a calibration approach, rather than a formal moment-based model estimation approach due to computational feasibility constraints. Specifically, for a given set of parameters, it takes about 2-3 seconds to solve our nonlinear HANK model for aggregate and distributional variables. To then compute model-implied moments, we need to simulate our model and each simulation with a sample size of 1,000 quarters takes approximately 45 minutes. Thus, it is unfortunately computationally infeasible to estimate model parameters.

(CPS). Following Heathcote et al. (2023) we rely on household data from the Annual Social and Economic (ASEC) supplement of the CPS from 1967 to 2019. We exclude households with zero or negative ASEC weight and households with no reference person or with no household member between 25 and 60 years of age. The labor income measure is constructed as the sum of wage earnings and income from self employment, divided by the number of adult equivalents in the household. Note that this is labor income before taxes and transfers; it is comparable to factor income as used in the HANK literature (see e.g. Bilbiie (2024)). The total income inequality measure includes capital income like interest income or dividends in addition to labor income.

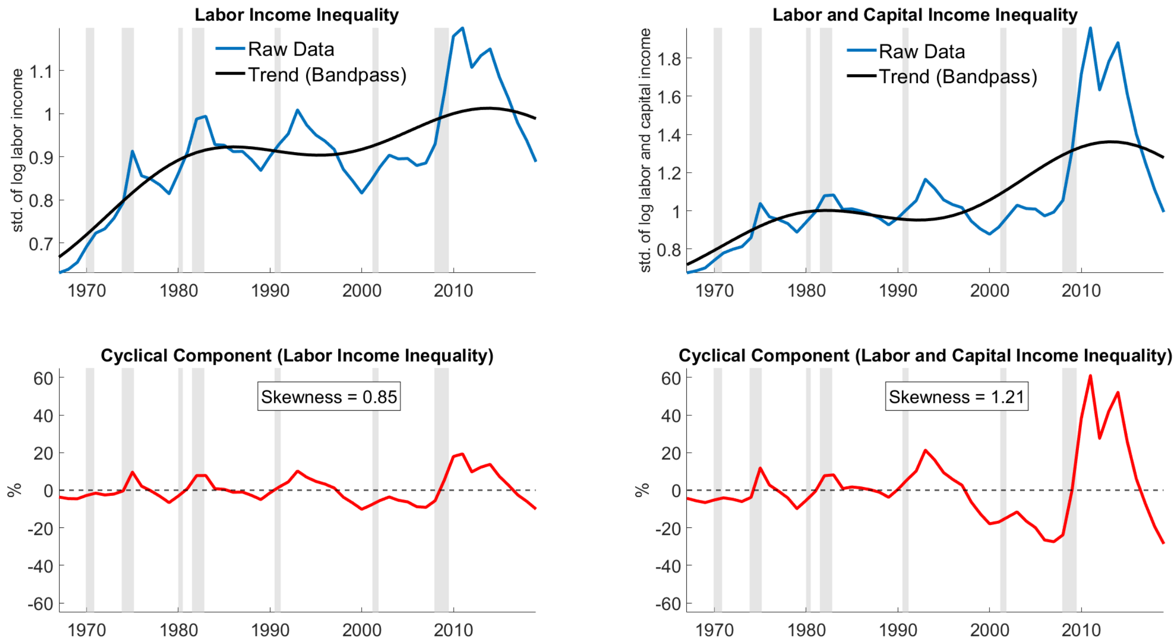

Figure 3: Cross-sectional standard deviation of log labor income and log labor + capital income (annual).

We use the standard deviation of log income to measure inequality in our sample. For each year in the sample, we discretize the income distribution using percentiles. We then take the natural logarithm of these percentiles and compute the weighted cross-sectional standard deviation in each year. Note that we exclude the lowest nine percentiles for labor income and the lowest seven percentiles for aggregate income, because income is equal to zero for them in some years. To still properly represent the bottom income percentiles, the tenth (eighth) percentile is weighted by factor ten (eight), while all higher percentiles are weighed by factor one. Note that the inequality data show a trend over time. Since we are interested to study business cycle dynamics in this paper, we detrend the inequality data. Following Bilbiie, Primiceri, and Tambalotti (2023) we use a band-pass filter that extracts fluctuations with periodicities less than or equal to 30 years. Figure 3 shows the raw data, trends, and business cycle components of the inequality data. Figure 3 shows that: i) inequality is countercyclical, i.e. it increases in recessions, ii) inequality is highly volatile, and iii) inequality is skewed, i.e. over the business cycle, it rises more than it falls, on average. In Appendix B, we show that these empirical features of inequality data are robust to alternative detrending methods.

3.2 Parameters

We adopt a quarterly calibration and set the parameter values as follows. The inverse elasticity of intertemporal substitution and the inverse Frisch elasticity are set to $\sigma = 1$ and $\nu = 2$, respectively. Further, we set the steady-state gross wage mark-up to $\mu_w = 1.1$. As a result, the elasticity of substitution is equal to $\varepsilon = 11$. We set the labor disutility parameter $\varphi = 1/\mu_w = 0.909$ such that steady-state labor and output are unity ($n = y = 1$, normalization). The Rotemberg wage adjustment cost parameter is set to $\phi = 208$. This parameter value is equivalent (in a linearized model version) to a Calvo wage stickiness parameter value that implies wage changes once every 1.5 years on average. Given the values of ε and ϕ, the steady-state slope of the Phillips curve is $\kappa = \frac{\varepsilon}{\phi} = 0.05$. We set the curvature parameter of the Phillips curve slope to $\chi = 50$ to match the positive skewness of inflation observed in the data. The steady-state gross price mark-up is set to $\mu_p = 1.2$. The Taylor rule parameters are set to $\phi_\pi = 1.5$ and $\phi_y = 0.2$, respectively. Steady-state government consumption spending is set to 20% of GDP. The steady-state net inflation rate is set to $\pi = 0$. The quarterly steady-state real interest rate is 1/2 of a percent, i.e. $r = 0.005$. The annual steady-state government debt-to-GDP is set to 70%. With these parameter values, asset market clearing in steady state results in a value for the household discount factor β of roughly 0.98.

We simulate our model using exogenous AR(1) processes: γ_t (demand) and ε_t (cost-push). We set the persistence parameters for both AR(1) processes to $\rho_\gamma = 0.9$ and $\rho_\varepsilon = 0.9$. Stochastic shocks to both AR(1) processes are assumed to follow $\epsilon_t^\gamma \sim \mathcal{N}(0, \sigma_\gamma^2)$ and $\epsilon_t^\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$, where the standard deviations are set to $100\sigma_\gamma = 0.0470$ and $100\sigma_\varepsilon = 0.0784$.

The heterogeneity in our model stems from heterogeneous idiosyncratic productivity states and heterogeneous asset holdings. There exist n_a idiosyncratic asset holding states, i.e. grid points on the asset grid, and n_e idiosyncratic productivity states. Following Auclert et al. (2021), the number of grid points is set to $n_a = 500$ for the asset distribution and $n_e = 11$ for the productivity grid. For calibration of the idiosyncratic income process we follow Auclert et al. (2021), with the values modified to be consistent with a quarterly frequency. Specifically, the 11-state Markov chain of idiosyncratic productivity has an implied AR(1) representation with the autocorrelation parameter equal to 0.98 and the standard deviation parameter of the innovations of 0.92. To match the moments of U.S. income inequality in the data, we set the cyclical labor income risk parameter to $\zeta_n = -4$ and the cyclical profit income risk parameter to $\zeta_d = -10$. This implies that both labor and profit income inequality and risk are countercyclical.

We solve the nonlinear HANK model using the Sequence-Space Jacobian software package developed by Auclert et al. (2021).[3] The simulation results in Table 1 and 2 are generated by a long model simulation over 1000 quarters, where random unexpected demand and cost-push shocks are drawn every quarter. It takes about 45 minutes to simulate our nonlinear model to obtain the model-implied moments. Comparing the model-implied moments with those in the data reveals that overall, the model accounts reasonably well for the moments in the data. We discuss further details in the next section.

4 Results

In this section, we report our results for the model simulation. First, we study how state dependency in the Phillips curve slope affects the propagation of shocks in our model. Section 4.1 compares impulse responses to a cost-push and a demand shock in our model to a model with a constant Phillips curve slope. Section 4.2 presents the results from a model simulation of randomly drawn demand and supply shocks to assess the ability of the model to match the features from the data discussed in Section 3.1.

4.1 Propagation of Shocks

4.1.1 Cost-push Shocks

The first column of Figure 4 shows the impulse responses of the nonlinear, state-dependent Phillips curve model to small and large cost-push shocks. The Figure reveals that the state-dependent Phillips curve generates asymmetries in the economy's response to cost-push shocks when these shocks are sizeable. Specifically, introducing a state-dependent slope parameter κ_t into the Phillips curve has no significant impact for small cost-push or cost-pull shocks, represented by the orange solid and green dotted lines. This is due to the approximate linearity of the Phillips curve close to the steady state (see Figure 1). However, for larger shocks the responses following equal-sized positive vs. negative cost-push shocks become increasingly asymmetric. In particular, the inflation increase following an adverse supply shock is amplified, whereas the decline after an equal-sized favorable supply shock is dampened.

When an adverse cost-push shock raises inflation, the stronger inflationary response induces the central bank to raise nominal interest rates more aggressively. This leads to a larger contraction in output compared to the model featuring a constant Phillips curve. This amplified drop in output exacerbates income inequality, as agents face higher countercyclical labor and profit inequality and income risk.

[3]Python replication codes are available on the authors' websites.

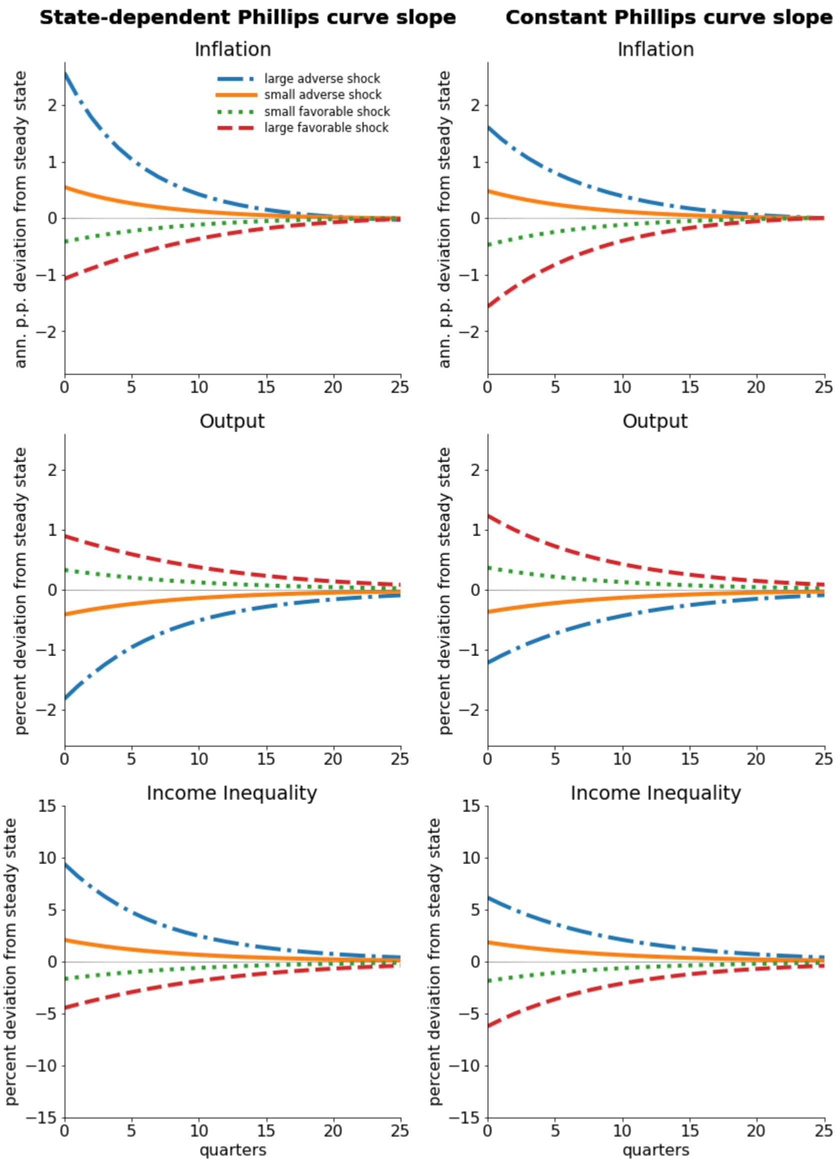

Figure 4: Impulse responses following a cost-push shock in our nonlinear HANK model.

Conversely, following a favorable cost-push shock (cost-pull shock), inflation falls, prompting the central bank to lower interest rates. However, since inflation declines less than in the model featuring a constant Phillips curve, the increase in output is also more muted. Consequently, the reduction in income inequality is less pronounced.

The second column of Figure 4 shows that using a constant Phillips curve slope, the impact of small and large cost-push shocks is symmetric, meaning that equal-sized positive and negative cost-push shocks cause equal-sized increases and decreases in inflation, output, and inequality.

4.1.2 Demand Shocks

The first column of Figure 5 shows the implications of a nonlinear, state-dependent Phillips curve slope for the propagation of demand shocks in the model. The Figure reveals that – similar to the case of cost-push shocks – the responses of inflation, output. and inequality are asymmetric when demand shocks become sizeable.

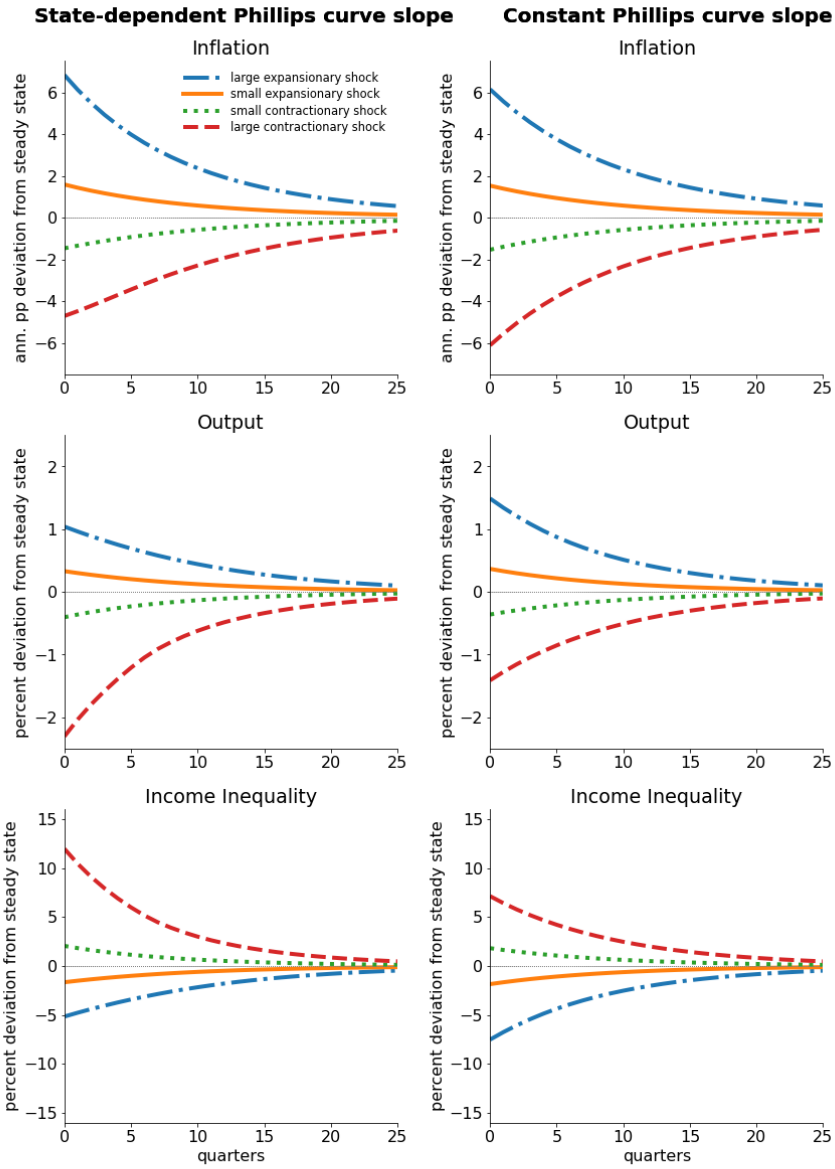

Figure 5: Impulse responses following a demand shock in our nonlinear HANK model.

The intuition for the asymmetries in the responses is that wages are adjusted more strongly when the economy is in a boom, i.e. when the output gap is positive, which leads to higher inflation. This dampens the upward adjustment in output following a positive demand shock. In a recession, i.e. when the output gap is negative, the drop in real wages is attenuated due to the nonlinearity of the Phillips curve slope, which in turn dampens the drop in inflation. In this case, the decrease in output is amplified as the central bank decreases the nominal interest rate less, as inflation falls less, and therefore the positive second-round effect of a lower real interest rate is smaller, and output decreases more strongly.

The response of the standard deviation of log income shows that this asymmetric response to shocks also carries over to our measure of inequality. Consistent with our empirical analysis, we use pre-tax income to construct our measure of inequality. Following a negative demand shock, output and inflation decrease. This decrease in output leads to a stronger dispersion of labor productivity states, as agents face countercyclical inequality and income risk. A negative demand shock thus increases the risk of becoming less productive, which leads to an increase in inequality. Similarly, the fall in output increases the risk of having lower profit income due to countercyclical profit income inequality and risk, which increases capital income inequality. In the nonlinear model, the fall in output is amplified due to the state-dependency of the Phillips curve discussed above. This, in turn, exacerbates the increase in income inequality. Conversely, following a positive demand shock, output, and inflation increase. Therefore, income inequality decreases, as agents face lower risk of becoming less productive and having lower profit income. As the increase in output is dampened in our nonlinear model, the decrease in inequality is dampened as well.

The second column of Figure 5 shows that using a constant Phillips curve slope, the impact of small and large demand shocks is symmetric, meaning that equal-sized positive and negative demand shocks cause equal-sized increases and decreases in inflation, output, and inequality.

All told, the impulse responses to cost-push and demand shocks reveal that inflation and inequality increase by more than they decrease, rendering the model potentially capable of accounting for the skewness of inflation and inequality observed in the data.

4.2 Model vs. Data Comparison

Table 1 compares the results from our model simulation of randomly drawn demand and cost-push shocks over 1000 quarters to the moments observed in the data. Overall, the nonlinear model with the state-dependent Phillips curve slope matches the standard deviations, skewness, and (auto-)correlations of inflation and GDP growth reasonably well. In contrast, the model with a constant Phillips curve slope matches the data considerably worse, especially the positive skewness of inflation.[4]

Figure 6 plots the simulation results of inflation, output, and inequality for a sample of 1000 consecutive demand and cost-push shocks in our model. The figure also shows

[4] One way to improve the fit of the constant Phillips curve slope is to increase the size of the shocks. Although this allows the constant Phillips curve slope model to better match the standard deviations of inflation, output, and inequality, that model still suffers from failing to account for the observed positive skewness in inflation and inequality.

Table 1: Model vs. Data Comparison

	Model		**Data**		
	Phillips curve slope				
	State-dependent	*Constant*	Mean	95% CI	
Standard deviation π_t	2.31	1.51	2.29	2.00	2.55
Skewness π_t	1.41	-0.17	1.23	0.93	1.53
Autocorrelation π_t	0.89	0.92	0.92	0.89	0.94
Standard deviation $\triangle y_t$	3.09	1.89	3.14	2.71	3.58
Skewness $\triangle y_t$	0.17	0.03	-0.26	-0.93	0.49
Autocorrelation $\triangle y_t$	-0.02	-0.01	0.30	0.15	0.45
Correlation π_t, $\triangle y_t$	-0.23	-0.19	-0.06	-0.25	0.14

the results using the same shocks in a version of our model with a constant Phillips curve slope. It shows that inflation surges are almost twice as large when the slope of the Phillips curve is state-dependent compared to the constant slope case. By contrast, deflationary pressures are dampened, at least to some extent. The asymmetry of inflation dynamics is more pronounced for larger shocks.

In addition, Figure 6 shows that when the economy is hit by supply and demand shocks, inequality tends to increase more strongly than it decreases. The standard deviation of log household income increases by almost twice as much in the state-dependent Phillips curve slope model compared to a model with a constant Phillips curve slope, while the reductions in boom periods are dampened. The decrease in output induces a stronger dispersion in productivity states, which leads to an increase in income inequality. Finally, Figure 6 shows that the increase in aggregate income inequality is even more pronounced in recessions due to countercyclical profit income inequality and risk.

Table 2 shows that the nonlinear, state-dependent Phillips curve slope allows the model to reproduce reasonably well the positive skewness of both labor and aggregate household income inequality observed in the data. In other words, income inequality increases more strongly in recessions than it falls in boom periods. A model with a constant Phillips curve slope predicts a skewness close to zero, which implies a symmetric response of inequality to equal-sized favorable and adverse shocks. The long simulation results show that our model is able to align closely with the empirical moments of labor income and labor and capital income inequality.[5]

In our framework, countercyclical labor and profit income inequality and risk, governed by the parameters ζ_n and ζ_d, are crucial to account for the pronounced volatility of inequality observed in the data. Countercyclical labor income risk implies that households are more likely to experience declines in productivity during recessions. This feature

[5] If anything, our model tends to slightly overshoot the skewness of labor income inequality relative to the 95 percent confidence interval of the data.

is intended to capture the elevated risk of unemployment during economic downturns, which may be an important driver of inequality fluctuations over the business cycle (see Chang and Schorfheide (2023)). Similarly, countercyclical profit income inequality and risk serve as a proxy for a risky asset, allowing for greater volatility in capital income inequality without requiring a full two-asset structure.

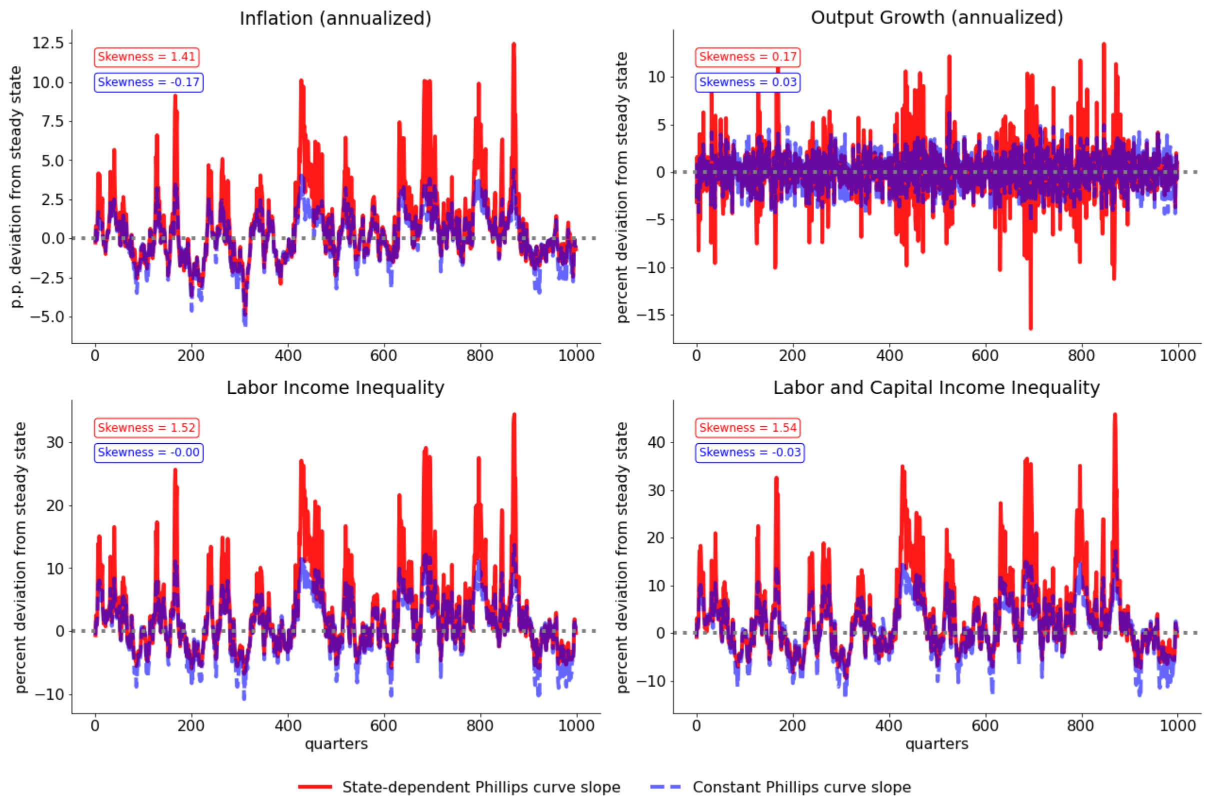

Figure 6: Long simulation in our nonlinear HANK model over 1000 quarters with demand and cost-push shocks.

Table 2: Income Inequality (std. of log income): Data vs. Model Comparison

	Model		**Data**		
	Phillips curve slope				
	State-dependent	*Constant*	Mean	95% CI	
Labor Income Inequality					
Mean	0.92	0.92	0.92	0.88	0.95
Skewness	1.51	0.00	0.85	0.32	1.34
Standard Deviation	0.06	0.04	0.07	0.05	0.08
Labor and Capital Income Inequality					
Mean	0.94	0.94	1.07	0.99	1.15
Skewness	1.55	-0.03	1.15	0.40	1.77
Standard Deviation	0.08	0.05	0.18	0.13	0.23

5 Conclusion

We introduce a nonlinear Phillips curve with a state-dependent slope into a standard Heterogeneous Agent New Keynesian (HANK) model. This modification enables our model to replicate the positive skewness of inequality and inflation observed in post-war U.S. data, a feature that a model with a constant Phillips curve slope fails to capture. Our results indicate that output declines lead to rising income inequality due to countercyclical labor and profit income inequality and risk. In response to demand and supply shocks, the state-dependent Phillips curve amplifies these effects, exacerbating inequality increases in recessions while dampening its reduction during booms.

References

Acharya, Sushant, Edouard Challe, and Keshav Dogra (2023). "Optimal monetary policy according to HANK". In: *American Economic Review* 113.7, pp. 1741–1782.

Acharya, Sushant and Keshav Dogra (2020). "Understanding HANK: Insights from a PRANK". In: *Econometrica* 88.3, pp. 1113–1158.

Alves, Felipe, Greg Kaplan, Benjamin Moll, and Giovanni L Violante (2020). "A further look at the propagation of monetary policy shocks in HANK". In: *Journal of Money, Credit and Banking* 52.S2, pp. 521–559.

Auclert, Adrien (2019). "Monetary policy and the redistribution channel". In: *American Economic Review* 109.6, pp. 2333–2367.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub (2021). "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models". In: *Econometrica* 89.5, pp. 2375–2408.

Auclert, Adrien, Hugo Monnery, Matthew Rognlie, and Ludwig Straub (2023). *Managing an Energy Shock: Fiscal and Monetary Policy*. Tech. rep. Working Paper.

Auclert, Adrien and Matthew Rognlie (2018). *Inequality and aggregate demand*. Tech. rep. National Bureau of Economic Research.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2020). *Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model*. Tech. rep. National Bureau of Economic Research.

— (2024a). *Fiscal and monetary policy with heterogeneous agents*. Tech. rep. National Bureau of Economic Research.

— (2024b). "The intertemporal keynesian cross". In: *Journal of Political Economy* 132.12, pp. 4068–4121.

Bayer, Christian, Benjamin Born, and Ralph Luetticke (2024). "Shocks, frictions, and inequality in US business cycles". In: *American Economic Review* 114.5, pp. 1211–1247.

Bayer, Christian, Ralph Luetticke, Lien Pham-Dao, and Volker Tjaden (2019). "Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk". In: *Econometrica* 87.1, pp. 255–290.

Benigno, Pierpaolo and Gauti B. Eggertsson (May 2024). "The Slanted-L Phillips Curve". In: *AEA Papers and Proceedings* 114, pp. 84–89. DOI: 10.1257/pandp.20241051. URL: https://www.aeaweb.org/articles?id=10.1257/pandp.20241051.

Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J Sargent (2021). "Inequality, Business Cycles, and Monetary-Fiscal Policy". In: *Econometrica* 89.6, pp. 2559–2599.

Bilbiie, Florin (2018). "Monetary policy and heterogeneity: An analytical framework". In: *CEPR Discussion Paper No. DP12601*.

Bilbiie, Florin, Giorgio Primiceri, and Andrea Tambalotti (2023). *Inequality and business cycles*. Tech. rep. National Bureau of Economic Research.

Bilbiie, Florin O (2019). "Optimal forward guidance". In: *American Economic Journal: Macroeconomics* 11.4, pp. 310–345.

— (2020). "The new Keynesian cross". In: *Journal of Monetary Economics* 114, pp. 90–108.

— (2024). "Monetary policy and heterogeneity: An analytical framework". In: *Review of Economic Studies*, rdae066.

Bilbiie, Florin O, Diego R Känzig, and Paolo Surico (2022). "Capital and income inequality: An aggregate-demand complementarity". In: *Journal of Monetary Economics* 126, pp. 154–169.

Bilbiie, Florin O, Tommaso Monacelli, and Roberto Perotti (2024). "Stabilization vs. Redistribution: The optimal monetary–fiscal mix". In: *Journal of Monetary Economics* 147, p. 103623.

Broer, Tobias, Niels-Jakob Harbo Hansen, Per Krusell, and Erik Öberg (2020). "The New Keynesian transmission mechanism: A heterogeneous-agent perspective". In: *The Review of Economic Studies* 87.1, pp. 77–101.

Cerrato, Andrea and Giulia Gitti (2025). "The Return of the Phillips Curve: Evidence from US Cities". In.

Chang, Minsu and Frank Schorfheide (2023). *On the Effets of Monetary Policy Shocks on Income and Consumption Heterogeneity*. Tech. rep. No 2877. NBER Working Paper.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia (2017). "Innocent Bystanders? Monetary policy and inequality". In: *Journal of Monetary Economics* 88, pp. 70–89.

Debortoli, Davide and Jordi Galí (2024). *Heterogeneity and aggregate fluctuations: Insights from tank models*. Tech. rep. National Bureau of Economic Research.

Del Canto, Felipe N, John R Grigsby, Eric Qian, and Conor Walsh (2023). *Are Inflationary Shocks Regressive? A Feasible Set Approach.* Tech. rep. National Bureau of Economic Research.

Erceg, Christopher J., Zoltan Jakab, and Jesper Lindé (2021). "Monetary policy strategies for the European Central Bank". In: *Journal of Economic Dynamics and Control* 132.104211.

Fernández-Villaverde, Jesús, Joël Marbet, Galo Nuño, and Omar Rachedi (2023). *Inequality and the zero lower bound.* Tech. rep. National Bureau of Economic Research.

Furceri, Davide, Prakash Loungani, and Aleksandra Zdzienicka (2018). "The effects of monetary policy shocks on inequality". In: *Journal of International Money and Finance* 85, pp. 168–186.

Hagedorn, Marcus, Jinfeng Luo, Iourii Manovskii, and Kurt Mitman (2019). "Forward guidance". In: *Journal of Monetary Economics* 102, pp. 1–23.

Hamilton, James D (2018). "Why you should never use the Hodrick-Prescott filter". In: *Review of Economics and Statistics* 100.5, pp. 831–843.

Harding, Martin, Jesper Lindé, and Mathias Trabandt (2022). "Resolving the missing deflation puzzle". In: *Journal of Monetary Economics* 126, pp. 15–34.

— (2023). "Understanding post-coving inflation dynamics". In: *Journal of Monetary Economics.*

Hazell, Jonathon, Juan Herreno, Emi Nakamura, and Jón Steinsson (2022). "The slope of the Phillips Curve: evidence from US states". In: *The Quarterly Journal of Economics* 137.3, pp. 1299–1344.

Heathcote, Jonathan, Fabrizio Perri, Giovanni L Violante, and Lichen Zhang (2023). "More unequal we stand? Inequality dynamics in the United States, 1967–2021". In: *Review of Economic Dynamics* 50, pp. 235–266.

Kaplan, Greg, Benjamin Moll, and Giovanni Violante (2018). "Monetary policy according to HANK". In: *American Economic Review* 108.3, pp. 697–743.

Kaplan, Greg and Giovanni L Violante (2018). "Microeconomic heterogeneity and macroeconomic shocks". In: *Journal of Economic Perspectives* 32.3, pp. 167–194.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson (2016). "The power of forward guidance revisited". In: *American Economic Review* 106.10, pp. 3133–3158.

McKay, Alisdair and Ricardo Reis (2016). "The role of automatic stabilizers in the US business cycle". In: *Econometrica* 84.1, pp. 141–194.

McKay, Alisdair and Christian Wolf (2022). *Optimal policy rules in HANK.* Tech. rep. Working Paper, FRB Minneapolis.

Moll, Benjamin (2014). "Productivity losses from financial frictions: Can self-financing undo capital misallocation?" In: *American Economic Review* 104.10, pp. 3186–3221.

Oh, Hyunseung and Ricardo Reis (2012). "Targeted transfers and the fiscal response to the great recession". In: *Journal of Monetary Economics* 59, S50–S64.

Pallotti, Filippo, Gonzalo Paz-Pardo, Jiri Slacalek, Oreste Tristani, and Giovanni L. Violante (2023). "Who bears the cost of inflation? Euro area households and the 2021-2022 shock". In: *ECB Working Paper Series* No 2877.

Ravn, Morten O and Harald Uhlig (2002). "On adjusting the Hodrick-Prescott filter for the frequency of observations". In: *Review of economics and statistics* 84.2, pp. 371–376.

Rotemberg, Julio J (1982). "Sticky prices in the United States". In: *Journal of political economy* 90.6, pp. 1187–1211.

A Derivation of the Wage Phillips Curve

In this section, we derive the nonlinear wage Phillips curve given by equation 6 in the main text. In our set-up, unions face quadratic nominal wage adjustment costs à la Rotemberg (1982). At time t, union j sets its wage $W_{j,t}$ to maximize the utility of its average worker as in Auclert, Rognlie, and Straub (2024b). Note that, following Auclert, Rognlie, and Straub (2024a), this implies that the union does not take the labor allocation rule in equation 4 into account when setting the wage. The maximization problem is defined as follows:

$$\max_{W_{j,t}} \sum_{s=0}^{\infty} \mathbb{E}_t \left[\frac{c_t^{1-\sigma}}{1-\sigma} - \varphi \frac{n_t^{1+\nu}}{1+\nu} - \frac{\phi}{2} \left(\frac{W_{j,t+s}}{W_{j,t+s-1}} - 1 \right)^2 \right], \tag{21}$$

$$\text{s.t.} \qquad n_{j,t} = \left(\frac{W_{j,t}}{W_t} \right)^{-\varepsilon} n_t. \tag{22}$$

Unions combine individual labor into tasks, which face demand given by (22).

Using (22), household post-tax real earnings are defined as follows:

$$z_t = (1-\tau_t) \int_0^1 \frac{W_{j,t}}{P_t} n_{j,t} dj = (1-\tau_t) \frac{1}{P_t} \int_0^1 W_{j,t} \left(\frac{W_{j,t}}{W_t} \right)^{-\varepsilon} n_t dj \tag{23}$$

We assume that all income from the union wage change is consumed immediately, which implies $\frac{\partial c_t}{\partial W_{j,t}} = \frac{\partial z_t}{\partial W_{j,t}}$ by the envelope theorem:

$$\frac{\partial c_t}{\partial W_{j,t}} = \frac{\partial z_t}{\partial W_{j,t}} = (1-\tau_t)(1-\varepsilon) \frac{1}{P_t} \left(\frac{W_{j,t}}{W_t} \right)^{-\varepsilon} n_t = (1-\tau_t)(1-\varepsilon) \frac{1}{P_t} n_{j,t} \tag{24}$$

The derivative of hours worked by household i (from equation 22) with respect to wage $W_{j,t}$ is given by:

$$\frac{\partial n_{i,t}}{\partial W_{j,t}} = -\varepsilon \frac{n_{j,t}}{W_{j,t}} \tag{25}$$

Using (24) and (25), we obtain the following first-order condition of the union:

$$c_t^{-\sigma}(1-\tau_t)(1-\varepsilon)\frac{1}{P_t}n_{j,t} + \varepsilon\varphi n_t^{\nu}\frac{n_{j,t}}{W_{j,t}} - \phi\frac{1}{W_{j,t-1}}\left(\frac{W_{j,t}}{W_{j,t-1}}-1\right) + \beta\phi\mathbb{E}_t\frac{W_{j,t+1}}{W_{j,t}^2}\left(\frac{W_{j,t+1}}{W_{j,t}}-1\right) = 0 \tag{26}$$

In equilibrium all unions set the same wage, which implies $W_{j,t} = W_t$ and $n_{j,t} = n_t$:

$$c_t^{-\sigma}(1-\tau_t)(1-\varepsilon)\frac{1}{P_t}n_t + \varepsilon\varphi n_t^{\nu}\frac{n_t}{W_t} - \phi\frac{1}{W_{t-1}}\left(\frac{W_t}{W_{t-1}}-1\right) + \beta\phi\mathbb{E}_t\frac{W_{t+1}}{W_t^2}\left(\frac{W_{t+1}}{W_t}-1\right) = 0 \tag{27}$$

Define wage inflation such that $\pi_t^w = \frac{W_t}{W_{t-1}} - 1$:

$$c_t^{-\sigma}(1-\tau_t)(1-\varepsilon)\frac{1}{P_t}n_t + \varepsilon\varphi n_t^{\nu}\frac{n_t}{W_t} - \phi\frac{1}{W_{t-1}}\pi_t^w + \beta\phi\mathbb{E}_t\frac{1}{W_t}(\pi_{t+1}^w+1)\pi_{t+1}^w = 0 \tag{28}$$

$$\Leftrightarrow \quad c_t^{-\sigma}(1-\tau_t)(1-\varepsilon)w_t n_t + \varepsilon\varphi n_t^{\nu} n_t - \phi(\pi_t^w+1)\pi_t^w + \beta\phi\mathbb{E}_t(\pi_{t+1}^w+1)\pi_{t+1}^w = 0, \tag{29}$$

where $w_t = \frac{W_t}{P_t}$ is the real wage.

Finally, this can be rearranged such that we obtain our nonlinear wage Phillips curve:

$$\pi_t^w(1+\pi_t^w) = \frac{\varepsilon}{\phi}\left(\varphi n_t^{\nu} - \frac{\varepsilon-1}{\varepsilon}(1-\tau_t)w_t c_t^{-\sigma}\right)n_t + \beta\mathbb{E}_t\left[\pi_{t+1}^w(1+\pi_{t+1}^w)\right]. \tag{30}$$

In our model we then introduce the state-dependent slope parameter $\kappa_t = \frac{\varepsilon}{\phi}e^{\chi(y_t-y)}$, which results in the following expression for our nonlinear Phillips curve in the main text:

$$\pi_t^w(1+\pi_t^w) = \kappa_t n_t\left(\varphi n_t^{\nu} - \frac{\varepsilon-1}{\varepsilon}(1-\tau_t)w_t c_t^{-\sigma}\right) + \beta\mathbb{E}_t\left[\pi_{t+1}^w(1+\pi_{t+1}^w)\right]. \tag{31}$$

B Robustness of De-trending Method

Table 3: Data moments income inequality (std. of logs) with different time series filters

	Band-pass 30y (baseline)	Band-pass 8y	HP filter	Hamilton
	Labor Income Inequality			
Skewness	0.76	0.58	0.52	1.31
Standard Deviation	0.05	0.03	0.03	0.07
	Labor and Capital Income Inequality			
Skewness	1.15	0.49	0.74	2.09
Standard Deviation	0.19	0.08	0.07	0.18

Table 3 presents the skewness and standard deviation of the standard deviation of log labor income (LI) and log labor plus capital income (LCI) in the dataset described in section 3.1 using four different de-trending methods to check the robustness of using the band-pass filter to extract fluctuations with periodicities lower than 30 years. First, we use a band-pass filter with an upper bound of 8 years, as often used for business cycle analysis. This yields a smaller, but still positive, skewness and a smaller standard deviation. Next, we use a two-sided Hodrick-Prescott (HP) filter, where we set the smoothing parameter to $\lambda = 6.25$, as is suggested for annual data by Ravn and Uhlig (2002). The results for the two-sided HP filter are close to our baseline results. Finally, we use a Hamilton filter (Hamilton (2018)), where we set the lead length to 2 and the lag length to 1, as recommended for annual data. The resulting cyclical component shows almost the same volatility as in the baseline analysis, but an even higher skewness. On average, all de-trending methods yield a skewness of 0.79 for labor income and 1.12 for labor and capital income.