



HAL
open science

Inequality and optimal monetary policy in the open economy

Sushant Acharya, Edouard Challe

► To cite this version:

Sushant Acharya, Edouard Challe. Inequality and optimal monetary policy in the open economy. *Journal of International Economics*, 2025, 155, pp.104076. <10.1016/j.jinteco.2025.104076>. <halshs-05148791>

HAL Id: halshs-05148791

<https://shs.hal.science/halshs-05148791v1>

Submitted on 4 Mar 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC 4.0 - Attribution - Non-commercial use - International License

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Journal of International Economics

journal homepage: www.elsevier.com/locate/jieInequality and optimal monetary policy in the open economy^{☆,☆☆,★}Sushant Acharya^{a,b}, Edouard Challe^{c,b,*}^a University of Melbourne, Australia^b CEPR, United Kingdom^c Paris School of Economics, France

ARTICLE INFO

Dataset link: [Inequality and Optimal Monetary Policy in the Open Economy \(Original data\)](#)

JEL classification:

E21
E30
E52
E63
F31
F41

Keywords:

Open economy
New Keynesian model
Inequality
Optimal policy

ABSTRACT

We study optimal monetary policy in a tractable Small Open Economy Heterogeneous-Agent New Keynesian (SOE-HANK) model in which households face uninsured idiosyncratic risk and unequal bond-market access. We derive conditions under which optimal policy in our SOE-HANK economy entails domestic producer price stability, extending the "open-economy divine coincidence" result of Galí and Monacelli (2005) beyond the Representative-Agent benchmark (SOE-RANK). Away from those conditions, inefficient fluctuations in consumption inequality generate monetary policy tradeoffs. Under plausible calibrations for the trade elasticities, the elasticity of intertemporal substitution, and the cyclical risk of income risk, the central bank stabilizes output and the exchange rate more than in SOE-RANK.

1. Introduction

How do aggregate shocks affect domestic inequality in a small open economy, and what should the central bank do about it? In this paper, we answer this question by constructing a Small-Open Economy Heterogeneous-Agent New Keynesian (SOE-HANK) model with rich cross-sectional heterogeneity. Following the recent SOE-HANK literature (De Ferra et al., 2020; Auclert et al., 2024; Druedahl et al., 2022; Guo et al., 2023), our baseline augments the workhorse model of Galí and Monacelli (2005) with incomplete markets. In particular, our model features households who face (i) uninsurable idiosyncratic income risk à la Aiyagari

[☆] The National Bureau of Economic Research has provided financial sponsorship to make this article open access and had no influence or involvement over the review or approval of any content.

^{☆☆} This article is part of a special issue entitled: ISOM 2024 published in Journal of International Economics.

[★] We thank the referees and the editor for their suggestions and guidance in revising the paper. We are grateful to Keshav Dogra, Tommaso Monacelli and Federica Romei for their comments and to Lukas Nord for excellent research assistance. We also received helpful feedback from seminar participants at Bocconi University, HEC Paris, Rennes SB, PSE, EUI, NBB, VIMM, Bank of Finland, Bank of Italy, FRB Philadelphia and the Universities of Copenhagen, London Queen Mary, Vienna, Tubingen, Goethe Frankfurt, Wisconsin, Cambridge, Venice, Cattolica Milano as well as from conference participants at SED 2022, ASSA 2023, ESSIM 2023, EEA 2023, the 2023 Rice-LEMMA Monetary Conference, STLAR 2023, the CRC TR 2024 workshop on Global Crises, Financial Markets, and Monetary Policy, the 2024 LSE-Oxford Workshop in International Macroeconomics & Finance, and the 2024 Macroeconomic Policies with Heterogenous Agents workshop at Université Paris-Dauphine. Edouard Challe acknowledges financial support from the French National Research Agency (ANR-20-CE26-0018-01) and the hospitality of the Bank of Finland.

* Corresponding author at: Paris School of Economics, France.

E-mail addresses: sushant.acharya@unimelb.edu.au (S. Acharya), edouard.challe@psemail.eu (E. Challe).

<https://doi.org/10.1016/j.jinteco.2025.104076>

Received 30 August 2024; Received in revised form 25 February 2025; Accepted 28 February 2025

Available online 11 March 2025

0022-1996/© 2025 Published by Elsevier B.V. This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).

(1994) and (ii) unequal access to financial markets, leading to hand-to-mouth behavior by some households. Relative to the Small-Open Economy Representative-Agent New Keynesian (SOE-RANK) benchmark, uninsurable risk results in time-varying precautionary savings, while hand-to-mouth behavior and heterogeneous marginal propensities to consume (MPCs) imply that households are differentially exposed to changes in national income. While the SOE-HANK literature has thus far focused on how these sources of household heterogeneity may alter the *positive* predictions of the SOE-RANK benchmark regarding the propagation of aggregate shocks, we take a *normative* perspective: our goal is to understand how uninsured risk and hand-to-mouth behavior affect optimal monetary policy relative to SOE-RANK.¹

Our critical insight is that, in SOE-HANK, aggregate shocks trigger potentially inefficient fluctuations in consumption inequality across households, which a benevolent Ramsey planner would want to stabilize. In our model, these fluctuations arise for two reasons. First, unequal access to bond markets generates fluctuations in *between-group* consumption inequality since the income and consumption of unconstrained and hand-to-mouth households respond differently to aggregate shocks. Second, uninsurable idiosyncratic risk generates fluctuations in *within-group* consumption inequality driven by the time-varying distribution of labor earnings. How much the planner effectively leans against such variations in consumption inequality ultimately depends on how strongly this objective conflicts with the traditional objectives of monetary policy in the absence of household heterogeneity. As is now well understood (and synthesized in Galí, 2015), in the basic closed-economy New Keynesian model, the central bank should pursue domestic price stability and ensure productive efficiency by stabilizing the output gap. In the open economy, an additional motive arises, namely the efficient manipulation of the terms of trade – made possible by the fact that the central bank exerts market power over the supply of domestically-produced goods. The concern for stabilizing consumption inequality in our SOE-HANK economy further complicates matters since stabilizing inequality may require departing from price stability, productive efficiency, and efficient terms-of-trade management – at least temporarily. Our analysis uncovers when and why such departures occur.

To understand the trade-offs introduced by the additional objective of stabilizing inequality as clearly and transparently as possible, we articulate our analysis around a benchmark scenario wherein consumption inequality is present but does *not* alter optimal monetary policy relative to the SOE-RANK benchmark. Quite sensibly, this scenario arises when monetary policy cannot affect fluctuations in inequality. In this situation, the best the central bank can do is focus on its other objectives and disregard inequality altogether, even though inequality is welfare-reducing. This property is similar in spirit to the *open-economy divine-coincidence* result of Galí and Monacelli (2005), who show that under the Cole and Obstfeld (1991) parameterization (when the elasticities of substitution between exported goods, between imported goods, and intertemporal, are all equal to one), optimal policy in SOE-RANK mimics that in the closed-economy benchmark in response to aggregate productivity shocks. Our paper shows that the same open-economy divine coincidence holds in SOE-HANK when, in addition to the Cole and Obstfeld (1991) parameterization, idiosyncratic income risk is *acyclical*.

This result is best understood in relation to the conditions for the optimality of domestic producer price stability under *complete* markets, as derived by Galí and Monacelli (2005). In their model, under the Cole and Obstfeld (1991) parameterization and aggregate productivity shocks, (i) there is no conflict between implementing stable domestic prices and closing the output gap in the corresponding closed economy, (ii) the exchange rate adjusts in the open economy in a way that renders movements in the terms of trade optimal. As it turns out, even with *incomplete* markets with respect to aggregate productivity shocks, under the Cole and Obstfeld (1991) parameterization, balanced trade is optimal, i.e., interest-rate and exchange-rate movements imply that unconstrained households as a whole optimally choose not to save or dis-save – even though it remains feasible. Consequently, the average consumption of unconstrained households is symmetric to that of hand-to-mouth households, eliminating time variations in between-group inequality. Furthermore, under incomplete markets against *idiosyncratic* income shocks, constant income risk implies that within-group consumption inequality is also time-invariant. Thus, under Cole-Obstfeld elasticities and acyclical idiosyncratic income risk, neither between- nor within-group inequality can be manipulated by the central bank, making it optimal to operate “as if” under complete markets.

Any departure from this benchmark scenario generates new tradeoffs for the central bank. To isolate the specific role that consumption risk and inequality play in those tradeoffs, the first departure from the conditions of the SOE-HANK divine coincidence we consider is by making the empirically relevant assumption that idiosyncratic income risk is *countercyclical*, rather than acyclical. This feature leads the central bank to depart from price stability, even though price stability remains feasible. For example, after a contraction in aggregate productivity, the central bank supports aggregate demand more than it would in SOE-RANK – else, consumption inequality would increase too much, bringing excess welfare losses –, which mitigates the recession and associated real appreciation but generates inflation – and the other way around in a productivity-driven expansion. Overall, the optimal policy implements significantly less output and exchange-rate volatility in SOE-HANK than in SOE-RANK after aggregate productivity shocks. This property remains true in our preferred calibration, where we also move away from the Cole and Obstfeld (1991) elasticities towards more realistic elasticities of trade and intertemporal substitution. Importantly, our analysis shows that countercyclical idiosyncratic risk *per se* warrants significant exchange rate stabilization – regardless of the degree of stabilization called for by the concern for between-group inequality due to hand-to-mouth behavior.

Next, we turn to the monetary policy response to capital-flow shocks, triggered by exogenous changes in the world interest rate. In contrast to aggregate productivity shocks, the optimal response to such shocks implies departures from domestic producer price stability even in SOE-RANK. For example, in our SOE-RANK benchmark, a capital inflow triggered by a fall in the world interest

¹ As reviewed below, earlier work has analyzed optimal monetary policy in open economies under hand-to-mouth behavior but ignoring uninsured idiosyncratic risk; our analysis activates both sources of household heterogeneity.

rate would generate a real appreciation and an output contraction (driven by expenditure switching) if the real interest rate at home remained unchanged. The central bank optimally responds to the shock by cutting interest rates at home to mitigate the real appreciation – but does not entirely undo it. This is because cutting rates so much as to match the world interest rate and fully stabilize the exchange rate would be inflationary, and the SOE-RANK central bank trades off a lower inflation gap against a lower output gap. Against this backdrop, we show that two forces tend to make exchange-rate stability relatively more desirable in SOE-HANK than in SOE-RANK. First, even in the absence of idiosyncratic risk, incomplete markets against aggregate shocks imply that the consumption of hand-to-mouth households is heavily exposed to exchange-rate fluctuations. This is because the exchange rate determines the relative price of home-produced goods and, hence, current national income (at a given output level), which hand-to-mouth households cannot smooth away via borrowing and lending. Second, in the presence of cyclical consumption risk, exchange-rate fluctuations and expenditure switching contribute to generating inefficient volatility in output and consumption risk. In our preferred calibration, the central bank implements a smoother exchange-rate response to capital flow shocks in SOE-HANK than in SOE-RANK.

In summary, the broad lesson from our analysis is that when we move away from the (implausible) conditions for SOE-HANK divine coincidence and towards more realistic calibrations for the trade elasticity, the elasticity of intertemporal substitution and the cyclical risk, then the central bank optimally stabilizes output and the exchange rate more in SOE-HANK than in SOE-RANK. Moreover, this conclusion holds regardless of the exact size of average income risk and the share of hand-to-mouth households in the economy (though those parameters may affect the quantitative differences between SOE-HANK and SOE-RANK.)

1.1. Literature review

Our paper bridges three strands of the literature: (i) one that analyzes the propagation of aggregate shocks in open-economy HANK models from a *positive* perspective; (ii) one studying optimal monetary policy in *closed-economy* HANK models; (iii) and the literature studying optimal monetary policy in SOE or two-country RANK models.

De Ferra et al. (2020) was the first paper to introduce within-country heterogeneity in Galí and Monacelli (2005)'s framework to study how capital-flow shocks affect output and national income when a subset of the households have high MPCs. The central theme of that paper is the contractionary Keynesian multiplier effects set off by a currency depreciation in the presence of foreign-currency debt. This was further explored by Auclert et al. (2024), who show how high MPCs magnify the real income channel of currency depreciations, working against expenditure switching and possibly triggering an output contraction. Other contributions following their lead include (Druehdahl et al., 2022), who study the propagation of foreign demand shocks, as well as Guo et al. (2023) and Oskolkov (2023), who study how international integration and exchange-rate regimes shape the propagation of aggregate shocks to domestic inequality. Acharya and Pesenti (2024) study monetary policy spillovers in a two-country HANK setting and show that while the real-income channel cannot change the sign of spillovers relative to RANK, cyclical risk can flip the sign of spillovers. Bayer et al. (2023) study international spillovers in a monetary union HANK model. None of these papers, however, study optimal monetary policy. Related to our line of work stand the contributions of Iyer (2017) and Chen et al. (2023), who study optimal monetary policy in an SOE and Two-country model, respectively, in the presence of hand-to-mouth households (SOE-TANK) and under Cole and Obstfeld (1991) elasticities. We nest some of their results when we shut down idiosyncratic risk, while bringing into the analysis ex-post (in addition to ex-ante) heterogeneity together with departures from Cole and Obstfeld (1991). For example, we also find that the presence of hand-to-mouth households may warrant stronger exchange-rate stabilization (as in Iyer, 2017) but does not by itself break the open-economy divine coincidence under aggregate labor productivity shocks (Chen et al., 2023). Another difference between these two papers and ours is that we do not assume full international risk sharing between unconstrained households of different countries as they do. While this does not matter in the special case of Cole and Obstfeld (1991) elasticities and no idiosyncratic risk, it does matter under any other parameterization.

The present paper naturally extends to the open economy the analysis of Ramsey-optimal monetary policy in *closed-economy* HANK models, as undertaken by Bhandari et al. (2021), Le Grand et al. (2022), McKay and Wolf (2022), Acharya et al. (2023) and Dávila and Schaab (2023). In contrast to the first three papers, which numerically solve for optimal policy, we here follow the approach of Acharya et al. (2023), who rely on the joint assumption of Constant Absolute Risk Aversion (CARA) utility for consumption and conditionally normally distributed income risk to achieve traceability. In the present paper, this analytical approach allows for the formal derivation of the SOE-HANK divine coincidence result, and also greatly facilitates the exploration of optimal policy away from the conditions of divine coincidence.

Last but not least, our paper is a continuation of the traditional literature that examines optimal monetary and exchange-rate policy in various versions of the open-economy RANK model. As discussed above, the starting point of our analysis is the open-economy divine coincidence result of Galí and Monacelli (2005) – which we generalize to incomplete markets (with respect to both aggregate and idiosyncratic shocks) and hand-to-mouth behavior. Related to Galí and Monacelli (2005), Clarida et al. (2001, 2002), Corsetti and Pesenti (2005), Benigno and Benigno (2003, 2006), Faia and Monacelli (2008) and De Paoli (2009a) study optimal monetary policy in SOE or two-country models under complete markets with respect to aggregate shocks (i.e., where the Backus and Smith (1993) condition relating the ratio of marginal consumption utilities to the real exchange rate holds), abstracting from uninsurable idiosyncratic shocks. Extending this line of work, Benigno (2009), De Paoli (2009b) and more recently Corsetti et al. (2023) and Egorov and Mukhin (2023) have examined optimal monetary policy with imperfect risk sharing *internationally* – i.e., disposing of the Backus-Smith condition – but still in the absence of idiosyncratic risk. By and large, the main focus of this literature has been to clarify when and how the central bank's willingness to manipulate the terms of trade may conflict with domestic price stability. Instead, our focus is on the monetary policy tradeoffs that uninsurable idiosyncratic risk and hand-to-mouth behavior may bring, potentially jointly with the tradeoffs induced by terms-of-trade manipulation.

2. Environment

The model is a variant of Galí and Monacelli (2005) that does not admit a representative household. Instead, the economy comprises heterogeneous households who face uninsurable idiosyncratic income risk à la Aiyagari (1994), while some households are excluded from asset markets altogether. Time is discrete. There is no aggregate risk – only unanticipated (persistent) aggregate shocks which occur at date zero, after which all agents in the economy have perfect foresight with respect to aggregate variables. The world economy consists of a continuum of small open economies $j \in [0, 1]$, one of which is "Home" and is accordingly indexed by H . We take the behavior of all other countries as given, and start by describing the behavior of economic agents in the Home country.

2.1. Households

The demographics at Home follow a perpetual-youth structure à la Blanchard-Yaari, in which all households face a survival probability of $\vartheta < 1$ at any date t . To ensure that Home population stays constant (normalized to 1), there are $1 - \vartheta$ "newborn" households, who replace the equal measure of households who did not survive. We introduce this perpetual-youth structure because population turnover ensures that the model always generates a stationary distribution of wealth.

Households differ in two aspects. (i) Since all households face uninsurable idiosyncratic shocks to their labor productivity, they differ in their realization of idiosyncratic productivity. (ii) Households also differ in their access to bond markets: at each date, a share θ of the newborn households cannot access bond markets and must consume their income at each subsequent date as long as they survive, i.e., they are "hand-to-mouth" and have a marginal propensity to consume (MPC) of 1. The other $1 - \theta$ share of newborn households have access to bond markets throughout their lives. We refer to these households as "unconstrained" since they can use asset markets to smooth the effects of shocks to current income on current consumption. In other words, they have an MPC less than 1.

The problem of household i born at time s can then be written as:

$$\max_{\{c_t^s(i), n_t^s(i)\}} \mathbb{E}_s \sum_{t=s}^{\infty} (\beta \vartheta)^{t-s} [u(c_t^s(i)) - v(n_t^s(i))] \tag{1}$$

where $\beta \in [0, 1]$ is the household subjective discount factor, $c_t^s(i)$ their consumption of the composite consumption good (described below) and $n_t^s(i)$ their labor supply.

We set $u(c) = -e^{-\gamma c} / \gamma$, where $\gamma > 0$ is the coefficient of absolute risk aversion (CARA), and $v(n) = \rho e^{(n-\bar{n})/\rho}$, where $\rho > 0$ denotes the Frisch elasticity and \bar{n} is a constant, which scales the dis-utility of labor. As in Acharya et al. (2023), the choice of the CARA utility alongside normally distributed idiosyncratic risk (see below) is key for the tractability of the planning problem, for two reasons. First, the CARA-Normal model delivers consumption functions which are linear in wealth. This implies that the economy permits linear aggregation, allowing us to characterize the aggregate dynamics of the economy without explicitly keeping track of the distribution of wealth. Second, the CARA-Normal framework allows aggregating the intertemporal utilities of infinitely many household types into a simple Social Welfare Function where the welfare cost of inequality is summarized by a one-dimensional sufficient statistic of the time-varying wealth distribution.

2.1.1. Unconstrained households

Unconstrained households can take unrestricted positions in actuarial bonds, which pay one unit of the Home consumption good in the next period conditional on survival. Upon "birth", unconstrained households are transferred an amount of asset wealth equal to the average beginning-of-period wealth of all surviving households (who are lump-sum taxed accordingly). These transfers eliminate the inter-cohort heterogeneity in asset holdings that the Blanchard-Yaari demographic structure entails whenever the economy's net foreign asset position differs from zero. These transfers ensure that, despite the overlapping-generations structure, idiosyncratic risk and access to bond markets are the only sources of household heterogeneity that matter, so that our economy effectively collapses to Galí and Monacelli (2005) under complete markets. The date- t flow budget constraint of an unconstrained household born at date $s \leq t$ is given by:

$$c_t^s(i, u) + (1 + \tau^*) \frac{\vartheta}{R_t} a_{t+1}^s(i) = a_t^s(i) + (1 - \tau^w) w_t n_t^s(i, u) e_t^s(i, u) + D_t + \mathcal{T}_t + \mathbb{T}_t, \tag{2}$$

where ϑ/R_t is the pre-tax date- t price of an actuarial bond, which pays off one unit of the Home consumption good at date $t + 1$. $a_t^s(i)$ is the date t wealth of household i born at date $s \leq t$. Since newborn households are transferred average beginning-of-period assets, we have $a_t^s(i) = a_t$, where

$$a_t = (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int a_t^s(i) di$$

is the mean (beginning-of-period) asset wealth of surviving households (where the \int operator averages over households of the same cohort and the \sum operator averages over cohorts). Next, w_t is the real wage per unit of effective labor and $e_t^s(i, u)$ is individual productivity. D_t denotes firms' dividends and \mathcal{T}_t a lump sum transfer from the government; both are distributed uniformly to all households – hand-to-mouth and unconstrained alike. The assumption that dividends and the lump-sum transfer are uniform is made for simplicity. Asymmetric transfers (due to, say, unequal distribution of dividends) would generate additional sources of

households' unequal exposure to aggregate shocks (see Bhandari et al., 2021 and Acharya et al., 2023, in the context of a closed economy). Finally, τ^* and τ^w are constant tax rates; the former affects the expected returns on savings and thus distorts savings decisions, while the latter affects post-tax labor earnings and thus distorts labor-supply decisions. As explained in Section 4 below, both tax rates are set optimally by the fiscal authority and are necessary for the constrained efficiency of the steady state, given the frictions plaguing the economy. \mathbb{T}_t is a lump-sum transfer received by unconstrained households only and composed of the rebate of the savings taxes ($\tau^* \vartheta a_{t+1}/R_t$) minus the taxes raised to finance the wealth transfers to the newborn $((1 - \vartheta)a_t)$:

$$\mathbb{T}_t = \tau^* \frac{\vartheta}{R_t} a_{t+1} - (1 - \vartheta) a_t \tag{3}$$

Individual labor productivity $e_t^s(i)$ evolves stochastically according to the following process:

$$e_t^s(i, u) = 1 + \sigma_t \xi_t^s(i, u), \quad \text{with} \quad \xi_t^s(i, u) = \xi_{t-1}^s(i, u) + v_t^s(i, u) \quad \text{and} \quad v_t^s(i, u) \sim \mathcal{N}(0, 1), \tag{4}$$

which implies that the conditional variance of next-period individual productivity is equal to σ_t^2 , while the conditional mean is $1 + \sigma_t \xi_t^s(i)$. Our assumption that individual productivity shocks follow a random walk is made for simplicity but also follows from the empirical finding that shocks to individual earnings are highly persistent – see, e.g., Storesletten et al. (2004). On the other hand, allowing for time variations in σ_t is necessary to generate cyclical labor-earnings risk, a robust feature of the data (Storesletten et al., 2004; Nakajima and Smirnyagin, 2019).

Because unconstrained households do not face hard borrowing constraints (only a no-Ponzi game condition), their optimal consumption-saving choice satisfies the individual Euler equation:

$$e^{-\gamma c_t^s(i, u)} = \left(\frac{\beta R_t}{1 + \tau^*} \right) \mathbb{E}_t e^{-\gamma c_{t+1}^s(i, u)} \tag{5}$$

where the survival rate ϑ has dropped out by virtue of the actuarial nature of real bonds.

2.1.2. Hand-to-mouth households

Hand-to-mouth households face idiosyncratic income risk just like unconstrained households. However, they cannot participate in asset markets and consume their current income in every period:

$$c_t^s(i, h) = (1 - \tau^w) w_t n_t^s(i, h) e_t^s(i, h) + D_t + \mathcal{T}_t, \tag{6}$$

where $e_t^s(i, h)$ is the same as in (4).

2.1.3. Labor supply

Under incomplete insurance and a frictionless labor market, households use labor supply for self-insurance against idiosyncratic shocks, working relatively more (less) when financial assets and consumption are low (high). As Auclert et al. (2023) have argued, however, if preferences are separable, this implies implausibly large wealth effects on labor supply. Accordingly, we assume that labor supply is chosen by a continuum of unions acting on behalf of the households. Unions mandate that all households work the same number of hours so that $n_t^s(i) = n_t$ for all (s, i) . This implies that individual hours are independent of individual wealth and, consequently, households cannot use labor supply for self-insurance.

Appendix A shows that the optimal labor supply condition is given by:

$$\mathcal{M}_w e^{(n_t - \bar{n})/\rho} = (1 - \tau^w) w_t e^{-\gamma c_t} \Sigma_t \tag{7}$$

where $\mathcal{M}_w > 1$ is the wage markup, c_t is aggregate consumption (so $e^{-\gamma c_t}$ is the marginal utility of average consumption), and Σ_t is the cross-sectional dispersion of marginal utilities in the population:

$$\Sigma_t = (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int e^{-\gamma(c_t^s(i) - c_t)} di \tag{8}$$

Under complete markets and without hand-to-mouth households, there is no consumption dispersion ($c_t^s(i) = c_t$) and consequently $\Sigma_t = 1$ in Eqs. (7)–(8). With idiosyncratic risk and hand-to-mouth households, Jensen's inequality implies that $\Sigma_t > 1$. Note that we assume that nominal wages are flexible, but will assume that prices are sticky as in Galí and Monacelli (2005).

Given the standard deviation of idiosyncratic productivity σ_t , the conditional standard deviation (as of time $t - 1$) of household time- t income – "time- t income risk" for short – is given by $\sigma_{y,t} = (1 - \tau^w) w_t n_t \sigma_t$. While σ_t is not directly observable, there is substantial evidence about income risk $\sigma_{y,t}$, both in the long run and over the business cycle. Accordingly, we let σ_t adjust in the background in such a way that realized income risk satisfies:

$$\sigma_{y,t} = \sigma_y e^{-\varphi \hat{y}_t}, \tag{9}$$

where \hat{y}_t is the proportional deviation of output from steady state $((y_t - y)/y)$ while σ_y (steady-state income risk) and φ (the cyclicality of income risk) are exogenous parameters. In particular, income risk is *countercyclical* (households face higher income risk during downturns) whenever $\varphi > 0$; $\varphi = 0$ corresponds to the case where income risk is *acyclical*.

2.2. Firms

2.2.1. Home retailers

A perfectly competitive retail sector produces the consumption good consumed by domestic households, using a CES aggregator of Home- and Foreign-produced goods:

$$c_t = \begin{cases} \left[(1-\alpha)^{\frac{1}{\eta}} c_{H,t}^{\frac{\eta-1}{\eta}} + \alpha^{\frac{1}{\eta}} c_{F,t}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}} & \text{for } \eta > 0, \eta \neq 1, \\ \left(\frac{c_{H,t}}{1-\alpha} \right)^{1-\alpha} \left(\frac{c_{F,t}}{\alpha} \right)^{\alpha} & \text{for } \eta = 1, \end{cases} \quad (10)$$

where $c_{H,t}$ is a bundle of Home-produced varieties, $c_{F,t}$ is a bundle of imported, foreign-produced varieties, and $1-\alpha \in [0, 1]$ is the degree of home bias in consumption. The nominal home-currency price of the consumption good c_t is denoted P_t , and the Home-currency prices of the Home- and Foreign-produced goods are denoted $P_{H,t}$ and $P_{F,t}$, respectively. The demand system (10) implies the standard demand curves for Home- and Foreign-produced goods:

$$c_{H,t} = (1-\alpha) p_{H,t}^{-\eta} c_t \quad \text{and} \quad c_{F,t} = \alpha p_{F,t}^{-\eta} c_t, \quad (11)$$

where $p_{H,t} = P_{H,t}/P_t$ and $p_{F,t} = P_{F,t}/P_t$ are the real prices (i.e., in terms of the home consumption good) of the Home- and Foreign-produced goods, respectively, and satisfy:

$$(1-\alpha) p_{H,t}^{1-\eta} + \alpha p_{F,t}^{1-\eta} = 1, \quad (12)$$

with the understanding that $p_{H,t}^{1-\alpha} p_{F,t}^{\alpha} = 1$ in the case where $\eta = 1$.

Under producer currency pricing (PCP), our maintained assumption throughout, we have $P_{F,t} = \mathcal{E}_t P_t^*$, where P_t^* denotes the price of imports expressed in foreign-currency units and \mathcal{E}_t is the nominal exchange rate between Home and Foreign currencies. $Q_t = \mathcal{E}_t P_t^*/P_t$ denotes the real exchange rate. Consequently, the real price of Foreign goods can be written as $p_{F,t} = Q_t$, and the real price of Home goods $p_{H,t}$ depends on the real exchange rate as follows:

$$p_H(Q_t) = \begin{cases} \left(\frac{1-\alpha Q_t^{1-\eta}}{1-\alpha} \right)^{\frac{1}{1-\eta}} & \text{for } \eta > 0, \eta \neq 1, \\ Q_t^{\frac{\alpha}{\alpha-1}} & \text{for } \eta = 1. \end{cases} \quad (13)$$

The Home-produced good $c_{H,t}$ is itself the output of a continuum $j \in [0, 1]$ of Home-produced varieties, which are combined according to the CES aggregator

$$c_{H,t} = \left[\int_0^1 c_{H,t}(j)^{\frac{\varepsilon-1}{\varepsilon}} dj \right]^{\frac{\varepsilon}{\varepsilon-1}}, \quad \varepsilon > 1, \quad (14)$$

which yields the standard demand curves for each variety:

$$c_{H,t}(j) = \left(\frac{P_{H,t}(j)}{P_{H,t}} \right)^{-\varepsilon} c_{H,t}, \quad (15)$$

where $P_{H,t}(j)$ is the nominal price of this variety, and $P_{H,t} = \left[\int_0^1 P_{H,t}^{1-\varepsilon}(j) dj \right]^{\frac{1}{1-\varepsilon}}$. On the other hand, the Foreign good $c_{F,t}$ is an aggregate of goods produced in a continuum of countries $m \in [0, 1]$, with each country-specific basket itself aggregating a continuum of country-specific varieties $j \in [0, 1]$. Formally, we have:

$$c_{F,t} = \begin{cases} \left[\int_0^1 c_{m,t}^{\frac{\nu-1}{\nu}} dm \right]^{\frac{\nu}{\nu-1}} & \text{for } \nu > 0, \nu \neq 1, \\ e^{\int_0^1 \ln c_{m,t} dm} & \text{for } \nu = 1, \end{cases} \quad \text{and} \quad c_{m,t} = \left[\int_0^1 c_{m,t}(j)^{\frac{\varepsilon-1}{\varepsilon}} dj \right]^{\frac{\varepsilon}{\varepsilon-1}}, \quad (16)$$

where $\varepsilon > 1$. Because all foreign prices are symmetric (equal to P_t^*), home retailers' demands for foreign varieties are symmetric across and within countries, i.e. for all (m, j) we have $c_{m,t}(j) = c_{F,t}$, given by Eq. (11).

2.2.2. Foreign retailers

Foreign retailers in every country $m \neq H$ behave symmetrically to those in the home economy, implying that the demand by any country $m \neq H$ for Home-produced variety j is given by:

$$c_{H,t}^*(j) = \alpha \left(\frac{P_{H,t}^*(j)}{P_{H,t}^*} \right)^{-\varepsilon} \left(\frac{P_{H,t}^*}{P_t^*} \right)^{-\nu} c^* \quad (17)$$

where $P_{H,t}^*(j)$ is the foreign-currency price of variety j from country H , $P_{H,t}^* = \left[\int_0^1 P_{H,t}^{*1-\varepsilon}(j) dj \right]^{\frac{1}{1-\varepsilon}}$ the foreign-currency price index for all goods from country H , and c^* is world consumption. Under PCP we have $P_{H,t}^*(j) = P_{H,t}(j)/\mathcal{E}_t$ for all j , while $Q_t = \mathcal{E}_t P_t^*/P_t$.

Hence, denoting Home's total volume of exports by $c_{H,t}^*$, Eq. (17) can be broken down as follows:

$$c_{H,t}^*(j) = \left(\frac{P_{H,t}(j)}{P_{H,t}} \right)^{-\varepsilon} c_{H,t}^*, \quad \text{where} \quad c_{H,t}^* = \alpha \left(\frac{p_H(Q_t)}{Q_t} \right)^{-\nu} c^*, \tag{18}$$

where $p_H(Q_t)$ is defined in (13).

2.2.3. Home producers

Home varieties are produced by monopolists using a linear production function turning one unit of labor into z_t units of the specialized good. Prices are sticky à la Rotemberg (1982): if firm j 's date $t-1$ price was $P_{H,t-1}(j)$, it costs the firm $\frac{\Psi}{2} \left(\ln \frac{P_{H,t}(j)}{P_{H,t-1}(j)} \right)^2 y_t$ to change the price at date t to $P_{H,t}(j)$. Ψ is a constant that scales the cost, and as is standard, the cost is assumed to be proportional to Home output, y_t . The output of firm j net of this cost can then be written as:

$$y_t(j) = z_t n_t(j) - \frac{\Psi}{2} \left(\ln \frac{P_{H,t}(j)}{P_{H,t-1}(j)} \right)^2 y_t \tag{19}$$

Firm j 's date- t real value of dividends (in terms of the home consumption basket) is given by $D_t(j) = p_H(Q_t) \left(\frac{P_{H,t}(j)}{P_{H,t}} \right) y_t(j) - (1 - \tau) w_t n_t(j)$, where τ denotes a payroll subsidy. The firm sets its price $P_{H,t}(j)$ to maximize the net present-discounted value of dividends, taking as given the demand curve

$$y_t(j) = \left(\frac{P_{H,t}(j)}{P_{H,t}} \right)^{-\varepsilon} (c_{H,t} + c_{H,t}^*) \tag{20}$$

where $c_{H,t}$ is given by (11) and $c_{H,t}^*$ by (18). Imposing symmetry on the firm's optimal pricing decision yields the standard New Keynesian Phillips curve:

$$\ln \Pi_{H,t} = \kappa \left[1 - \left(\frac{\chi - 1}{\chi - 1 + \alpha} \right) \frac{p_H(Q_t) z_t}{w_t} \right] + \beta \left(\frac{z_t w_{t+1} y_{t+1}}{z_{t+1} w_t y_t} \right) \ln \Pi_{H,t+1}, \tag{21}$$

where $\kappa = \varepsilon/\Psi$ is the slope of the NKPC and $\chi = \eta(1 - \alpha) + \nu$ is (minus) the sum of the price elasticity of imports ($-\eta(1 - \alpha)$) and the price elasticity of exports ($-\nu$) – the "trade elasticity" for short.

Equivalently, χ is the elasticity of the demand for Home goods with respect to their real price $p_{H,t}$. Importantly, in (21), the payroll subsidy τ has been set to

$$1 - \tau = \left(1 - \frac{1}{\varepsilon} \right) \left(\frac{\chi - 1 + \alpha}{\chi - 1} \right), \tag{22}$$

which achieves the optimal balance, in steady state, between reducing the monopolistic distortion induced by producers' market power and the domestic planner's desire to exploit its market power over the international price of Home goods (see Farhi and Werning (2012) for a derivation of the formula). We follow Farhi and Werning (2012) in assuming that every country sets its payroll subsidy optimally, taking as given the payroll subsidies set by the other countries. The symmetric equilibrium yields (22), as well as a steady-state real exchange rate of $Q = 1$. Intuitively, in the closed economy ($\alpha = 0$), Eq. (22) collapses to $\tau = \varepsilon^{-1}$ and the subsidy eliminates the incidence of producers' market power, rendering the steady state optimal from the planner's perspective. In the open economy ($\alpha \in (0, 1)$), the fiscal authority finds it optimal to lower steady-state Home output relative to the closed economy in order to raise the relative price of Home goods and, ultimately, Home national income. Therefore, the subsidy that renders the steady state optimal from the planner's perspective is lower than in a closed economy.

Finally, since producers behave symmetrically (so that $P_{H,t}(j) = P_{H,t}$ for all j), Eq. (19) implies that Home output is given by

$$y_t = \frac{z_t n_t}{1 + \frac{\Psi}{2} (\ln \Pi_{H,t})^2} \tag{23}$$

The demand for home output is determined according to Eqs. (11), (18) and (20). Imposing symmetry, the market-clearing condition for Home goods simplifies to

$$y_t = (1 - \alpha) \left(p_H(Q_t) \right)^{-\eta} c_t + \alpha \left(\frac{Q_t}{p_H(Q_t)} \right)^{\nu} c^* \tag{24}$$

2.2.4. Financial intermediaries

Because of imperfect international risk sharing, the Home's net foreign asset position (NFA) responds to aggregate shocks. We assume that capital flows are intermediated by a competitive financial sector. These financial intermediaries hold positions in international bonds (paying off units of the composite consumption goods of the various countries) while trading actuarial bonds (paying off units of the domestic consumption basket in case of survival) with home households. In the spirit of Schmitt-Grohé and Uribe (2003), the technology for turning domestic savings into international lending (or international borrowing into domestic consumption) involves transaction costs that create an (arbitrarily small) wedge between the home country's savings and its NFA, ensuring that the Home's NFA eventually returns to steady state after the aggregate shock dissipates.

To be more specific, at time t intermediaries issue $(1 - \theta) a_{t+1}$ real actuarial bonds (i.e., a_{t+1} per unconstrained household) at a price ϑ/R_t , which will pay one unit of the domestic composite consumption good to the ϑ surviving households at time $t + 1$. Total

savings $(\vartheta/R_t)(1-\theta)a_{t+1}$ are turned into $\Phi((1-\theta)a_{t+1})(\vartheta/R_t)(1-\theta)a_{t+1}$ good units available for international lending, where $\Phi(\cdot)$ is a C^2 function satisfying $\Phi(0) = 1, \Phi' < 0, \Phi'' > 0$.² The date- t price of a Foreign bond (i.e., paying one unit of the Foreign consumption basket at date $t+1$) is $1/R_t^*$, where R_t^* is the world interest rate, which is taken as given by the Home country. Letting $A_{F,t+1}$ denote the net quantity of Foreign bonds purchased by home intermediaries, the relation between the home country's savings and its NFA is given by³:

$$\underbrace{(1/R_t^*)Q_t A_{F,t+1}}_{\text{NFA}} = \underbrace{\Phi[(1-\theta)a_{t+1}]}_{\text{transaction costs}} \times \underbrace{(\vartheta/R_t)(1-\theta)a_{t+1}}_{\text{Home savings}} \tag{25}$$

At time $t+1$, Home intermediaries collect from their bond portfolio the amount $Q_{t+1}A_{F,t+1}$ and pay out $(1-\theta)\vartheta a_{t+1}$ to domestic survivors. Since intermediaries make no profits, we have:

$$Q_{t+1}A_{F,t+1} = (1-\theta)\vartheta a_{t+1} \tag{26}$$

Combining (25) and (26) gives the following real interest rate parity equation:

$$R_t = \frac{Q_{t+1}}{Q_t} R_t^* \Phi[(1-\theta)a_{t+1}] \tag{27}$$

In what follows, we assume that $\Phi(x) = e^{-\Gamma x/(1-\theta)}$, where $\Gamma \geq 0$ is an arbitrary constant. Finally, aggregating household savings within and across cohorts, we show in Appendix B the economy-wide asset accumulation equation to take the form:

$$(1-\theta)\vartheta a_{t+1} = R_t [(1-\theta)\vartheta a_t + p_H(Q_t)y_t - c_t], \tag{28}$$

where $(1-\theta)\vartheta a_{t+1}$ are the liabilities of financial intermediaries (see Eq. (26).)

2.3. Monetary and fiscal policy

We assume that Home monetary policy directly controls the path of real interest rate R_t , which is chosen optimally. Similarly, fiscal policy sets the constant labor-income tax τ^w , the savings tax τ^* , as well as the wealth tax $\{\tau_t^a\}_{t=0}^\infty$. These fiscal instruments are set optimally, as we discuss in Section 4 below. Unconstrained households enjoy the lump-sum rebate (3), while the lump sum transfer \mathcal{T}_t received by all households corresponds to the rebate of labor income taxes ($\tau^w w_t n_t$) minus the subsidies to firms ($\tau w_t n_t$) so that

$$\mathcal{T}_t = (\tau^w - \tau) w_t n_t. \tag{29}$$

3. Equilibrium

3.1. Unconstrained households

As in Acharya and Dogra (2020) and Acharya et al. (2023), the CARA-Normal structure of our model allows us to linearly aggregate the individual consumption decisions of unconstrained households into a single aggregate Euler equation for this group. As we show in Appendix C, this Euler equation takes the following form:

$$\Delta c_{t+1}(u) = \frac{1}{\gamma} \ln \left(\frac{\beta R_t}{1 + \tau^*} \right) + \frac{\gamma}{2} \sigma_{c_u,t+1}^2, \tag{30}$$

where $c_t(u)$ is the mean consumption of unconstrained households and $\sigma_{c_u,t+1}$ the conditional standard deviation (as of time t) of next-period consumption – "consumption risk".

The first term on the RHS of (30) reflects the fact that a higher post-tax real interest rate incentivizes households to save more, which causes consumption growth going forward to increase. The second term on the RHS reflects unconstrained households' precautionary motive in the face of consumption risk. Higher consumption risk $\sigma_{c_u,t+1}^2$ causes households to increase their desired precautionary savings, also causing consumption growth to increase. As we also show in Appendix C, under our assumed risk process, consumption risk evolves according to the following joint forward recursion:

$$\sigma_{c_u,t} = \mu_t \sigma_{y,t} + (1 - \mu_t) \sigma_{c_u,t+1} \tag{31}$$

² If the country is a creditor (i.e., $(a_{t+1} > 0)$, then $1 - \Phi((1-\theta)a_{t+1}) > 0$ per unit of savings is sunk into intermediation services, and hence only the amount $(\vartheta/R_t)(1-\theta)a_{t+1} [1 - (1 - \Phi((1-\theta)a_{t+1}))] = \Phi((1-\theta)a_{t+1})(\vartheta/R_t)(1-\theta)a_{t+1} < (\vartheta/R_t)(1-\theta)a_{t+1}$ is effectively turned into international lending. Conversely, if the country is a debtor ($a_{t+1} < 0$), then every unit of the domestic good made available for domestic consumption over domestic output sinks an additional quantity $\Phi((1-\theta)a_{t+1}) - 1 > 0$ of goods, i.e., it takes $(\vartheta/R_t)(1-\theta)a_{t+1} [1 + (\Phi((1-\theta)a_{t+1}) - 1)] = \Phi((1-\theta)a_{t+1})(\vartheta/R_t)(1-\theta)a_{t+1} > (\vartheta/R_t)(1-\theta)a_{t+1}$ of international lending to make $(\vartheta/R_t)(1-\theta)a_{t+1}$ available to domestic consumers.

³ Since foreign countries are symmetric, the bond they issue are perfect substitutes, which warrants a single foreign bond price $1/R_t^*$. Moreover, as there is no home bias in asset holdings and the home economy is vanishingly small relative to the rest of the world, bonds denominated in the home consumption basket have negligible space in home intermediaries' portfolios. That is why they do not appear in Eq. (25).

where

$$\mu_t^{-1} = 1 + \frac{\vartheta(1 + \tau^*)}{R_t} \mu_{t+1}^{-1} \quad (32)$$

Consumption risk at date t depends not just on current income risk $\sigma_{y,t}$ but also on how much of it passes through to consumption risk, μ_t . This is captured by the first term on the right-hand side of (31). Furthermore, since individual productivity is persistent, it is not just today's income risk and passthrough that matter for consumption risk; instead, it is the entire path of future expected income risk and passthrough. This is captured by the last term on the right-hand side of (31).

Taken together, Eqs. (30) to (32) encompass the key role of time-varying precautionary savings in determining aggregate demand in SOE-HANK. For example, an increase in current or future income risk tends to raise current consumption risk, urging unconstrained households to precautionary save, which contracts aggregate demand. If income risk is countercyclical, this precautionary savings channel magnifies the effect of aggregate shocks affecting aggregate demand.

3.2. Hand-to-mouth households

The individual consumption of hand-to-mouth households is given by (6). Since aggregate dividends are $D_t = p_{H,t}y_t - (1 - \tau)w_t n_t$, the aggregate lump sum transfer is $\mathcal{T}_t = (\tau^w - \tau)w_t n_t$, and given the individual productivity process (4), we can express the consumption of a HtM household i , born at date s , as:

$$c_t^s(i, h) = p_H(Q_t)y_t + \sigma_{y,t} \varepsilon_t^s(i, h) \quad (33)$$

Aggregating over all HtM households, their average consumption is simply

$$c_t(h) = p_H(Q_t)y_t, \quad (34)$$

Eqs. (13) and (33)–(34) show how exposed hand-to-mouth households' consumption is to changes in real income triggered by exchange rate movements – the "real income channel" emphasized by Auclert et al. (2024). At any given output level y_t , a domestic currency depreciation (say, triggered by capital outflows) leading to a fall in the relative price of home-produced goods hits the income and consumption of hand-to-mouth households as a whole one for one. This is in contrast to unconstrained households, which can reduce the effect of exchange-rate movements on their current consumption via borrowing and lending. To be more specific, holding everything else constant, by Eq. (34), a real exchange rate variation $dQ_t > 0$ triggers a change in $c_t(h)$ of size $dc_t(h) = p'_H(Q_t)y_t dQ_t < 0$ – since HtM households' MPC is equal to 1. On the other hand, unconstrained households as a whole behave like a permanent-income consumer with Euler equation (30) and budget constraint (45) in Appendix B. Their MPC in response to a transitory income change is thus less than 1, implying that $dc_t(h) < dc_t(u) < 0$ after a one-off shock $dQ_t > 0$.

3.3. Aggregate Euler equation

Ultimately, aggregate consumption c_t is the weighted average of the two groups, i.e.,

$$c_t = (1 - \theta)c_t(u) + \theta c_t(h) \quad (35)$$

Combining (30), (34) and (35), we get the following aggregate Euler equation determining aggregate consumption growth in our economy:

$$\Delta c_{t+1} = (1 - \theta) \underbrace{\left[\ln \left(\frac{\beta R_t}{1 + \tau^*} \right) + \frac{\gamma}{2} \sigma_{c_u, t+1}^2 \right]}_{\text{unconstrained households}} + \theta \underbrace{\left(p_H(Q_{t+1})y_{t+1} - p_H(Q_t)y_t \right)}_{\text{hand-to-mouth households}} \quad (36)$$

3.4. Welfare cost of consumption dispersion

While there are various measures of income and consumption inequality that one could compute, our framework identifies a unique welfare relevant measure of inequality, the exact form of which depends on the objective function we endow the planner with. As in Acharya et al. (2023), we assume that the planner is utilitarian and puts equal weight on the expected discounted lifetime utility of all households alive at date 0, and puts a Pareto weight of β^t on the lifetime value of households born at date $t > 0$. In this case, Appendix D shows that the planner's objective function can be written as $\mathbb{W}_0 = \sum_{t=0}^{\infty} \beta^t \mathbb{U}_t$, where the date- t total utility flow to the planner \mathbb{U}_t is given by:

$$\mathbb{U}_t = u(c_t)\Sigma_t - v(n_t), \quad (37)$$

and $\Sigma_t \geq 1$, given by Eq. (8) above, measures the social cost associated with the dispersion in the marginal utility of consumption across the entire population. In SOE-RANK, we have $\Sigma_t = 1$ and thus $\mathbb{U}_t = u(c_t) - v(n_t)$: the flow utility is simply that brought to the planner by the Representative Agent. In contrast, in SOE-HANK, a mean-preserving increase in consumption dispersion is costly to the planner, which is manifested in the fact that $\Sigma_t > 1$. Recall that because of our assumption of CARA utility, $u(c) < 0$, and so an increase in Σ_t lowers \mathbb{U}_t .

To the extent that Σ_t captures consumption dispersion over the population as a whole, it encompasses both ex-ante heterogeneity (in access to bond markets) and ex-post heterogeneity (due to idiosyncratic shocks). We further show in Appendix D that Σ_t can be broken down as follows:

$$\Sigma_t = (1 - \theta) e^{-\gamma\theta Y_t} \Sigma_{u,t} + \theta e^{\gamma(1-\theta)Y_t} \Sigma_{h,t} \quad (38)$$

where $Y_t = c_t(u) - c_t(h)$ captures consumption dispersion *between* groups and the $\Sigma_{k,t}$ s capture (the welfare cost of) consumption dispersion *within* each group:

$$\Sigma_{k,t} = (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int e^{-\gamma(c_t^s(i,k) - c_t(k))} di, \quad k = u, h. \quad (39)$$

Eqs. (38)–(39) encode how between- and within-group inequality jointly contribute to the overall welfare cost of inequality. Notice that when there are between-group inequality in favor, say, of the unconstrained (i.e., $Y_t > 0$), then the planner gives more weight on inequality within the group of hand-to-mouth households ($e^{(1-\theta)Y_t}$ goes up) and less on inequality within the group of unconstrained households ($e^{-\theta Y_t}$ goes down). The opposite occurs when $Y_t < 0$. Of course, $\Sigma_{k,t} = 1$ whenever households are symmetric within their group, i.e., when idiosyncratic risk is shut down. In this case, the SOE-HANK model is simply a Two-Agent (SOE-TANK) model in which only between-group inequality matter.

Within-group inequality are driven by idiosyncratic risk, to which unconstrained and hand-to-mouth households respond very differently. On the one hand, unconstrained households are permanent-income consumers who can freely borrow or lend and whose individual consumption (as a deviation from the mean consumption of this group) follows a random walk (see Appendix D.1). As a result, consumption dispersion within this group following the repeated occurrence of idiosyncratic shocks is persistent. This in turn implies that, for this group, (the welfare cost of) consumption dispersion $\Sigma_{u,t}$ is a slow-moving variable obeying a quasi-AR(1) process driven by consumption risk (see Appendix D.2 for the derivation):

$$\ln \Sigma_{u,t} = \frac{\gamma^2 \sigma_{c_{u,t}}^2}{2} + \ln(1 - \vartheta + \vartheta \Sigma_{u,t-1}) \quad (40)$$

For $\Sigma_{u,t}$ to be stationary, it must be that the dispersion of marginal utility among the unconstrained households have a mean reverting tendency. As in Acharya et al. (2023), this requires the turnover rate $1 - \vartheta$ to be sufficiently large: formally, we require that $-2 \ln \vartheta > \gamma^2 \sigma_y^2$.

On the other hand, hand-to-mouth households cannot borrow or lend, which implies that their consumption risk is equal to their income risk at every point in time, i.e., $\sigma_{c_{h,t}} = \sigma_{y,t}$. Aggregating consumption dispersion across cohorts for this group, we find $\Sigma_{h,t}$ to be (see Appendix D.2 for details):

$$\Sigma_{h,t} = \frac{1 - \vartheta}{e^{-\frac{\gamma^2 \sigma_{y,t}^2}{2}} - \vartheta} \quad (41)$$

4. Optimal policy

The planner maximizes $\mathbb{W}_0 = \sum_{t=0}^{\infty} \beta^t U_t$, where U_t is defined by (37)–(38), subject to the optimal labor-supply condition (7), the New Keynesian Phillips curve (21), the output Eq. (23), the market-clearing condition (24), the real interest rate parity (27), the NFA dynamics (28) the aggregate Euler equation (36) (together with the joint recursion (31)–(32)) and the within-group cost-of-inequality indices (40) and (41).

The Lagrangian function corresponding to this welfare maximization problem is formulated in Appendix E. We solve the optimization problem in two steps. First, we compute the optimal values of the (time-invariant) fiscal instruments τ^* and τ^w ; this ensures the constrained efficiency of the steady state of the Ramsey plan. Next, we characterize the optimal monetary policy response to aggregate shocks, given optimal fiscal policy. In so doing, we set the lagged values (i.e., dated $t = -1$) of the Lagrange multipliers on the forward looking constraints to their steady-state counterparts, as is common in the New Keynesian literature.

4.1. Optimal fiscal policy

Since consumption inequality lowers average welfare, the HANK planner unambiguously prefers *lower* consumption inequality, and would like to engineer some redistribution. If monetary policy was the only policy instrument the planner had, they would use monetary policy to induce some redistribution and reduce the average level of inequality. However, monetary policy is arguably not the perfect tool to engineer such a redistribution, and fiscal policy is likely a better way to achieve it. In contrast, it is commonly believed that monetary is relatively more agile compared to fiscal policy, and can respond quickly to cyclical fluctuations.

To strike a balance between these concerns, we take the following approach. As in Acharya et al. (2023), we limit the burden on monetary policy by assuming that the planner sets fiscal policy optimally absent aggregate shocks. The fiscal instruments are set so that in steady state, the benefit of lowering inequality is exactly balanced by the cost of productive inefficiency, implying that fiscal policy achieves the *constrained*-efficient level of steady-state inequality. As a result, absent any aggregate shock, it is optimal for monetary policy not to try to reduce inequality further, for example, by engineering a surprise interest-rate cut. Consequently, in our specification of the optimal policy problem, fiscal policy determines the *average* level of consumption inequality, leaving monetary policy to address the cyclical fluctuations of inequality around this average, constrained-efficient level. Appendix E.2 provides formal details about how fiscal policy optimally sets the savings tax τ^* , labor-income τ^w , and the date-0 wealth tax τ_0^a ; we just discuss the forces that shape these choices here.

Savings tax. Since the rest of the world is assumed to be risk-neutral while sharing the same subjective discount factor as Home households, the steady state world interest rate is $R^* = \beta^{-1}$. At that subjective discount factor and in the absence of a savings tax, domestic households, who precautionary save against uninsured idiosyncratic risk, accumulate claims on the rest of the world and accordingly incur the deadweight loss generated by the transaction cost (see Section 2.2.4). Formally, Eqs. (27), (30) and (31) imply that $e^{-\Gamma a} = (1 + \tau^*)e^{-\frac{\gamma^2 \sigma_y^2}{2}}$ in any steady state, so that $a > 0$, and hence $R < R^*$, at $\tau^* = 0$. The optimal savings tax $\tau^* = e^{\frac{\gamma^2 \sigma_y^2}{2}} - 1$ deters those inefficient capital outflows and eliminates the corresponding deadweight loss. Without the savings tax, the planner would attempt to address the inefficiency of steady state capital outflows by stimulating aggregate demand, but this would inefficiently distort production and create inflation.

Labor-income tax. The planner sets a nonzero labor income tax τ^w , for two reasons. First, the income tax is used to correct the labor-market distortions induced by the market power of unions (see Section 2.1.3); its value in the absence of risk is simply $\tau^w = 1 - \mathcal{M}_w < 0$, i.e., a subsidy prevents labor supply from falling short of its efficient level (Erceg et al., 2000). Second, to the extent that the labor-income tax affects labor supply, it affects the steady state level of output and hence income risk (by (9)). The optimal labor income tax (or subsidy) is that which optimally balances those two forces in steady state. Note that if income risk is acyclical ($\varphi = 0$ in (9)), then the second motive is moot, and the optimal income tax is the same as in the absence of risk. However, if income risk is countercyclical ($\varphi > 0$), $\tau^w = 1 - \mathcal{M}_w$ is no longer optimal in our SOE-HANK economy. Instead, at the margin, the planner would like to increase labor supply and hence steady state output to reduce the level of income risk faced by households. In the absence of the labor-income tax, or if the latter were constrained only to correct the labor-market distortion, the planner would attempt to increase steady-state output in order to lower risk. It would do so by running the economy hot, which would result in positive steady-state inflation. Just as the savings tax allowed the planner to avoid capital outflows without overheating the economy, the labor income tax allows the planner to achieve its desired level of steady-state output without steady-state inflation.

Wealth tax. Rather than focusing on the redistributive role of monetary policy, we focus on how optimal monetary policy in SOE-HANK differs from that in SOE-RANK due to the planner's desire to compensate for missing insurance markets. We do so by endowing the planner with a tax on total wealth at date 0, which they can use to achieve their desired level of redistribution at date 0. As in Acharya et al. (2023), this tax is levied on (demeaned) financial wealth, as well as the predictable component of human wealth as of the end of date -1 . As discussed in Acharya et al. (2023), this tax allows the planner to achieve their desired level of redistribution without using monetary policy, thus leaving monetary policy unencumbered. Without this wealth tax, even when no aggregate shock hits the economy, the planner would implement a surprise change in monetary policy to achieve the desired level of redistribution at date 0. Since we study the problem of a utilitarian planner, who puts equal Pareto weights on all households alive at date 0, the planner optimally uses this wealth tax to eliminate all pre-existing wealth inequality at date 0 by setting $\tau_0^a = 1$.⁴ Given the optimal time-0 wealth tax, it is without loss of generality to set $\tau_t^a = 0$ for all $t > 0$, since only the post-tax real interest rate $R_t(1 - \tau_t^a)/(1 + \tau^*)$ matters for allocations as all agents have perfect foresight about aggregate variables as of date 0.

4.2. Optimal monetary policy

Unlike optimal fiscal policy, which is set optimally absent aggregate shocks, the planner's choice of optimal monetary policy can be described as a sequence of real interest rates $\{R_t\}_{t=0}^{\infty}$ conditional on the aggregate shock that realizes at date 0. In what follows, we study the optimal monetary policy responses to a (i) temporary aggregate productivity shock $\{z_t\}$ and (ii) a temporary shock to the world interest rate $\{R_t^*\}$. Studying the optimal response to aggregate productivity shocks allows us to compare our findings to much of the SOE literature, most notably Galí and Monacelli (2005), and to highlight how introducing uninsurable income risk and hand-to-mouth behavior affects optimal policy. Shocks to the Foreign interest rate, on the other hand, have been a critical motivation for the recent development of *positive* SOE-HANK models (De Ferra et al., 2020; Auclert et al., 2024), and we offer a normative perspective on how a small open economy should optimally respond to such shocks.

Calibration. The impulse-response functions below plot the elasticities of the endogenous variables with respect to one-off, persistent changes in aggregate labor productivity or the world interest rate. The persistence parameter for both shocks is set to 0.9. Next, we set the steady state Foreign interest rate to $R^* = 1/\beta = 1.04$, and we set $\theta = 0.85$, following, e.g., Nisticò (2016), Farhi and Werning (2019) and Acharya et al. (2023). We set the Frisch elasticity of the median Home household to $\rho = 1/3$. Regarding the supply side, we set the elasticity of substitution between varieties to $\varepsilon = 10$ and the slope of the NKPC to $\kappa = 0.1$; the inflation cost parameter $\Psi = \varepsilon/\kappa$ follows. Steady-state earnings risk is set to $\sigma_y = 0.1$. Further, following standard calibrations in the SOE literature, we set $\alpha = 0.4$ and $\Gamma = 0.1$. Our baseline value for the share of hand-to-mouth households is $\theta = 0.3$, also in line with the literature. All the above parameters are held constant across the figures. We then consider several configurations for the elasticities of substitutions (intertemporal $1/\gamma$; and across consumption bundles, η and ν) as well as the cyclicity of income risk φ , as we describe below.

⁴ In Acharya et al. (2023), we also considered planners whose desired level of redistribution was less than 100%, and showed how that affected the optimal conduct of monetary policy. In the context of that paper, we showed that while the monetary policy response is qualitatively similar (in fact, the desire to stabilize output is further strengthened), quantitatively, the difference between the optimal policy chosen by the utilitarian planner and non-utilitarian planners is small. Thus, we focus on the problem of the utilitarian planner in this paper.

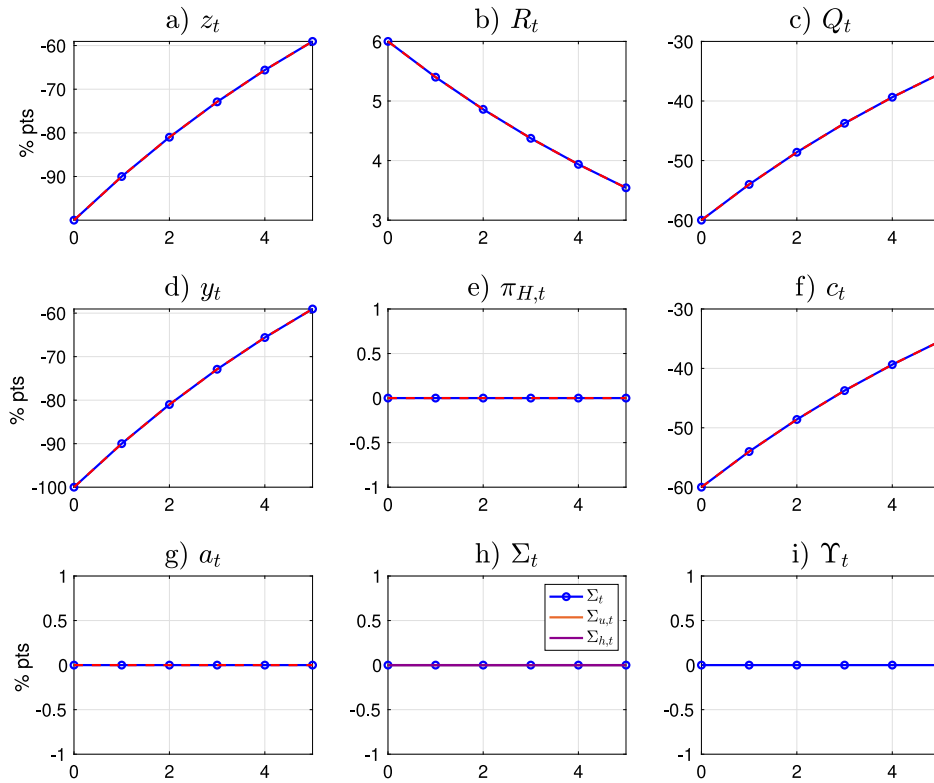


Fig. 1. SOE-HANK (solid blue with circles) vs. SOE-RANK (dashed red) under Cole-Obstfeld elasticities ($\gamma = \eta = \nu = 1$) and acyclical risk ($\varphi = 0$). Price stability is optimal in both SOE-HANK and SOE-RANK in response to aggregate productivity shocks. The exact values of θ (the share of hand-to-mouth households) and σ_y (the size of earnings risk) do not matter for this result.

4.2.1. Responses to domestic productivity shocks

We first characterize the optimal monetary policy response to aggregate productivity shocks. Before showing how uninsurable risk and hand-to-mouth behavior affect the optimal conduct of monetary policy, it is useful to highlight that merely the presence of uninsurable idiosyncratic risk or hand-to-mouth households does not necessitate a change in optimal monetary policy relative to the SOE-RANK benchmark. Proposition 1 below formalizes this.

Proposition 1. Under the Cole and Obstfeld (1991) parameterization ($\gamma = \eta = \nu = 1$), suppose households face income risk ($\sigma_y^2 > 0$), but income risk is acyclical ($\varphi = 0$) and there is a fraction $\theta \in [0, 1)$ of hand-to-mouth households. Then, despite the presence of idiosyncratic income risk and hand-to-mouth households, the optimal monetary policy in SOE-HANK in response to an aggregate productivity shock is the same as in SOE-RANK: optimal policy perfectly stabilizes domestic producer price inflation at each date. Moreover, this policy is optimal regardless of the share of hand-to-mouth households θ or the size of earnings risk σ_y^2 .

Proof. See Appendix F. \square

Proposition 1 is similar in spirit to the main result obtained by Galí and Monacelli (2005) under complete markets and is best explained in relation to their findings. First, recall that in the closed-economy RANK model, producer price stability is optimal following an aggregate productivity shock – a result often referred to as the “divine coincidence” since Blanchard and Galí (2007). In short, the planner can implement both price stability and productive efficiency in response to aggregate productivity shocks, because the latter do not generate a monetary policy tradeoff. In the open economy, there is potentially another margin that monetary policy must internalize, namely how changes in monetary policy affect the nominal exchange rate and ultimately the terms of trade. Galí and Monacelli (2005) showed, however, that under the Cole and Obstfeld (1991) parameterization, there is an “open-economy divine coincidence”: the very same policy of domestic producer price stability as in the closed-economy remains optimal, as it achieves the planner’s desired movements in the terms of trade (in addition to ensuring productive efficiency.)

The open-economy divine coincidence SOE-RANK is depicted in Fig. 1 (dashed red line). After a contraction in aggregate productivity, the optimal policy achieves zero inflation – thereby replicating the flexible price (or natural) allocation. In this allocation, the supply of the Home-produced goods decreases and its relative price rises, causing an appreciation of the domestic consumption basket (i.e., a fall in the real exchange rate Q_t). Aggregate domestic consumption falls, and it perfectly tracks aggregate

income so that the Home country's savings remain at zero all along the transition path. The optimal policy is implemented via a hike in the policy rate, which also replicates the path of the natural interest rate. According to Proposition 1 above, a similar open-economy divine coincidence holds in SOE-HANK as long as income risk is acyclical ($\varphi = 0$), even though the planner now also must worry about how its actions affect consumption inequality. To understand why this is the case, let us start with the case in which there are no hand-to-mouth households ($\theta = 0$), i.e., all households are unconstrained. In this simple scenario, even though there is *within*-group inequality (due to uninsured idiosyncratic risk), there is (trivially) no *between*-group inequality. To see how monetary policy can affect (the welfare cost of) inequality, recall from Eq. (40) that by affecting the level of consumption risk ($\sigma_{c_{u,t}}^2$) that unconstrained households face, monetary policy can affect the evolution of $\Sigma_{u,t}$ (which in this case is equal to Σ_t , since we have set $\theta = 0$). Log-linearizing (9) and (31), we observe that, to first order, consumption risk evolves as follows:

$$\hat{\sigma}_{c_{u,t}} = -\varphi \hat{y}_t + \frac{\vartheta(1 - \tau^*)}{R} \hat{\sigma}_{c_{u,t+1}}$$

When $\varphi > 0$, monetary policy can reduce consumption risk via the income risk channel: higher aggregate output reduces the level of income and consumption risk that unconstrained households face. However, when $\varphi = 0$, the above expression implies that $\hat{\sigma}_{c_{u,t}} = 0$ for all t , and so (40) implies that monetary policy cannot affect the evolution of the welfare cost of inequality ($\Sigma_{u,t} = \Sigma_t$ stays constant). Therefore, while the monetary authority does dislike consumption inequality, it cannot affect it, and so the best it can do is focus on the objectives of price stability and productive efficiency. Ultimately, divine coincidence continues to hold in SOE-HANK when risk is acyclical are there are no hand-to-mouth households.

Perhaps surprisingly, according to Proposition 1, adding a fraction of hand-to-mouth households ($\theta > 0$), does not change optimal policy under Cole and Obstfeld (1991) elasticities. Even under acyclical income risk (so that within-group inequality is unaffected by monetary policy), we could expect a policy of domestic price stability to interfere with the planner's objective of stabilizing between-group inequality $Y_t = c_t(u) - c_t(h)$. However, under the Cole and Obstfeld (1991) elasticities and acyclical income risk, the concern for between-group inequality does not generate a monetary policy tradeoff either. To understand why this is the case, recall that in the complete-market benchmark of Galí and Monacelli (2005), efficient risk-sharing coupled with Cole and Obstfeld (1991) elasticities imply that Home consumption equals the real exchange rate – a direct implication of the Backus and Smith (1993) efficient risk-sharing condition. The exact same property also holds under incomplete markets in response to aggregate productivity shocks, and it implies that, under the optimal plan, the Home economy does not borrow from or lend to the rest of the world, i.e., a SOE which starts with a zero net foreign asset position maintains that position forever under the policy of domestic producer price stability. In other words, the unconstrained households in our SOE-HANK with acyclical risk do not accumulate or decumulate any foreign assets under the optimal plan, despite the lack of international risk sharing against domestic productivity shocks. Thus, under a policy of domestic producer price stability, even though each individual unconstrained household may borrow or save in response to idiosyncratic income shocks, unconstrained households as a whole do not borrow or save and on average consume their share of national income in every period – just like hand-to-mouth households on average do, by construction (see Eq. (34)). Consequently, unconstrained households as a whole behave symmetrically to hand-to-mouth households as a whole, and there is no between-group inequality, whatever the share of hand-to-mouth households in the economy.

To summarize, in our SOE-HANK, acyclical income risk implies that the monetary authority cannot affect within-group inequality, while Cole-Obstfeld elasticities imply that there are no fluctuations in between-group inequality. Eventually, under this parameter configuration monetary authority disregard inequality altogether (despite its potential welfare costs) and focuses on the other objectives of producer price stability and productive efficiency. Since the latter two do not conflict under Cole-Obstfeld elasticities, the SOE-HANK divine coincidence prevails, no matter how large the fraction of hand-to-mouth households or the extent of idiosyncratic income risk.

Fig. 1 illustrates this correspondence between optimal monetary policy in SOE-HANK with acyclical risk and SOE-RANK. To highlight how monetary policy optimally responds to aggregate shocks, all the figures plot generalized IRFs depicting those responses minus the optimal responses absent shock. The solid blue curves depict the optimal response to a negative domestic productivity shock in our SOE-HANK with acyclical risk, while the dashed-red lines depict the optimal dynamics in response to the same shock in an SOE-RANK. As the figure shows, both lines overlap in all panels, showing that a purely inward-looking policy focusing exclusively on stabilizing domestic producer prices can be optimal even when markets are incomplete.

However, away from the knife-edge case just discussed, the planner's willingness to stabilize inequality does produce additional policy trade-offs. This implies that optimal monetary policy in SOE-HANK is no longer the same as in SOE-RANK in empirically relevant cases. To illustrate this point as transparently as possible, we first depart from the special case described above by maintaining the Cole and Obstfeld (1991) elasticities, but assuming, realistically, that income risk is countercyclical rather than acyclical (Storesletten et al., 2004; Nakajima and Smirnyagin, 2019). Fig. 2 depicts the differences between the optimal dynamics in response to a negative productivity shock in our SOE-HANK with countercyclical risk (solid-blue lines) and in SOE-RANK (dashed-red lines). The income-risk cyclicity parameter is set to $\varphi = 5$ (instead of $\varphi = 0$), following Acharya et al. (2023). Optimal policy in SOE-HANK now implements a smaller decline in output and the real exchange rate relative to SOE-RANK (the blue lines are higher than the dashed red lines in Panels (d) and (c)). As a result, inflation temporarily increases, i.e., it is optimal to depart from strict producer price stability (Panel (e)). This policy is achieved by raising the interest rate significantly less than in SOE-RANK (Panel (b)). The reason why this more accommodative policy is optimal is that, under countercyclical risk, raising the interest rate as much as in SOE-RANK – and consequently appreciating the domestic currency as much – would generate too large a collapse in the demand for Home goods, hence too large an increase in inequality domestically. In contrast, a smaller appreciation supports higher demand for Home goods. It prevents output from falling as much as in SOE-RANK, mitigating the adverse impact on inequality in

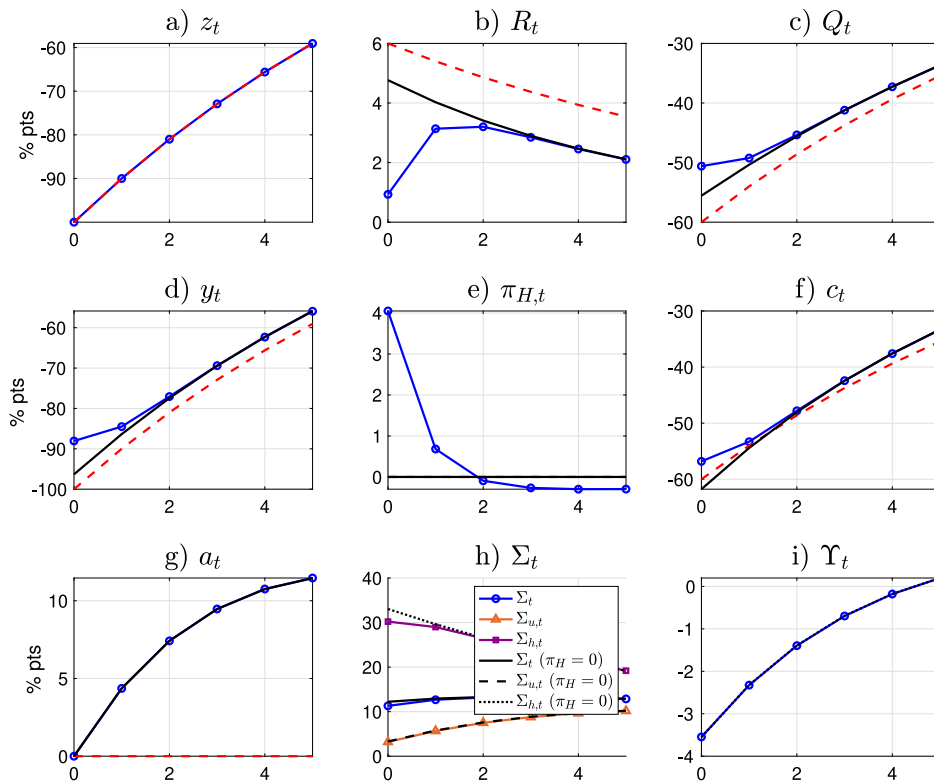


Fig. 2. SOE-HANK (solid blue with circles) vs. SOE-RANK (dashed red) vs. SOE-HANK under price stability (solid black), all under Cole-Obstfeld elasticities ($\gamma = \eta = \nu = 1$) and countercyclical risk ($\varphi = 5$). Price stability is no longer optimal in SOE-HANK. If imposed in that model, it creates excess volatility in output, the exchange rate, and inequality.

SOE-HANK. The milder output response in SOE-HANK comes at the cost of higher inflation on impact, which the planner is willing to tolerate as it prevents an even larger increase in the welfare cost of inequality Σ , (see Panel (h) of Fig. 2).

To further highlight the difference in optimal policy between SOE-HANK and SOE-RANK, Fig. 2 also displays the aggregate dynamics in SOE-HANK if producer price stability (i.e., the optimal policy in SOE-RANK) were inefficiently imposed in SOE-HANK (see the solid black lines). By virtue of the Cole-Obstfeld elasticities, between-group inequality would be unaffected (the black and blue lines lie exactly on top of each other in Panel (i)). However, within-group inequality would increase by more (Panel (h)), leading to welfare losses.

While we have restricted our attention to the Cole-Obstfeld elasticities so far, the property that optimal policy stabilizes output and the real exchange rate more in SOE-HANK than in SOE-RANK after an aggregate productivity shock when idiosyncratic risk is countercyclical is robust across plausible trade elasticities and elasticities of intertemporal substitution. To illustrate this, Fig. 3 plots those optimal responses under the standard parameterization adopted by, e.g., Egorov and Mukhin (2023), who set $\gamma = 2, \eta = 1.5$ and $\nu = 4$. While Egorov and Mukhin (2023) assume CRRA preferences, the local CRRA and CARA coefficients are identical under our normalization of the steady state. Unsurprisingly, price stability is no longer optimal even in SOE-RANK under this parameterization. This is because the terms-of-trade manipulation motive now kicks in and clashes with the objectives of stabilizing producer prices and the output gap, breaking the open-economy divine coincidence. To elaborate, under Cole-Obstfeld elasticities ($\chi = 2 - \alpha$), national income (to first order) is given by $\hat{p}_{H,t} + \hat{y}_t = (1 - \alpha)\hat{c}_t + \alpha\hat{Q}_t$, and optimal policy exactly replicates the flexible-price allocation wherein $\hat{c}_t = \hat{Q}_t = \hat{p}_{H,t} + \hat{y}_t$ (see Appendix F for details). Whenever the trade elasticity is larger ($\chi > 2 - \alpha$), strict producer price stability would still replicate the flexible-price allocation but, in so doing, would produce excess expenditure-switching away from home-produced goods after a negative aggregate productivity shock. Instead, the central bank engages in terms-of-trade manipulation and implements more exchange-rate stability than in the flexible-price allocation. Against this backdrop, stabilizing inefficient fluctuations in inequality in SOE-HANK provides, in this calibration, an additional motive for stabilizing the real exchange rate. In fact, this motive plays out more strongly at trade elasticities higher than Cole-Obstfeld, as can be seen by comparing the differences between the red-dashed and blue lines of Panel (c) in Figs. 2 and 3. This is because when income risk is countercyclical, output and consumption inequality become more sensitive to exchange rate fluctuations when trade elasticities are high.

Finally, Figs. 2 and 3 show that the impact of aggregate productivity shocks on inequality under countercyclical risk works primarily through within-group inequality (Panel (h)). In comparison, between-group inequality (Panel (i)) move very little, even away from Cole-Obstfeld elasticities. This is unlike what happens after a world interest rate shock, as we show next.

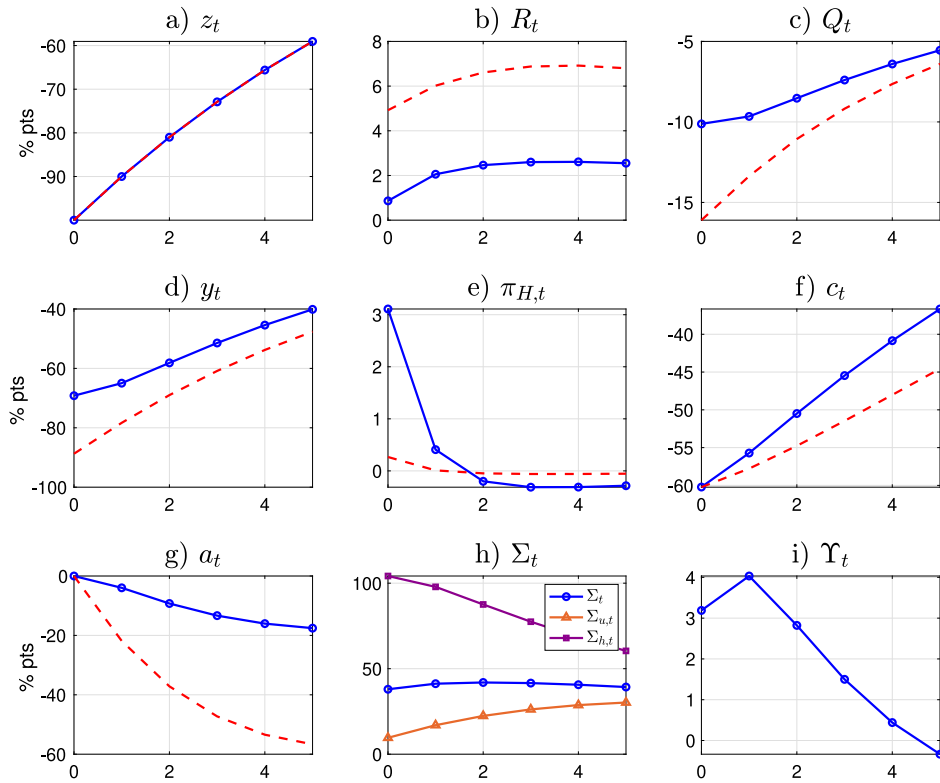


Fig. 3. SOE-HANK (solid blue with circles) vs. SOE-RANK (dashed red) under non-Cole-Obstfeld elasticities ($\gamma = 2, \eta = 1.5, \nu = 4$) and countercyclical risk ($\varphi = 5$). The difference in optimal policies is qualitatively similar but quantitatively magnified relative to the Cole-Obstfeld case.

4.2.2. World interest-rate shock

We may now turn to the effects of changes in the world interest rate R_t^* . Before examining optimal monetary policy, let us briefly elaborate on the transmission of those shocks to the domestic economy, holding domestic monetary policy unchanged. Consider a fall in R_t^* , whose direct effect (holding the domestic real interest rate unchanged) is an inflow of capital into the domestic economy. This inflow causes a real appreciation (see (27)) and thereby an increase in the relative price of home-produced goods (see (13)). In SOE-RANK, this lowers the real marginal cost of firms and is thus deflationary (see Eq. (21)), urging the central bank to accommodate the shock by cutting the domestic interest rate. Furthermore, in SOE-HANK, under plausible trade elasticities ($\chi > 1$), expenditure switching away from home-produced goods generates a fall in national income that disproportionately lowers the consumption of hand-to-mouth households via the real-income channel – furthering the central bank’s incentives to cut interest rate in order to mitigate the increase in between-group inequality. Ultimately, while the central bank cuts interest rates in response to this shock in both SOE-RANK and in SOE-HANK, the cut is more pronounced in SOE-HANK. That is, monetary policy in SOE-HANK optimally chooses to implement a smaller interest rate differential, which in turn results in a smaller appreciation of the Home currency.

Fig. 4 depicts the optimal dynamics graphically in the case with non-Cole-Obstfeld elasticities and countercyclical risk. The blue lines depict dynamics in SOE-HANK with countercyclical risk, the red dashed lines depict dynamics in SOE-RANK, and the dotted blue curves depict optimal dynamics in SOE-HANK with acyclical risk. As with the aggregate productivity shock under countercyclical income risk, the optimal policy in SOE-HANK prevents the exchange rate from appreciating as much as in SOE-RANK. This policy is implemented by cutting domestic interest rates more to reduce the interest rate differential between the domestic economy and the rest of the world. As Panel (h) and (i) show, both *within-group* and *between-group* inequality now play an important role in explaining the difference in the degree of exchange-rate stabilization between SOE-HANK and SOE-RANK. This is in contrast with what happened after a productivity shock, where between-group inequality played almost no role. The reason for this difference in inequality responses across aggregate shocks is that the world interest rate shock directly hits the real exchange rate. These changes in real exchange rates result in large fluctuations in the consumption of hand-to-mouth households via the real-income channel.

5. Conclusion

In this paper, we have presented a tractable framework to study the additional tradeoffs faced by the central bank in a small open economy wherein households are heterogeneous due to uninsured idiosyncratic risk and unequal access to bond markets.

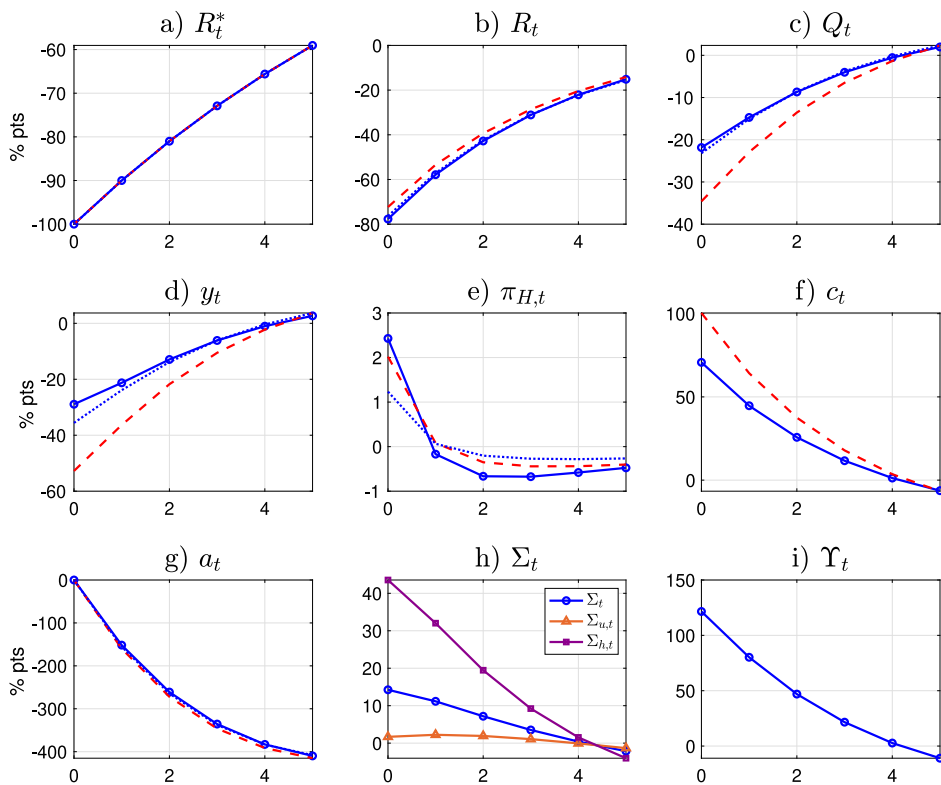


Fig. 4. SOE-HANK (blue) vs. SOE-RANK (dashed red) under non-Cole-Obstfeld elasticities ($\gamma = 2, \eta = 1.5, \nu = 4$) and countercyclical risk ($\varphi = 5$, solid blue with circles) or acyclical risk ($\varphi = 0$, dotted blue). The difference in optimal policies is qualitatively similar but quantitatively magnified relative to the Cole-Obstfeld case.

We first showed that those features are not sufficient, by themselves, to generate differences in optimal monetary policy between the baseline SOE-HANK model and the SOE-RANK model. In fact, under Cole-Obstfeld elasticities and acyclical earnings risk, the optimal policy implements strict domestic producer price stability after aggregate productivity shocks in both models. However, we also found that moving away from those assumptions and adopting more realistic values of the elasticities of substitution and the cyclical risk could lead to substantial differences between SOE-RANK and SOE-HANK in terms of optimal policy and aggregate outcomes: in those scenarios, the optimal policy implemented significantly less exchange-rate volatility than in the SOE-RANK benchmark.

Declaration of competing interest

Authors declare that they have no relevant or material financial interests related to the research described in this paper.

Acknowledgments

Edouard Challe acknowledges financial support from the French National Research Agency (ANR-20-CE26-0018-01) and the hospitality of the Bank of Finland.

Appendix A. Supplementary data

The technical appendix to this paper, containing derivation details and the proof of Proposition 1, can be found online at <https://doi.org/10.1016/j.jinteco.2025.104076>.

Data availability

Inequality and Optimal Monetary Policy in the Open Economy (Original data) (Mendeley Data)

References

- Acharya, Sushant, Challe, Edouard, Dogra, Keshav, 2023. Optimal monetary policy according to HANK. *Am. Econ. Rev.* 113 (7), 1741–1782.
- Acharya, Sushant, Dogra, Keshav, 2020. Understanding HANK: Insights from a PRANK. *Econometrica* 88 (3), 1113–1158.
- Acharya, Sushant, Pesenti, Paolo A., 2024. Spillovers and Spillbacks. Technical Report, National Bureau of Economic Research.
- Aiyagari, S. Rao, 1994. Uninsured idiosyncratic risk and aggregate saving. *Q. J. Econ.* 109 (3), 659–684.
- Auclert, Adrien, Bardóczy, Bence, Rognlie, Matthew, 2023. MPCs, MPEs, and multipliers: A trilemma for New Keynesian models. *Rev. Econ. Stat.* 105 (3), 700–712.
- Auclert, Adrien, Rognlie, Matthew, Souchier, Martin, Straub, Ludwig, 2024. Exchange rates and monetary policy with heterogeneous agents: Sizing up the real income channel. Working Paper Series 28872, National Bureau of Economic Research.
- Backus, David K., Smith, Gregor W., 1993. Consumption and real exchange rates in dynamic economies with non-traded goods. *J. Int. Econ.* 35 (3–4), 297–316.
- Bayer, Christian, Müller, Gernot, Kriwoluzky, Alexander, Seyrich, Fabian, 2023. A HANK2 model of monetary unions. Technical Report, CEPR Discussion Papers.
- Benigno, Pierpaolo, 2009. Price stability with imperfect financial integration. *J. Money, Credit, Bank.* 41, 121–149.
- Benigno, Gianluca, Benigno, Pierpaolo, 2003. Price stability in open economies. *Rev. Econ. Stud.* 70 (4), 743–764.
- Benigno, Gianluca, Benigno, Pierpaolo, 2006. Designing targeting rules for international monetary policy cooperation. *J. Monet. Econ.* 53 (3), 473–506.
- Bhandari, Anmol, Evans, David, Golosov, Mikhail, Sargent, Thomas J., 2021. Inequality, business cycles, and monetary-fiscal policy. *Econometrica* 89 (6), 2559–2599.
- Blanchard, Olivier Jean, Galí, Jordi, 2007. Real wage rigidities and the new keynesian model. *J. Money, Credit, Bank.* 39, 35–65.
- Chen, Sihao, Devereux, Michael B., Shi, Kang, Xu, Juanyi, 2023. Consumption heterogeneity and monetary policy in an open economy. *J. Monet. Econ.*
- Clarida, Richard, Galí, Jordi, Gertler, Mark, 2001. Optimal monetary policy in open versus closed economies: an integrated approach. *Am. Econ. Rev.* 91 (2), 248–252.
- Clarida, Richard, Galí, Jordi, Gertler, Mark, 2002. A simple framework for international monetary policy analysis. *J. Monet. Econ.* 49 (5), 879–904.
- Cole, Harold L., Obstfeld, Maurice, 1991. Commodity trade and international risk sharing: How much do financial markets matter? *J. Monet. Econ.* 28 (1), 3–24.
- Corsetti, Giancarlo, Dedola, Luca, Leduc, Sylvain, 2023. Exchange rate misalignment and external imbalances: What is the optimal monetary policy response? *J. Int. Econ.* 103771.
- Corsetti, Giancarlo, Pesenti, Paolo, 2005. International dimensions of optimal monetary policy. *J. Monet. Econ.* 52 (2), 281–305.
- Dávila, Eduardo, Schaab, Andreas, 2023. Optimal Monetary Policy with Heterogeneous Agents: Discretion, Commitment, and Timeless Policy. Technical Report, National Bureau of Economic Research.
- De Ferra, Sergio, Mitman, Kurt, Romei, Federica, 2020. Household heterogeneity and the transmission of foreign shocks. *J. Int. Econ.* 124, 103303.
- De Paoli, Bianca, 2009a. Monetary policy and welfare in a small open economy. *J. Int. Econ.* 77 (1), 11–22.
- De Paoli, Bianca, 2009b. Monetary policy under alternative asset market structures: The case of a small open economy. *J. Money, Credit, Bank.* 41 (7), 1301–1330.
- Druedahl, Jeppe, Ravn, Søren Hove, Sunder-Plassmann, Laura, Sundram, Jacob Marott, Waldstrøm, Nicolai, 2022. The transmission of foreign demand shocks. Technical Report, Working Paper 2022.
- Egorov, Konstantin, Mukhin, Dmitry, 2023. Optimal policy under dollar pricing. *Am. Econ. Rev.* 113 (7), 1783–1824.
- Erceg, Christopher J., Henderson, Dale W., Levin, Andrew T., 2000. Optimal monetary policy with staggered wage and price contracts. *J. Monet. Econ.* 46 (2), 281–313.
- Faia, Ester, Monacelli, Tommaso, 2008. Optimal monetary policy in a small open economy with home bias. *J. Money, Credit, Bank.* 40 (4), 721–750.
- Farhi, Emmanuel, Werning, Ivan, 2012. Dealing with the Trilemma: Optimal Capital Controls with Fixed Exchange Rates. Technical Report, National Bureau of Economic Research.
- Farhi, Emmanuel, Werning, Iván, 2019. Monetary policy, bounded rationality, and incomplete markets. *Am. Econ. Rev.* 109 (11), 3887–3928.
- Galí, Jordi, 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*, second ed. Princeton University Press.
- Galí, Jordi, Monacelli, Tommaso, 2005. Monetary policy and exchange rate volatility in a small open economy. *Rev. Econ. Stud.* 72 (3), 707–734.
- Guo, Xing, Ottonello, Pablo, Perez, Diego J., 2023. Monetary policy and redistribution in open economies. *J. Political Econ. Macroecon.* 1 (1), 191–241.
- Iyer, Tara, 2017. Optimal monetary policy in an open emerging market economy. Technical Report, FRB of Chicago Working Paper No. WP-2016-6.
- Le Grand, François, Martin-Baillon, Alais, Ragot, Xavier, 2022. Should monetary policy care about redistribution? Optimal fiscal and monetary policy with heterogeneous agents. *Sci. Po Work. Pap.*
- McKay, Alisdair, Wolf, Christian, 2022. Optimal policy rules in HANK. Technical Report, Working Paper, FRB Minneapolis.
- Nakajima, Makoto, Smirnyagin, Vladimir, 2019. Cyclical labor income risk. Working Paper 2019.
- Nisticò, Salvatore, 2016. Optimal monetary policy and financial stability in a Non-Ricardian economy. *J. Eur. Econ. Assoc.* 14 (5), 1225–1252.
- Oskolkov, Aleksei, 2023. Exchange rate policy and heterogeneity in small open economies. *J. Int. Econ.* 142, 103750.
- Rotemberg, Julio J., 1982. Monopolistic price adjustment and aggregate output. *Rev. Econ. Stud.* 49 (4), 517–531.
- Schmitt-Grohé, Stephanie, Uribe, Martin, 2003. Closing small open economy models. *J. Int. Econ.* 61 (1), 163–185.
- Storesletten, Kjetil, Telmer, Chris, Yaron, Amir, 2004. Cyclical dynamics in idiosyncratic labor market risk. *J. Political Econ.* 112 (3), 695–717.