

Household Heterogeneity and the Performance of Monetary Policy Frameworks

by Edouard Djeutem,¹ Mario He,² Abeer Reza² and Yang Zhang²

¹International Economic Analysis Department

²Canadian Economic Analysis Department

Bank of Canada

djee@bankofcanada.ca, MHe@bankofcanada.ca, yangzhang@bankofcanada.ca



Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

DOI: <https://doi.org/10.34989/swp-2022-12> | ISSN 1701-9397

©2022 Bank of Canada

Acknowledgements

The authors would like to acknowledge helpful suggestions from Rhys Mendes, Jose Dorich, Edouard Challe, Sushant Acharya, Geoffrey Dunbar, Julien Bengui and seminar participants at the Bank of Canada and the Canadian Economics Association Annual Conference. Opinions expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Canada or its staff. Any remaining errors are ours.

Abstract

We compare the performance of alternative monetary policy frameworks (inflation targeting, average inflation targeting, price level targeting and nominal GDP level targeting) in a tractable HANK model where incomplete financial markets and idiosyncratic earnings risk introduce precautionary savings and consumption inequality. Financial market incompleteness generates an additional source of societal welfare loss due to cyclical fluctuations in inequality on top of those from inflation and output volatility. We find that history-dependent policies are preferred in this framework. However, if central banks put a high weight on curbing inequality, AIT and IT can be preferred over PLT.

Topics: Monetary policy framework; Monetary policy transmission; Monetary policy and uncertainty; Economic models

JEL codes: D31, D52, E21, E31, E52, E58

1 Introduction

There is a growing discussion in the literature on the role monetary policy plays in affecting economic inequality. Empirical research has demonstrated that contractionary monetary policy increases inequality in the U.S. (Coibion, Gorodnichenko, Kueng and Silvia 2017) and internationally (Furceri, Loungani and Zdzienicka 2018).¹ At the same time, a number of theoretical papers have suggested that optimal monetary policy seeks to reduce inequality in addition to stabilizing inflation and the output gap (Acharya, Challe and Dogra 2021, Bhandari, Evans, Golosov and Sargent 2018). Since inequality increases during recessions (Nakajima and Smirnyagin 2019), the interaction between monetary policy strategy and inequality can be more salient when policy rates become constrained by the effective lower bound (ELB). The secular decline in the neutral rates of interest in industrial economies suggest that these economies could face a new challenge of encountering the ELB constraint more frequently in the future. This has prompted a review of alternative monetary policy strategies in a number of central banks, notably, the U.S. Federal Reserve (Svensson 2020), the European Central Bank (Cecioni et al., 2021), and the Bank of Canada (Dorich, Mendes and Zhang 2020).

Among alternative monetary policy frameworks, history-dependent policy frameworks or “makeup” strategies have recently been highlighted for their promise in weathering ELB episodes (Nakata, Schmidt and Budianto 2020, Amano, Gnocchi, Leduc and Wagner 2020). These frameworks ensure that past inflation shortfalls are followed by temporarily higher future inflation and through the expectations channel. This results in better inflation stabilization, specifically at the ELB. However, most of the discussions on alternative policy frameworks have so far focused on the inflation and output trade-off facing the central bank, with little concern about inequality.²

This paper contributes to the discussion by focusing on analyzing the effects of history-dependent policy frameworks in an environment featuring household heterogeneity and consumption inequality. In particular, we contribute to this literature by introducing an additional trade-off involving inequality that originates from incomplete financial markets. We adopt the tractable HANK model of Acharya et al. (2021), where incomplete financial markets and idiosyncratic earnings risk lead to consumption inequality in response to aggregate shocks at business cycle frequencies. From a societal perspective, inequality arising from a market failure is costly. We compare the performance of flexible inflation targeting (IT) with history-dependent policy frameworks, such as price-level targeting (PLT) and average-inflation targeting (AIT) in a modelled economy with inequality. We also consider nominal GDP level (NGDPL) targeting.

¹Coibion et al. (2017) show that contractionary monetary policy systematically increases inequality in labour earnings, income, and consumption in the U.S. Furceri et al. (2018) show that contractionary monetary policy shocks increase income inequality across 32 advanced and emerging market economies. Holm, Paul and Tischbirek (2020) show that Norwegian households’ consumption and income responses to monetary policy shocks differ across their distribution of asset holdings.

²A notable exception is Feiveson, Gornemann, Hotchkiss, Mertens and Sim (2020), which suggests that stabilization benefits from makeup strategies could generate disproportionate improvements for historically disadvantaged households.

Two crucial assumptions keep the model tractable despite household heterogeneity introduced through uninsurable idiosyncratic earnings risk. First, idiosyncratic shocks are assumed to be log-normally distributed. Second, utility in this model follows a constant absolute risk aversion (CARA) functional form. The combination of the CARA-normal assumptions allows linear aggregation of individual policy functions. As a result, there is no need to track an individual’s history of past shocks in order to infer aggregate dynamics in the model. Rather, the aggregate dynamics of the model can be expressed with a four-equation New Keynesian (NK) model that can then be solved using perturbation techniques. These assumptions therefore significantly lessen the computational burden while capturing the essence of household heterogeneity.

In addition to the usual inter-temporal substitution channel, monetary policy in this model can influence aggregate economic outcomes through two more channels that affect precautionary savings. First, the *income risk channel* directly relates episodes of lower output to those with higher uninsurable idiosyncratic income risk, which motivates precautionary savings. Policy accommodation to lower output fluctuations can also lower precautionary savings by reducing income risk. Second, the *self-insurance channel* allows agents to partially self-insure against their idiosyncratic income risk by participating in the bond market. Lower interest rates reduce precautionary savings through this channel by making it easier for agents to use the aggregate bond for partial self-insurance. In other words, when interest rates are low, there is reduced pass-through from idiosyncratic income risk to consumption risk. Combined with the inter-temporal substitution channel, these two novel channels amplify monetary policy impact in the HANK economy relative to that in a Representative Agent New Keynesian (RANK) model where these two channels are absent.³ Moreover, in this model, monetary easing reduces inequality while tightening increases inequality, consistent with empirical evidence (Coibion et al. 2017, Furceri et al. 2018).

The aggregate social welfare loss function in this model is tractable and dependent on two parts—one capturing average welfare across households, and another capturing the cost of consumption inequality across households that can be summarized using a sufficient statistic. This allows us to examine the effect of alternative policy frameworks on social welfare in the presence of consumption inequality. Specifically, the social welfare loss function includes welfare costs from two market imperfections. The first is the standard New Keynesian market failure due to the presence of nominal price rigidities under monopolistically competitive markets. The second is the cost of consumption inequality that stems from imperfections in financial markets that restrict the ability of agents to insure their idiosyncratic risks.

We implement alternative policy frameworks (IT, AIT, PLT, and NGDPL) by characterizing each framework using a policy rule specified according to its respective

³More specifically, in a RANK economy, financial markets are complete and all idiosyncratic risks are insured against. This leaves the inter-temporal substitution channel as the only channel for monetary policy transmission. This amplification property of the tractable HANK model is qualitatively consistent with the findings in Kaplan and Violante (2018).

target object (inflation, average inflation, the price level, and nominal GDP levels, respectively), similar to Amano et al. (2020) and Feiveson et al. (2020). We then optimize the policy rules over a grid of coefficient values in the policy reaction function to minimize a corresponding loss function of the central bank. To account for uncertainty around the central bank preferences, we consider three classes of loss functions. The first is the model-consistent, *social welfare-based loss* function that includes the societal cost of consumption inequality. We complement welfare analysis by considering a second class of loss functions including (i) a simple *ad hoc loss* function that depends on the volatilities of inflation and the output gap, and (ii) an alternative *ad hoc loss* function that includes an additional term representing consumption inequality. The first class of ad hoc loss is easy to monitor and communicate to the public, and has been extensively used in the literature in evaluating alternative monetary policy strategies (Bernanke, Kiley and Roberts 2019, Svensson 2020). The second loss function is useful to provide additional intuition about the interaction between inequality and alternative policy frameworks. The second class of loss functions is regime-specific functions delegated to the central bank, as in Svensson (2020) and Dorich et al. (2020). In this case, we use the delegated loss functions only to choose the optimized parameters of the policy rule. We then evaluate the frameworks using volatilities of several key economic variables, while looking for alternative frameworks that stabilize a broad range of variables. The advantage of our tractable HANK model is that we can compare welfare loss under delegated loss functions.

Our main finding is that history-dependent policies dominate in the HANK economy when the central bank considers the social welfare-based loss function. Two main properties of the HANK model underpin this result. First, even though the HANK economy features amplified output responses to aggregate shocks, the relationship between output and inflation via the New Keynesian Phillips Curve (NKPC) remains the same as the standard RANK model. Therefore, the amplified transmission channels of monetary policy recently highlighted in the HANK literature (Kaplan, Moll and Violante 2018, Acharya and Dogra 2020, Bilbiie 2019, Ravn and Sterk 2016) do not alter the trade-off between inflation and the output gap compared to the standard model. In this case, the seminal findings of Nessen and Vestin (2005) and Vestin (2006) also apply to the HANK model—history-dependent or “makeup” strategies result in lower inflation volatility through the expectation channel. We demonstrate this result clearly by considering an ad hoc loss function with inflation and output gap only.⁴

Second, for the benchmark calibration considered in this paper, the welfare costs due to incomplete financial markets are small compared to that due to price rigidities. In particular, for a second-order approximation of the social welfare loss, the weight on inequality arising from a failure in financial markets is small compared to the weight on inflation stemming from nominal rigidities. The benefits from history-dependent frameworks accruing to a reduction in inflation outweigh any possible trade-offs arising from inequality. Acharya et al. (2021) and Bhandari et al. (2018)

⁴Note that we do not include rule-of-thumb price setters in our analysis. See Nessen and Vestin (2005) and Dorich et al. (2020) for analysis involving rule-of-thumb price setters.

find that the optimal policy in the HANK model is more accommodating during recessions compared to the RANK case specifically because the central bank realizes that a boost in output also improves inequality. Although our model encompasses this result for the benchmark calibration, it does not change the ranking of alternative policy frameworks.

Makeup strategies also dominate the ranking when considering a simple ad hoc loss function involving inflation and output gap volatility. However, if the central bank faces an ad hoc loss function that includes a large weight on consumption inequality on top of output gap and inflation, IT and AIT become more favourable than PLT. This is because the central bank faces a trade-off in the income risk and pass-through channels that generate consumption inequality. While income risk is lowered by stabilizing the output gap, the pass-through to consumption is lowered when interest rates are low and stable. PLT involves higher volatility in interest rates while stabilizing prices and the output gap, vis-à-vis IT. In other words, agents find it harder to partially self-insure themselves against idiosyncratic risk in the PLT regime. This leads to unstable pass-through from income risk to consumption risk and hence higher consumption inequality in PLT. At one point, the higher cost of inequality in PLT outweighs its benefits from price and output stabilization vis-à-vis IT and AIT in this special case.

Introducing an ELB constraint preserves the relative ranking of the monetary policy frameworks despite the increase in the volatility of cyclical fluctuations. In fact, history-dependent policies are less costly vis-à-vis IT (i.e., corresponds to lower loss than IT) under the ELB. However, in the special case where the central bank considers an ad hoc loss function with output gap and inflation and a high weight on consumption inequality, AIT and IT continue to outperform PLT.

Our paper contributes to two recent strands of literature on monetary policy: inequality and the lower bound. First, Nakata et al. (2020) and Benchimol and Bounader (2019) consider the performance of history-dependent policy frameworks at the ELB in models with bounded rationality. Amano et al. (2020) provides a similar analysis in a Two-Agent New Keynesian model with a focus of AIT at the ELB. However, these models do not address the performance of alternative regimes in the presence of inequality. We contribute to this literature by comparing the performance of alternative policy frameworks at the ELB where the central bank loss function directly includes the welfare-cost of inequality.

Second, Acharya et al. (2021) and Bhandari et al. (2018) consider optimal monetary policy under commitment in HANK settings. They find that optimal policy seeks to mitigate increases in inequality when incomplete financial markets keep agents from insuring their idiosyncratic risks. On the other hand, Le Grand, Martin-Baillon and Ragot (2020) argues that the redistributive role of monetary policy is limited when a rich set of fiscal tools are available. We contribute to this literature by considering specific policy frameworks (IT, AIT, PLT) implemented through optimized simple policy rules, and analyze the performance of tractable HANK models under the effects of the ELB.

Our paper is closest to the contributions of Feiveson et al. (2020), which also implements alternative policy frameworks using simple policy rules in a HANK model. However, we improve upon their findings by (a) directly comparing social welfare costs including costs to inequality across various frameworks, (b) controlling for policy rule uncertainty by considering a large combination of policy parameter space through efficient frontiers, and (c) exploring uncertainty around central bank loss functions characterization by including framework-specific delegated loss functions and societal welfare-based loss function.

The rest of the paper is organized as follows. Section 2 lays out the core of the model. Section 3 develops the second-order approximation of the social welfare function with costly consumption inequality. Section 4 provides the calibration strategy. Section 5 demonstrates the properties of the model using impulse response functions under alternative policy frameworks. Section 6 discusses ranking of alternative policy frameworks using both ad hoc loss and model-consistent welfare loss functions, and section 7 concludes.

2 A tractable HANK model

In this section, we introduce a tractable New Keynesian model with CARA preferences and idiosyncratic income risk following Acharya and Dogra (2020) and Acharya et al. (2021), and extend it to include both household preference and price markup shocks. While firms and policymakers solve the same problem as in the standard model, the crucial innovation is on the demand block of the model in which households are heterogeneous due to idiosyncratic income risk. This class of models shares the key economic mechanisms of a complex quantitative HANK model (as in Kaplan et al. (2018)), while at the same time keeps the tractability of a prototypical New Keynesian model.

2.1 Environment

Household heterogeneity. Time is discrete and indexed by t . The economy is inhabited by a continuum of households of measure one indexed by $i \in [0, 1]$. Each household dies with probability $1 - \vartheta$, to be replaced by a cohort of newborn of equal size $1 - \vartheta$ and indexed by s . A household derives utility from consuming goods, $c_{s,t}^i$, and minimizes the dis-utility from labor. Each household supplies labour $l_{s,t}^i$ at the prevalent wage rate w_t . The per-period utility function takes the CARA functional form, separable in consumption and labor as follows:

$$U(c_{s,t}^i, l_{s,t}^i, \xi_{s,t}^i, v_t) = v_t \left(-\frac{1}{\gamma} e^{-\gamma c_{s,t}^i} - \rho e^{-\frac{1}{\rho} (\xi_{s,t}^i - l_{s,t}^i)} \right) \quad (1)$$

where $\gamma > 0$ denotes the coefficient of absolute risk aversion, $\rho > 0$ is related to Frisch elasticity of labour supply, v_t denotes preference shocks as in Galì (2015), which is assumed to follow a Gaussian AR(1) stochastic process in logs, and $\xi_{s,t}^i$ denotes the

household-specific labour endowment shock. We refer to $\zeta_{s,t}^i$ as an uninsurable idiosyncratic income shock because this shock effectively distorts a household's labour supply decision. More formally, this shock is normally distributed with constant mean $\bar{\zeta}$ and time-varying variance σ_t^2 , so that $\zeta_{s,t}^i \sim \mathcal{N}(\bar{\zeta}, \sigma_t^2)$, where the time-varying variance is determined endogenously in the model. The choice of modelling uninsurable income risk as the only source of heterogeneity can be motivated by the fact that individual differences in employment status are arguably the largest source of earnings dispersion.

Financial markets are incomplete in this environment as only risk-free government bonds can be traded to insure against labour income risk. These imperfections in the financial market imply that households are unable to insure against idiosyncratic risk by trading financial assets. As such, market incompleteness changes households' incentives to consume and save by modifying the maximal amount of liquidity available at time t for consumption.

The household budget constraint at date t is:

$$c_{s,t}^i + q_t a_{s,t+1}^i = a_{s,t}^i + w_t l_{s,t}^i + T_t. \quad (2)$$

Income on the right-hand side consists of current holdings of risk-free bonds, $a_{s,t}^i$, labour income at wage w_t , government lump-sum transfers and net taxes, T_t . On the left-hand side, income is used to support consumption expenditure $c_{s,t}^i$ and the new holdings of financial assets at the actuarially fair price of risk-free bond, q_t , issued by financial intermediaries.

Financial intermediaries. There are a large number of financial intermediaries operating in perfectly competitive markets following Yaari (1965), Blanchard (1985), and Acharya et al. (2021). The balance sheet of these intermediaries is as follows. On the liability sides, financial intermediaries issue risk-free assets sold to households at price q_t . On the asset side, households buy government debt, B_{t+1} , at price $1/R_t$. Households only receive interest payments when they are alive. As such, their profit is given by $-\vartheta a_{t+1} + B_{t+1}$ and the budget constraint, by $-q_t a_{t+1} + \frac{B_{t+1}}{R_t} \leq 0$. The zero profit condition arising from the perfect competition assumption implies that:

$$q_t = \frac{\vartheta}{R_t}. \quad (3)$$

Households optimal choice. Households' optimal consumption and labour supply satisfy the standard Euler equation and intratemporal conditions given by:

$$1 = \mathbb{E}_t \left(\beta R_t \frac{v_{t+1}}{v_t} \frac{e^{-\gamma c_{s,t+1}^i}}{e^{-\gamma c_{s,t}^i}} \right),$$

$$w_t = \frac{e^{-\frac{1}{\rho}(\bar{\zeta}_{s,t}^i - l_{s,t}^i)}}{e^{-\gamma c_{s,t}^i}}. \quad (4)$$

Two important assumptions make it possible to aggregate the above Euler equation in a linear fashion: (1) CARA preference, and (ii) normally distributed idiosyncratic income risk. This means that infinite dimensional distributions of consumption and hours worked can be summarized by their cross-sectional averages. To see this, define the *cash-on-hand* as $x_{s,t}^i = a_{s,t}^i + w_t(\zeta_{s,t}^i - \bar{\zeta})$.⁵ An increase in cash-on-hand hence results either from larger savings or a positive draw from the idiosyncratic income shock distribution. As shown by Acharya et al. (2021), the optimal consumption and labour supply are linear in $x_{s,t}^i$ and given by:

$$\begin{aligned} c_t(x_{s,t}^i) &= c_t + \mu_t x_{s,t}^i \\ l_t(x_{s,t}^i) &= (\rho \log w_t - \gamma \rho c_t) - \gamma \rho \mu_t x_{s,t}^i + \zeta_{s,t}^i \end{aligned} \quad (5)$$

where c_t denotes the aggregate consumption and μ_t the income risk pass-through defined by difference equations determined in equilibrium.

The first equation in (5) shows that an individual household's consumption depends on aggregate consumption, c_t , the household's idiosyncratic cash-on-hand, $x_{s,t}^i$, and a common rate, μ_t , at which all households' income risk passes through to consumption. Due to market incompleteness, idiosyncratic shocks to cash-on-hand lead therefore to excess consumption variability. This feature is absent in a RANK model with complete markets where consumption depends on lifetime income instead of period-specific cash-on-hand, since idiosyncratic income is fully insured in a RANK model and does not affect consumption at time t . In addition, the identical income risk pass-through across households stems from the fact that households draw their idiosyncratic income risk from the same distribution and that this model does not have illiquid assets.

The second equation in (5) shows optimal labour supply as a function of wages, cash-on-hand, and idiosyncratic risk. We see that a household that receives an extra dollar at time t in terms of cash-on-hand optimally chooses to reduce labour supply by $\gamma \rho \mu_t$ as dictated by the wealth effects. The reduction in labour supply, however, is lower than what would prevail in a complete market RANK economy because of the missing insurance market.

Firms. The rest of the model is fairly standard. More details of the full model are provided in Appendix A. On the supply side, perfectly competitive firms produce a homogeneous final good, y_t , using differentiated intermediate goods, $y_t(j)$ for $j \in [0, 1]$, given by:

$$y_t = \left(\int_0^1 y_t(j)^{\frac{\varepsilon_t - 1}{\varepsilon_t}} dj \right)^{\frac{\varepsilon_t}{\varepsilon_t - 1}}$$

where ε_t denotes the stochastic elasticity of substitution between the differentiated intermediate goods. The demand of each specific variety j and the price index of the intermediate bundle are respectively:

$$y_t(j) = \left(\frac{p_t(j)}{P_t} \right)^{-\varepsilon_t} y_t; \quad P_t = \left(\int_0^1 p_t(j)^{1-\varepsilon_t} dj \right)^{\frac{1}{1-\varepsilon_t}}. \quad (6)$$

⁵This concept, first introduced by Deaton (1991), represents intuitively the maximal household-specific liquidity available to finance expenditures.

Intermediate goods are produced by monopolistically competitive firms endowed with a linear technology. Labour is the only input to production, $y_t(j) = z_t n_t(j)$, where z_t stands for the aggregate productivity level.⁶ Each intermediate good producer sets its prices facing a quadratic adjustment cost modelled along the lines of Rotemberg (1982). The quantities sold, net of adjustment cost, is given by:

$$y_t(j) = z_t n_t(j) - \frac{\Phi}{2} \left(\frac{p_t(j)}{p_{t-1}(j)} - 1 \right)^2 y_t. \quad (7)$$

The intermediate goods producer's problem can be written as:

$$\max_{p_t(j), n_t(j), y_t(j)} \mathbb{E}_t \sum_{t=0}^{\infty} Q_{t|0} \left\{ \frac{p_t(j)}{P_t} y_t(j) - (1 - \tau) w_t n_t(j) \right\}$$

subject to the production function in (7) and the demand function in (6), where $Q_{t|0} = \prod_{s=0}^{t-1} \frac{1}{R_s}$ is the pricing kernel, $\Phi > 0$ is the parameter that governs the nominal price rigidities in form of quadratic cost, and τ is the time-invariant payroll subsidy rate. The firm's pricing problem leads to the following expression of the New Keynesian Philips' Curve:

$$(\Pi_t - 1)\Pi_t = \frac{\varepsilon_t}{\Phi} \left(1 - \frac{\varepsilon_t - 1}{\varepsilon_t} \frac{z_t}{(1 - \tau)w_t} \right) + \mathbb{E}_t \left(\frac{1}{R_t} \frac{y_{t+1} z_{t+1} w_{t+1}}{y_t z_{t+1} w_t} (\Pi_{t+1} - 1)\Pi_{t+1} \right) \quad (8)$$

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

Policymakers. There are two policymakers: a fiscal authority and a monetary authority. The fiscal authority issues short-term debt, spends, and runs a tax-and-transfer system to satisfy a period-by-period government budget constraint:

$$\frac{B_{t+1}}{1 + i_t} = P_t T_t + \tau P_t w_t n_t + B_t. \quad (9)$$

where $n_t = \int_0^1 n_t(j) dj$ is the aggregate labour demand. In this paper, we abstract from issues related to policies that aim at changing tax instruments over the business cycle to achieve a given redistributive goal. A monetary policy authority sets the interest rate according to a simple policy rule that is subject to an effective lower bound on the policy rate. We consider policy rules that differ across alternative policy frameworks described in details in section 5.1.

Aggregation and market clearing. Let $m_{s,t}^i$ be a generic variable for a household i of a cohort s at time t . We denote the aggregate counterpart, m_t of $m_{s,t}^i$, as follow:

$$m_t = \sum_{s=-\infty}^t (1 - \vartheta) \vartheta^{s-t} \int_i m_{s,t}^i di \quad (10)$$

so as to sum over all households within each cohort. For example, aggregate labour supply is defined as $l_t = \sum_{s=-\infty}^t (1 - \vartheta) \vartheta^{s-t} \int_i l_{s,t}^i di$.

⁶In absence of productivity shock, we set z_t to the steady state value, z , for all t .

In equilibrium, markets clear at all dates. This means that demand equals supply in all markets including the labour market, $n_t = l_t$, the good market, $c_t = y_t$, and the bond market, $a_t = 0$. Here n_t refers to labour demanded by all firms, l_t aggregate hours supplied by all households alive at time t , and a_t aggregate savings in risk-free assets. For simplicity, we further assume, as in Acharya et al. (2021), that government debt is in zero net supply, so that $B_t = 0$ for all t .

Aggregate and individual shocks. There are aggregate demand shocks and aggregate supply shocks. Aggregate supply shocks are markup shocks, ε_t and aggregate productivity shocks z_t given by:

$$\begin{aligned}\log \varepsilon_t &= \rho_\varepsilon \log \varepsilon_{t-1} + u_{\varepsilon,t}, \\ \log z_t &= \rho_z \log z_{t-1} + u_{z,t},\end{aligned}\tag{11}$$

where $\rho_z \in (0,1)$, $\rho_\varepsilon \in (0,1)$, and $u_{z,t}$ and $u_{\varepsilon,t}$ are i.i.d. mean-zero innovations and standard deviations σ_ε and σ_z respectively. Note that in the remainder of the paper, we consider that demand and markup shocks contribute to business cycles of the Canadian economy and abstract from productivity shocks in this version of the model.

We model aggregate demand shocks by preference shocks given by the process:

$$\log v_t = \rho_v \log v_{t-1} + u_{v,t}.\tag{12}$$

where $\rho_v \in (0,1)$, $u_{v,t}$ is an i.i.d. mean-zero innovation and standard deviation σ_v . Note that all the innovation processes, $u_{z,t}$, $u_{\varepsilon,t}$, and $u_{v,t}$, are uncorrelated with each other over time.

Each household's idiosyncratic risk is drawn from the same distribution with a variance that moves endogenously with the business cycle. In particular, income risk evolves according to the following equation:

$$w_t^2 \sigma_t^2 = w^2 \sigma^2 e^{2\phi(y_t - y)}\tag{13}$$

where w_t is the wage rate, σ_t^2 is the variance of income risk, y_t is the aggregate output at time t , and w , σ^2 , and y are their steady state counterparts. The coefficient ϕ measures the cyclicalness of income risk. When $\phi < 0$, income risk is countercyclical, implying that income risk is higher in times of low economic activity.

2.2 Linearized model

In equilibrium, despite the presence of aggregate shocks, this model boils down to a system of four equations characterizing the evolution of the key macroeconomic

variables. The log-linearized model around the steady state is given by:⁷

$$\begin{aligned}
\hat{y}_t &= \left(1 - \frac{\phi\Lambda}{\gamma}\right) \mathbb{E}_t \hat{y}_{t+1} - \frac{1}{\gamma} \left(i_t - \mathbb{E}_t \pi_{t+1}\right) - \frac{\Lambda}{\gamma} \mathbb{E}_t \hat{\mu}_{t+1} + \frac{1}{\gamma} (1 - \rho_v) \hat{v}_t, \\
\hat{\mu}_t &= -\frac{(1 - \tilde{\beta})\gamma\rho w}{1 + \gamma\rho w} \frac{\hat{w}_t}{w} + \tilde{\beta} (\mathbb{E}_t \hat{\mu}_{t+1} + i_t - \mathbb{E}_t \pi_{t+1}), \\
\hat{y}_t &= \frac{\rho}{1 + \gamma\rho} \frac{\hat{w}_t}{w} + \frac{y}{1 + \gamma\rho} \hat{z}_t, \\
\pi_t &= \frac{\tilde{\beta}}{\vartheta} \mathbb{E}_t \pi_{t+1} + \kappa (\hat{y}_t - \hat{y}_t^n) + \hat{\varepsilon}_t,
\end{aligned} \tag{14}$$

where $\Lambda = \gamma^2 \mu^2 w^2 \sigma^2$ is the steady state consumption risk, R is the steady state real interest rate, $\tilde{\beta} = \vartheta/R$ is the steady state price of an actuarial bond, $\kappa = \frac{\varepsilon}{\Phi} \frac{1+\gamma\rho}{\rho}$ is the slope of the Phillips curve, $\pi_t = \log \Pi_t$ denotes the inflation rate, and $\hat{y}_t^n = \frac{\rho+y}{1+\gamma\rho} \hat{z}_t$ is the log-linearized flexible price output. The first two equations are log-linearized forms of equations ((11) and (12) from Acharya et al. (2021)), after imposing the market-clearing conditions. The third equation expresses the production function, taking into account labour market clearing conditions. The fourth equation is the New Keynesian Philips' curve.

Uninsurable idiosyncratic risk creates precautionary savings through two distinct channels. To see this, we rewrite the first equation in (14) as follow:

$$\hat{y}_t = \underbrace{\mathbb{E}_t \hat{y}_{t+1} - \frac{1}{\gamma} \mathbb{E}_t (i_t - \pi_{t+1})}_{\text{inter-temporal substitution}} - \underbrace{\frac{\Lambda}{\gamma} \mathbb{E}_t \hat{\mu}_{t+1}}_{\text{self-insurance}} - \underbrace{\frac{\phi\Lambda}{\gamma} \mathbb{E}_t \hat{y}_{t+1}}_{\text{income risk}} + \frac{1}{\gamma} (1 - \rho_v) \hat{v}_t \tag{15}$$

Precautionary saving

First, the *income risk channel*, $\frac{\phi\Lambda}{\gamma} \hat{y}_{t+1}$, directly relates expected output movements to the magnitude of precautionary savings through income risk. When income risk is countercyclical, i.e., $\phi < 0$, a recession implies a *rise* in income risk. In the absence of complete markets, forward-looking agents start saving today in anticipation of higher income risk in the next period. This increase in precautionary savings puts downward pressure on output, *amplifying* the effect of the recession vis-à-vis the RANK case.⁸

Second, the *self-insurance channel*, $\frac{1}{\gamma} \Lambda \hat{\mu}_{t+1}$, depends on the evolution of the pass-through from income risk to consumption risk, μ_t . To understand this channel, we can re-write the equation for the dynamics of the pass-through variable μ_t from equation (14) as follows:

$$\hat{\mu}_t = \mathbb{E}_t \sum_{j=0}^{\infty} \tilde{\beta}^j \left[\tilde{\beta} (i_{t+j} + \pi_{t+j+1}) - (1 - \tilde{\beta}) \frac{\gamma\rho w}{1 + \gamma\rho w} \frac{\hat{w}_{t+j}}{w} \right]$$

⁷Appendix A describes in detail the derivation of the linearized model. Tables (10) and (11), in the appendix, summarize the model.

⁸In contrast, when income risk is procyclical, i.e., $\phi > 0$, the reverse occurs, and income risk *dampens* the effect of the economic contraction vis-à-vis the RANK case. In the acyclical case of $\phi = 0$, this channel is absent.

In this economy, agents can partially self-insure against idiosyncratic risk by purchasing or selling an aggregate state-contingent asset. When participating in this savings vehicle is easy, agents feel less of a need to save for precautionary reasons. When real rates are lower, it is easier for households to borrow using the aggregate savings vehicle. At the same time, when expected future wages are higher, it becomes easier for agents who borrow today to repay the debt in the future. Both these conditions (lower real rate and higher expected future wages) increase borrowing in the aggregate asset, and reduce precautionary savings today. Consequently a lower portion of earnings risk is passed on to consumption, and μ_t falls.

2.3 Nested RANK model

This class of tractable HANK model naturally nests a RANK version as a special case obtained by setting the steady state consumption risk, Λ , to zero. Recall that the only source of heterogeneity comes from idiosyncratic income risk. By eliminating income risk, the demand side of the economy could be effectively characterized with a representative household. In this case, the system of equations boils down to:

$$\begin{aligned}\hat{y}_t &= \mathbb{E}_t \hat{y}_{t+1} - \frac{1}{\gamma} \left(i_t - \mathbb{E}_t \pi_{t+1} \right) + \frac{1}{\gamma} (1 - \rho_v) \hat{v}_t \\ \hat{y}_t &= \frac{\rho}{1 + \gamma \rho} \frac{\hat{w}_t}{w} + \frac{y}{1 + \gamma \rho} \hat{z}_t \\ \pi_t &= \frac{\tilde{\beta}}{\vartheta} \mathbb{E}_t \pi_{t+1} + \kappa (\hat{y}_t - \hat{y}_t^n) + \hat{\varepsilon}_t\end{aligned}\tag{16}$$

By setting Λ to zero in (14), we see that the remaining endogenous variables in this system are independent of $\hat{\mu}_t$. As such, this equation is irrelevant for the nested RANK model. This leaves the inter-temporal substitution channel as the only channel of monetary policy transmission operative in the RANK model, whereby households borrow more or save less in response to an interest change, leading to an increase or a decrease in aggregate demand. Note that the firm pricing decision encapsulated in the NKPC remains the same in the RANK and HANK models.

3 Inequality and welfare

In this section, we first characterize the utilitarian social welfare function that can be expressed in closed-form despite the fact that households are heterogeneous. We then derive a second-order approximation of this social welfare that depends on the welfare cost of consumption inequality. This expression has the disadvantage of being unobservable in the data. We then derive an equation for the evolution of consumption inequality, which is potentially observable in the data.

3.1 Micro-founded welfare function

Individual welfare. The individual welfare of a household i , from cohort s at time 0, denoted by $W_{s,0}^i$, corresponds to the present discount value of utility flows. Using

the consumption and labour supply in (5) and plugging them into the lifetime utility yields:

$$\mathbb{W}_{s,0}^i = \mathbb{E}_0 \sum_{t=0}^{\infty} (\beta \vartheta)^t \underbrace{\left[v_t \left(-\frac{1}{\gamma} e^{-\gamma y_t} - \rho e^{\frac{1}{\rho}(\bar{\xi} - n_t)} \right) \right]}_{\text{Common}} \underbrace{e^{-\mu_t x_{s,t}^i}}_{\text{Household specific}} \quad (17)$$

where β denotes the subjective discount factor. Equation (17) shows that lifetime utility changes with aggregate variables (output, y_t , aggregate labour, n_t) and individual marginal utility of consumption captured by $e^{-\mu_t x_{s,t}^i}$. As such, two households i and j only differ in this economy to the extent that they draw different sequences of uninsurable idiosyncratic income shocks (which in turn lead to differences in cash-on-hand). The welfare costs of these shocks are captured by $e^{-\mu_t x_{s,t}^i}$.

Social welfare. Acharya, Challe, and Dogra (2020) show that in this economy, a utilitarian social welfare function, \mathbb{W}_0 , is conveniently expressible as

$$\mathbb{W}_0 = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[v_t \left(-\frac{1}{\gamma} e^{-\gamma y_t} - \rho e^{\frac{1}{\rho}(n_t - \bar{\xi})} \right) \right] \Sigma_t \quad (18)$$

where the first part of the per-period utility in square bracket is the common utility across households given above, and the second part conveniently summarizes the cross-sectional variation in marginal utility of consumption across households using a simple autoregressive summary statistic, Σ_t , with the linearized equation given by:

$$\hat{\Sigma}_t = \Lambda(\hat{\mu}_t + \phi y \hat{y}_t) + \frac{\tilde{\beta}}{\beta} \hat{\Sigma}_{t-1}. \quad (19)$$

From equation (18), we see that an increase in inequality, Σ_t , reduces social welfare. Furthermore, the welfare cost of inequality, $\hat{\Sigma}_t$, is a persistent process where the innovation depends on consumption risk $\Lambda(\hat{\mu}_t + \phi y \hat{y}_t)$.

Second-order welfare approximation. We derive a quadratic loss function by taking a second-order Taylor approximation of the expected social welfare function. More precisely, we take an approximation of each variable affecting welfare in deviation to its steady state value.⁹

We further consider stationary equilibria so that the unconditional expectation of a log-linearized variable with respect to steady state is zero in the model without the ELB. The main results of our approximation are summarized in the following proposition.

⁹In a classical RANK model, the removal of monopolistic competition friction in steady state by an appropriate choice of production subsidy, τ , renders zero inflation optimal (Gali 2015, see section 4.2). This need not to be true in our HANK model. Given that higher output contributes to lower income risk, the social planner has an incentive during booms to deviate from price stability in order to reduce inequality. The HANK planner chooses a higher subsidy ($\tau > \frac{1}{\varepsilon}$) and wages ($w > 1$) compared to the social planner of a RANK economy in order to implement a zero inflation in steady state (Acharya et al. 2021).

Proposition 1 *The second-order approximation of the expected per-period social welfare is:*

$$\begin{aligned}\mathbb{L}_0 = & Var(\pi_t) + \lambda_{\tilde{y}} Var(\tilde{y}_t) + \lambda_{\Sigma} Var(\hat{\Sigma}_t) + \lambda_y Var(\hat{y}_t) \\ & + \lambda_{y,\Sigma} Cov(\hat{y}_t, \hat{\Sigma}_t) + \lambda_{z,\Sigma} Cov(\hat{z}_t, \hat{\Sigma}_t) \\ & + \lambda_{y,\nu} Cov(\hat{y}_t, \hat{\nu}_t) + \lambda_{\Sigma,\nu} Cov(\hat{\Sigma}_t, \hat{\nu}_t).\end{aligned}\tag{20}$$

where the parameters are given by:

$$\begin{aligned}\lambda_y &= \frac{\gamma\rho(y-w) + \rho(w-1) + w(y-1)}{\Phi\rho w} \\ \lambda_{y,\Sigma} &= \lambda_{y,\nu} = 2\frac{w-1}{\Phi w} \\ \lambda_{\tilde{y}} &= \frac{1+\gamma\rho}{\Phi\rho} = \frac{\bar{\xi}}{\Phi\rho} \\ \lambda_{\Sigma} &= \frac{1+\gamma\rho w}{\Phi\gamma w} \\ \lambda_{\Sigma,\nu} &= 2\lambda_{\Sigma} \\ \lambda_{z,\Sigma} &= -\frac{2}{\Phi}\end{aligned}\tag{21}$$

and $\tilde{y}_t = \hat{y}_t - \hat{y}_t^n$ is the output gap, defined as the difference between the sticky price output and its flexible price counterpart; and Φ measures the degree of nominal price rigidity introduced by the cost of price adjustment.

Let us unpack this proposition. First of all, notice that in a RANK economy where $y = w = 1$, $\hat{\Sigma}_t = 0, \forall t$ and consequently $\lambda_y = \lambda_{\Sigma} = \lambda_{y,\Sigma} = \lambda_{z,\Sigma} = \lambda_{y,\nu} = 0$, we obtain the usual expected per-period average welfare where only fluctuations in inflation and output gap contribute to reducing welfare. Second, the approximation (20) depends on two different definitions of output gaps. Given that the steady states of the sticky price allocation and the flexible price allocation are identical, the only distortion affecting the difference between \hat{y}_t and \hat{y}_t^n is the price dispersion associated with nominal rigidities. As such, $Var(\tilde{y}_t)$ measures the pure production efficiency loss associated with nominal rigidities. In contrast, $Var(\hat{y}_t)$ captures the additional effect of idiosyncratic risk on social welfare. Third, household heterogeneity and uninsurable idiosyncratic income risk affect the loss function (20) through two channels: (i) their combined effects on the coefficients of the social welfare approximation, and (ii) an introduction of an additional trade-off in the loss function between low and stable inflation and high inequality measured by $\hat{\Sigma}_t$.

Changes in the steady state wages w impact the relative weights of the second-moments of the macroeconomics variables in the loss function. The coefficient on the variance of consumption inequality λ_{Σ} captures the sensitivity of the welfare loss due to the cross-sectional variability in consumption inequality. We can show analytically that $\partial\lambda_{\Sigma}/\partial w < 0$, indicating that the larger the steady state wage needed to achieve zero inflation, the lower the weight on the welfare cost of inequality. This coefficient further represents the desire of the social planner to reduce consumption inequality in the economy. The term $\lambda_{\tilde{y}}$ on the variance of output gap \tilde{y} is related to the disutility of supplying labour. Greater variance of output gap, $Var(\tilde{y}_t)$, increases the

dis-utility of labour. Similarly, a higher mean of the distribution of the uninsurable idiosyncratic shocks, $\bar{\xi}$, or a lower Frish elasticity of labour supply, ρ , would both contribute to a larger dis-utility of labour.

More importantly, our approximation showcases a novel trade-off between low and stable inflation and the variability of the welfare cost of consumption inequality encoded by the second-moments involving $\hat{\Sigma}_t$.¹⁰ When the cyclical costs of inequality, $Var(\hat{\Sigma}_t)$, are high, the monetary authority can sacrifice price stability in order to lower inequality. This mechanism is absent in the RANK model where household heterogeneity is absent. Acharya et al. (2021) show how this new trade-off affects optimal monetary policy.

3.2 Consumption inequality

In this section, we derive a measure of consumption inequality by aggregating individual consumption variance. Using the consumption policy function $c_t(x_{s,t}^i) = c_t + \mu_t x_{s,t}^i$ together with the evolution of cash-on-hand, we show that individual consumption volatility is a random walk where the innovation is the income risk. More details of the derivation can be found in Appendix C.

At the aggregate level, the random death and birth of cohorts stabilizes the model and generates a stationary consumption inequality process given by:

$$\sigma_{c,t}^2 = \mu_t^2 w_t^2 \sigma_t^2 + \vartheta \sigma_{c,t-1}^2. \quad (22)$$

The log-linearized dynamics of consumption inequality is summarized in the following proposition.

Proposition 2 *Aggregate consumption inequality evolves according to:*

$$\hat{\sigma}_{c,t} = (1 - \vartheta)(\hat{\mu}_t + \phi y \hat{y}_t) + \vartheta \hat{\sigma}_{c,t-1} \quad (23)$$

where the steady state consumption inequality is:

$$\sigma_c^2 = \frac{\mu^2 w^2 \sigma^2}{1 - \vartheta} = \frac{\Lambda}{\gamma^2(1 - \vartheta)} \quad (24)$$

Proposition 2 not only provides an equation describing how consumption inequality evolves in the model, but also highlights the channels through which monetary policy affects consumption inequality. Specifically, consumption inequality at any given point in time depends on the existing inequality accumulated over the years and on the income risk faced by the new cohort entering the economy. As illustrated earlier, through the self-insurance channel, encoded in $\hat{\mu}_t$ in equation (23), an increase in expected real rates or a decrease in expected real wages increases the income risk pass-through and consequently increases consumption inequality. Through the income risk channel, encoded by $\phi y \hat{y}_t$ in equation (23), a positive output gap decreases inequality in a model with countercyclical income risk.

¹⁰In absence of technological shocks, the last coefficient $\lambda_{z,\Sigma} = 0$.

It is useful to further emphasize that monetary policy affects consumption inequality in this model by precisely enabling well-functioning labour market and credit markets. For instance, following a negative income shock, households can consume less or work more. Thus, changes in the policy rate affects households by influencing the optimal choices in these two dimensions.

Finally, the expression of consumption inequality, (23), shows that $\sigma_{c,t}^2$ is a proxy for the unobservable welfare cost of inequality, $\hat{\Sigma}_t$ in equation (19). These two equations only differ in how they weigh the innovation and the autoregressive components. A comparison of their coefficient of autoregression shows that $\frac{\tilde{\beta}}{\beta} = \frac{\vartheta}{\beta R} > \vartheta$, so that the welfare cost of inequality process is marginally more persistent than the consumption inequality process.

4 Calibration

We calibrate the model at a quarterly frequency. We set the coefficient of absolute risk aversion parameter, γ , to be 1, and ρ to be $1/5$.¹¹ We set the elasticity of substitution between varieties, ε , to 9, implying a 12.5 percent markup in steady-state. The slope of the Phillips curve, κ , is calibrated to be 0.023, somewhat lower than textbook calibrations. A relatively flat Phillips curve is consistent with estimated models using Canadian data (Corrigan et al., 2021; Gervais and Gosselin 2014) as well as the observed lack of disinflation in recent U.S. and Canadian experiences during their respective ELB episodes. This implies a value for the Rotemberg adjustment cost parameter $\Phi = 2314$. We set the mean of the distribution of idiosyncratic risk, $\bar{\xi}$, to $1 + \gamma\rho$ to normalize the efficient level of output in a steady state to 1.

Since the global financial crisis, estimates of the neutral rate have shifted downward, and the current Canadian nominal neutral rate is in the range of 1.75 to 2.75 percent, with a midpoint of 2.25 percent (Brouillette et al., 2021). We calibrate the discount factor β to generate an annualized natural rate of interest of 2.25 percent, following the latest assessment of neutral rate for Canada. For the HANK case, the steady-state relationship, $R = \beta^{-1}e^{-\frac{\Lambda}{2}}$, results in a value of $\beta = 0.985$, given the parameter values chosen below. We set the effective lower bound for monetary policy to 25 bps.

The persistence parameters for the preference and cost-push shocks are set to 0.8. Standard deviations for the shocks are calibrated to match the general property that two-thirds of movements in the Canadian output gap in the RANK case are explained by demand shocks, while the remaining one-third are explained by supply (cost-push) shocks. This breakdown between supply and demand shocks is consistent with the estimated results from Bank of Canada's main DSGE model, ToTEM III (Corrigan

¹¹Under the efficient steady-state, $c = y = n = w = 1$, which implies a coefficient of relative risk aversion, $-\frac{cU_{cc}}{U_c} = c\gamma$ of 1, and the Frisch elasticity of labour supply, $\frac{U_l}{lU_{ll} - l\frac{U_{lc}^2}{U_{cc}}} = l\rho$ of $1/5$, as in Galí (2015).

et al., 2021). With this calibration, along with the choice of neutral rate at 2.25 percent, the nested RANK model reaches the effective lower bound at a frequency of around 14 percent.

Following Acharya et al. (2021), we set the survival probability, ϑ , to be 0.85. We calibrate the standard deviation of the distribution of idiosyncratic risk, σ , using the cross-sectional variance of income data from Canadian tax filers from the Longitudinal Administrative Database (LAD), a panel comprising a 20 percent sample of annual tax filings between 1982 and 2016.¹² This approach is also in line with Kaplan et al. (2018), that calibrated earnings processes using moments from U.S. tax filer data as reported by Guvenen, Karahan, Ozkan and Song (2015).¹³ The variance of log-earnings in the LAD database fluctuate between 0.61 to 0.80 over our sample period of A to B. We choose a value of $\sigma = 0.74$, which corresponds to a cross-sectional variance of wage income of 0.76 in the steady-state.

The constant semi-elasticity of the variance of cash-on-hand, $\phi = \frac{\partial \ln V(x)}{\partial y}$, can be modelled to capture procyclical ($\phi > 0$), countercyclical ($\phi < 0$), or acyclical ($\phi = 0$) earnings risk. Empirical evidence suggests that earnings risk is strongly countercyclical in Canada (Bowlus, Gouin-Bonenfant, Liu, Lochner and Park 2020, Brzozowski, Gervais, Klein and Suzuki 2010) as well as in the U.S., whether measured by the variance of idiosyncratic risk (Storesletten, Telmer and Yaron 2004, Nakajima and Smirnyagin 2019), or by the skewness (Guvenen et al. 2015). We perform estimation and provide evidence that the cross-sectional variance in earnings risk in Canada is countercyclical in both administrative tax-filer data (LAD) as well as survey-based panel data (SLID). In particular, we obtain estimates for the semi-elasticity of the variance in earnings with respect to output within ranges from -0.47 to -2.62.¹⁴ More details on the empirical estimation can be found in Appendix B. In particular, we calibrate the cyclicity of earnings risk, ϕ to be -2.62, the highest in this range.¹⁵

The transmission mechanism in HANK models depends crucially on the calibration of the income risk parameters σ and ϕ . Appendix E therefore explores the

¹²Available HANK models (Acharya and Dogra (2020), Acharya et al. (2021)) as well as the empirical literature on income heterogeneity provide distributional parameters at annual frequencies. In calibrating the quarterly versions of these parameters, we invoke the property that the variance of an average of i.i.d. normal distributions is the average of the variances. Under the assumption of i.i.d. normal idiosyncratic shocks, it follows that the annual variance of idiosyncratic earnings is simply the average of the quarterly variance of annualized idiosyncratic earnings, and that the magnitudes of volatility across the two frequencies are comparable. This means that empirical moments using annual data are still valid in informing our quarterly model. Note that this convenient property is no longer valid when idiosyncratic shocks are persistent.

¹³Following Guvenen et al. (2015), we only consider earnings data for working age males with earnings above minimum wage in the previous year to ensure sufficient ties to the labour market. Karibzhanov (2020) show that key moments from the LAD database are comparable to those obtained from U.S. tax filer data.

¹⁴We only consider coefficients that are statistically significant at least at the 10 percent level. Appendix B also infers values of ϕ based on job-loss dynamics, assuming that discrete probabilities of job losses are the salient risk to earnings. This gives us values of ϕ between -0.04 and -7.4.

¹⁵We have performed sensitivity analysis around this calibration and found that the qualitative conclusion remains robust to values within this range.

Table 1: Key Structural Parameters

Param	Description	Source	Value
γ	Absolute risk aversion	Gali (2015)	1
ρ	Frisch labour elasticity	Gali (2015)	1/5
ε	Elasticity of substitution	Gali (2015)	9
κ	Slope of Philips' curve		0.023
ϑ	Survival probability	ACD(2020)	0.85
σ	Idiosyncratic labour risk	LAD data	0.74
ϕ	Cyclicalty of risk	LAD and SLID data	-2.62
β	Discount factor	Target 2.25% neutral rate	0.985
R	Real risk-free rate		1.0056
$\tilde{\beta}$	Steady-state price of bond	$\tilde{\beta} = \frac{\vartheta}{R}$	0.845
w	Steady state wage		1.62
μ	Income risk pass-through to consumption in steady-state	$\mu = \frac{1-\tilde{\beta}}{1+\gamma\rho w}$	0.116
Λ	Steady-state consumption risk	$\Lambda = (\gamma\mu w\sigma)^2$	0.019
$\bar{\xi}$	Mean risk in steady-sate	$\bar{\xi} = 1 + \gamma\rho$	1.2
ρ_ν	Demand shock persistence		0.8
ρ_ε	Cost-push shock persistence		0.8
σ_ν	Demand shock standard deviation		0.0169
σ_ε	Cost-push shock standard deviation		0.0002696

transmission of shocks under alternative calibrations of these two parameters. In particular, we show that higher absolute values of σ and ϕ result in higher degrees of amplification of shocks in the HANK model. However, adopting these alternative values does not alter the main message of this paper.

5 Model dynamics under alternative policy frameworks

The main objective of this paper is to examine the performance of alternative monetary policy frameworks with varying degrees of history dependence in a HANK model. Before we proceed to the results presented in Section 6, it is instructive to examine the key dynamics in the HANK model under alternative monetary policy frameworks. In this section, we first define the frameworks we implement via policy rules, and then examine the dynamics of shock transmission under two illustrative cases of inflation targeting and price-level targeting.

5.1 Monetary Policy Rules

We implement alternative monetary policy frameworks using different classes of policy rules described below. This approach of representing policy frameworks using policy rules is consistent with the recent analysis in Amano et al. (2020) and Feiveson et al. (2020).

The alternative policy frameworks differ in the degree of history dependence applied to the nominal variable targeted by the central bank. Consider a general formulation of a nominal target variable, P_t^* :

$$P_t^* = \sum_{j=0}^N (\pi_{t-j} - \bar{\pi})$$

where $\pi_t = p_t - p_{t-1}$ is the quarterly inflation rate, p_t is the log price level, and $\bar{\pi}$ is the target inflation rate. As N increases, the central bank responds to a larger window of inflationary deviations in history compared to the target $\bar{\pi}$. With a large window of N , the Central Bank will attempt to react to, or ‘make up for’ misses in inflation π_t from its target $\bar{\pi}$ that occurred farther in the past.

In this general setup, the inflation targeting framework corresponds to a case where $N = 4$, making year-over-year inflation the target for the central bank. When $N \rightarrow \infty$, P_t^* converges to the price level, since the latter is the accumulation of all past inflation. Intermediate cases of $\infty > N > 4$ represent average inflation targeting frameworks, where the central bank reacts to a defined window of past inflation deviations and treats the remaining deviations as bygones.

We implement the different frameworks using a simple policy rule of the following form:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) [i^* + \theta_\pi P_t^* + \theta_y \tilde{y}_t] \quad (25)$$

where i_t is the nominal policy interest rate, i^* is the nominal neutral rate of interest, and \tilde{y}_t is the output gap.

Two issues regarding the parameters of equation (25) need to be addressed before moving on. First, we consider values for the coefficient on the output gap, θ_y , such that $\theta_y \geq 0$. Note that $\theta_y = 0$ corresponds to a case of strict targeting framework, while $\theta_y > 0$ corresponds to flexible targeting. Inflation targeting practised in Canada and elsewhere is flexible inflation targeting, as central banks usually respond to the real side of the economy, regardless of the degree of explicit focus on the real side in their communicated framework. In what follows, we do not restrict ourselves to either fixed or flexible targeting, but rather consider both possibilities.

Second, we set the smoothing parameter, $\rho_i = 0.85$. This value is broadly in line with estimates of simple monetary policy rules for Canada (c.f. Corrigan et al., 2021).¹⁶ We keep this parameter value fixed to allow a cleaner comparison across different frameworks.

For concreteness, we consider four cases:

Flexible Inflation Targeting rule (IT).

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) [i^* + \theta_\pi (\pi_{t+4}^{yy} - \bar{\pi}) + \theta_y \tilde{y}_t] \quad (26)$$

¹⁶Woodford (2003) argue that delegating the central bank a loss function that includes interest-rate smoothing can be welfare improving when the central bank operates under discretion.

where π_{t+4}^{yy} is the four-quarter ahead year-over-year inflation rate.

Average Inflation Targeting rule (AIT).

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) \left[i^* + \theta_\pi (\pi_{t+4}^{3y} - \bar{\pi}) + \theta_y \tilde{y}_t \right] \quad (27)$$

where (π_{t+4}^{3y}) stands for the four-quarter ahead three-year average inflation rate. Since this rule effectively includes two years of history, any deviations to inflation that occurred more than two years ago are treated as bygones.

Price-Level Targeting rule (PLT).

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) \left[i^* + \theta_\pi (p_{t+4} - \bar{p}_t) + \theta_y \tilde{y}_t \right] \quad (28)$$

where p_{t+4} denotes the four-quarter ahead log price level and \bar{p}_t denotes the log of the target price path such that

$$\begin{aligned} p_t &= p_{t-1} + \pi_t \\ \bar{p}_t &= \bar{p}_{t-1} + \bar{\pi} \end{aligned}$$

where $\bar{\pi}$ is the target rate of inflation.

Nominal GDP Level Targeting rule (NGDPL).

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) \left[i^* + \theta_y \{ (p_{t+4} + y_{t+4}) - (\bar{p}_t + \bar{y}_t) \} \right] \quad (29)$$

where $p_{t+4} + y_{t+4}$ denotes the four-quarter ahead log nominal GDP level and $\bar{p}_t + \bar{y}_t$ denotes the log of the target nominal GDP level.

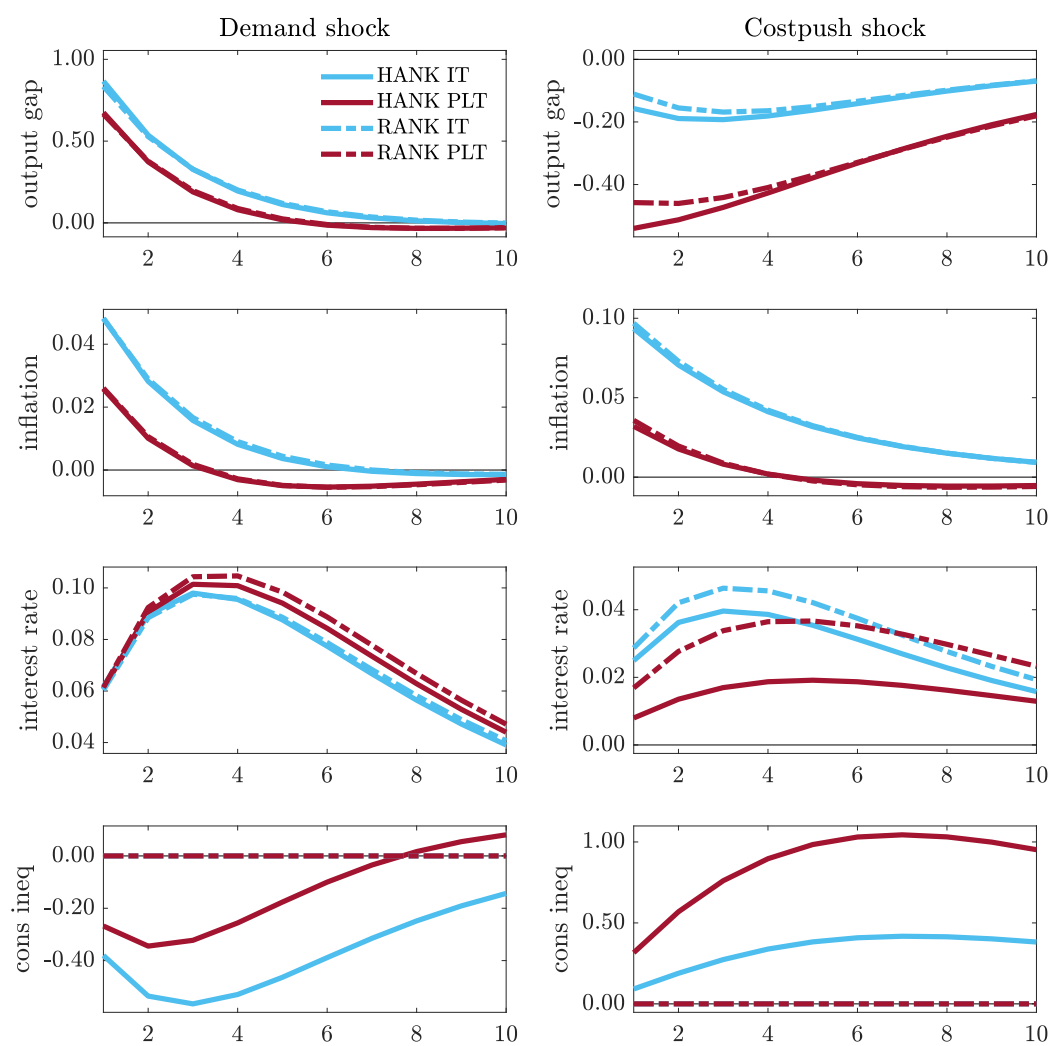
Finally, we perform our analysis where the HANK model is log-linearized around a zero-inflation steady-state. In other words, we set $\bar{\pi} = 0$.

5.2 Transmission of shocks under IT and PLT

At each period of time, our model economy is hit by two aggregate shocks—a demand shock, modelled as an exogenous shock to preferences, and a cost-push shock. In this subsection, we explain how uninsurable risk in the HANK model amplifies business cycles and generates consumption inequality in response to these two shocks. We also show how these responses change across the IT and PLT frameworks.

The dynamics of the real economy and consumption inequality are shown in Figure 1 using impulse responses for demand and cost-push shocks for the HANK and embedded RANK models. In this exercise, we assume that the monetary authority follows either of two policy rules. The inflation targeting case, represented in blue, is implemented with the rule $i_t = 2.5\pi_t$. The price level targeting case, represented in red, is implemented with a similar rule $i_t = 2.5p_t$. For ease of exposure, we consider only strict targeting frameworks in this example. The HANK model is represented in solid lines, and the embedded RANK case is represented in dashed lines.

Figure 1: Impulse responses to demand and cost-push shocks across different policy frameworks



The first column in Figure 1 shows impulse responses for a persistent positive demand shock. Consider first the IT case. A positive demand shock generates a boom in output and a corresponding rise in inflation. In the HANK case, a boom in output directly lowers the cross-sectional dispersion in income. As a result, precautionary savings fall through the *income risk* channel, which amplifies the initial boom in output in the HANK model vis-à-vis the RANK case. At the same time, the decline in cross-sectional income dispersion is translated into a reduction in consumption inequality.

Note, however, that the interest rate hike makes it harder for agents to partially self-insure themselves against their idiosyncratic risk using aggregate bonds. Since it becomes harder to hedge against risk, precautionary savings rise through the *self-insurance* channel, offsetting some of the amplification effects described above.

In a PLT framework, the policy rule stipulates that the initial rise in inflation due to the positive demand shock will be followed by a period of negative inflation to bring the price level back to target. Anticipating the period of negative inflation, forward-looking price setters increase their prices by less than the IT case. As a result, the rise in inflation (and output) is less pronounced under PLT than under IT. At the same time, a smaller output response translates to a smaller decline in consumption inequality under PLT.

The second column in Figure 1 shows responses for a persistent cost-push shock, modelled as a shock to the NKPC. Central banks interested in price stability respond to this higher inflation by raising rates and generating a recession to bring inflation back to target.

Again, consider the IT case first. The fall in output leads to higher *income risk* and hence more precautionary savings, which *amplifies* the recession. At the same time, the rise in rates result in higher precautionary savings through the *self-insurance* channel, further amplifying the downturn. In other words, incomplete markets generate a larger downturn in the HANK economy vis-à-vis the RANK case. This downturn leads to an increase in consumption inequality. Importantly, the central bank now faces an additional trade-off. In order to bring inflation back to target following a cost-push shock, central banks now not only have to generate a recession, but also have to accept higher inequality as a consequence of that recession.

In a PLT framework, the central bank has to generate a period of negative inflation following the initial response to the inflationary cost-push shock to bring price levels back to target. This is achieved by generating a larger recession than in the IT case. In the HANK model, the recession is again amplified through the channels described above. A larger downturn under PLT also leads to a bigger increase in consumption inequality.

6 Ranking of monetary policy frameworks

In this section, we compare the performance of alternative history-dependent monetary policy frameworks on aggregate outcomes in our calibrated HANK model. Technical details of the model-based evaluation of macroeconomic performance are broadly in line with those described in Sections 3 and 4 in Dorich, Mendes, and Zhang (2021). See also Swarbrick and Zhang (forthcoming). We describe the simulation setup, then present the main findings under two cases: (i) when the ELB does not pose a binding constraint for monetary policy, and (ii) when it does. The former can be thought of as a situation where the unconventional policy tools available to the central banker can perfectly compensate for the ELB, and the latter, a situation where these tools are completely ineffective (c.f. the discussion in Dorich et al. (2020)). These extreme cases therefore provide an upper and lower bound of possibilities facing the central bank.

In each of these cases, we first present our main finding based on the social welfare-based loss function, which has arguments on inflation, output, and the welfare cost of inequality with precise weights on each from the deep parameters calibrated in Section 4. We then consider a number of simple ad hoc loss functions, which help illustrate our main findings when accounting for uncertainty in central bank preferences. Finally, we discuss outcomes when the central bank is delegated a framework-specific loss function.

6.1 Simulation design

We evaluate the performance of alternative policy frameworks through the lens of loss functions considered by the central bank. Once a loss function is defined, the remaining exercise consists of finding the combination of θ_π and θ_y parameters in the policy rule that yields the minimum loss across different policy frameworks implemented via policy rules described in section 5.1.

We perform the analysis using stochastic simulations under each regime, allowing different coefficient values in respective policy rules. Each framework in the horse race is assumed also to incorporate an equivalent element of interest rate smoothing. In this sense, all frameworks include some history dependence. In a policy rule with interest rate smoothing, current policy interest rates are a function of the lagged policy interest rate (past conditions), implying a degree of history dependence. In the horse race in the HANK model, all policy rules are subject to the same degree of gradualism in adjustment to the policy interest rate. The simulation of each rule uses random errors that embody the estimated historical distribution of shocks in the model. The simulated moments of relevant variables therefore depend on the shock structure described in Section 4. We then evaluate the performance of these policy frameworks across a broad range of criteria (such as unconditional mean and volatility of the variables, as well as the regime's performance during the ELB).

Table 2: Welfare loss in RANK and HANK across different policy frameworks

Frameworks	RANK	HANK	(HANK - RANK)
Inflation Targeting	0.088	0.160	0.072
Average Inflation Targeting (3y)	0.069	0.123	0.054
Price Level Targeting	0.065	0.115	0.050
Nominal-GDP level targeting	0.176	0.300	0.124

6.2 Simulation results without the ELB

6.2.1 Results with a welfare-based loss function

We first evaluate the performance of alternative frameworks using the social welfare-based loss function derived in proposition 1. Normalizing the weight on inflation in the expected loss function to one, we get the following reduced-form expression, given parameter and steady-state values provided in table 1.

$$L_t^{SWF} = Var(\pi_t) + 0.003Var(\tilde{y}_t) + 0.0003Var(\hat{\Sigma}_t) + 0.0007E(\hat{\Sigma}_t) + 0.0002Var(\hat{y}_t) \\ + 0.0003E(\hat{y}_t) + 0.0003Cov(\hat{y}_t, \hat{\Sigma}_t) + 0.0003Cov(\hat{y}_t, \hat{v}_t) + 0.0006Cov(\hat{\Sigma}_t, \hat{v}_t)$$

Two important intuitions are worth highlighting. First, inequality generates a loss in expected social welfare in the HANK model. Second, while being relevant, the weight of inequality on welfare loss is small, and inflation remains the main driving force for welfare. This suggests that, under the benchmark calibration, welfare loss generated by nominal rigidities far outweighs the loss stemming from financial market incompleteness.

We find the minimum expected social-welfare loss by optimizing the policy rule parameters for the alternative frameworks described in section 5.1. These losses are expressed as a percentage of steady-state output and are presented in Table 2. The first column represents the losses from the reference RANK model, where sticky prices are the only distortion in the economy since markets are assumed to be complete so that idiosyncratic risk is perfectly insured against. The expected loss in the RANK model simply admits the first two arguments to the right-hand side (RHS) of equation (30). The second column considers the HANK case with inequality through the introduction of uninsurable idiosyncratic risk and market incompleteness. The third column measures the difference between the first two, which can be interpreted as the marginal loss introduced by market incompleteness. Table 3 provides further details for the HANK case by reporting standard-deviations of key variables for the alternative frameworks.

We obtain three key results under welfare-based comparison. First, history-dependent frameworks dominate for the welfare-based loss function. PLT produces the lowest

Table 3: Simulated moments in HANK across different policy frameworks

Frameworks	$std(\pi^{yy})$	$std(\tilde{y})$	$std(\Delta i)$	$std(\sigma^c)$	Welfare	Relative Loss (vs. IT)
Inflation Targeting	0.159	1.546	0.872	4.006	0.160	100%
Average Inflation Targeting (3y)	0.186	1.215	1.078	3.520	0.123	77%
Price Level Targeting	0.173	1.198	1.305	3.563	0.115	72%
Nominal-GDP Level Targeting	0.436	1.122	0.284	2.134	0.300	188%

welfare loss, followed by AIT and IT, while NGDP level targeting produces the highest welfare loss. Looking at decomposed components of the welfare through Table 3, we see that none of the frameworks strictly dominates in stabilizing all key macroeconomic variables. When augmented with smoothing, IT yields competitive performance as PLT in stabilizing inflation. In contrast, AIT yields comparable output gap stabilization as PLT without generating excessive interest rate volatilities. History-dependent rules also demonstrate superior performance in stabilizing consumption inequality.

Second, in our benchmark calibration, welfare losses due to nominal rigidities and incomplete financial markets are small—within 0.06 to 0.30 percent of steady-state output. Nonetheless, the HANK setup increases the importance of history dependence. In the RANK model, moving from IT to PLT improves social welfare by approximately 0.02 percent of steady-state output. Compared to that, moving from IT to PLT in the HANK model improves social welfare by approximately 0.04 percent of steady-state output.¹⁷

Lastly, welfare loss in the HANK model is always greater than that in the RANK model. The uninsurable idiosyncratic risk and the resulting fluctuation in consumption inequality therefore accounts for welfare loss worth 0.05 to 0.12 percent of steady-state output in our model.

6.2.2 Results with ad hoc loss functions

The baseline social welfare-based loss considered above contains welfare costs from inflation, output, and inequality. In this subsection, we complement the baseline results by considering some uncertainty around the characterization of central bank loss functions. We use several simple ad hoc loss functions to illustrate the conditionality of relative performance of alternative frameworks to the specification of the loss function.

¹⁷Kryvtsov, Shukayev and Ueberfeldt (2008) find that moving from IT to PLT is equivalent to reducing the standard-deviation of inflation by 0.045 percentage points. For reference, the observed standard deviation in CPI inflation was 0.4 percentage point in the 1996 Q1–2007 Q2 periods. All else being equal, this is equivalent to a reduction in the variance of inflation of 0.034 percentage points, which translates to a welfare loss of 0.034 percentage points of steady-state output following proposition 1. In other words, our benchmark calibration provides a welfare gain of moving to PLT that is slightly higher than in Kryvtsov et al. (2008).

Simple loss functions allow us to be agnostic about society's values imposed on the central bank, while keeping the model transmission mechanisms intact. Rather than one particular loss function with a specific weight on its arguments, we can consider the performance of policy frameworks across a broad set of possible characterizations. The framework that dominates across a larger subset of loss functions is robust to the specification of the loss function.

Loss function with inflation and output gap. Consider first an ad hoc loss function of the form:

$$L_t = (\pi_t^{yy} - \bar{\pi}_t)^2 + \lambda_y (\tilde{y}_t)^2 \quad (30)$$

where π_t^{yy} represents annualized year-over-year inflation, \tilde{y}_t denotes the output gap, and λ_y denotes the relative weight on the output gap. Loss functions of this form have been widely used as a benchmark in comparing outcomes under alternative models (Bernanke et al. (2019), Dorich et al. (2020)).¹⁸


Instead of restricting λ_y to its welfare-consistent value, we take an agnostic approach and consider a range of values of $\lambda_y \in [0, 1]$, following Vestin (2006). Conditional on a chosen value of λ_y , we optimize policy coefficients θ_π and θ_y for each policy framework to find the set of coefficients that minimize the average of the loss given in equation (30). The performance of each framework is then represented by the unconditional mean squared errors of the output gap and annualized year-over-year inflation deviations from target. We repeat the exercise for each value of λ_y to form a loss plot.

The results are presented in Figure 2 through an inflation-output gap efficient frontier for each framework. Since both demand and supply shocks contribute to business cycles of the HANK economy, an efficient frontier of the policy rule implicitly accounts for the necessary adjustment of the nominal interest rate in response to deviations of inflation from target. In this simple framework, the trade-off between inflation and output gap variability depends largely on the choice of the relative inflation weight in the policy reaction function. As the weight on inflation increases, the nominal cash rate becomes more variable and required policy changes become larger, thereby increasing interest rate variability. When there is no ELB, PLT yields the best performance in price stability without generating excessive real fluctuation. The efficient frontier of PLT situates closest to the origin, followed by that of AIT and IT.

To understand this result, consider first the findings in Nessen and Vestin (2005) and Vestin (2006) in a RANK setup. Policy actions under PLT or AIT are set to undo the effects of past shocks on the price-level, which involves correcting past inflation deviations. History-dependent policies lower the volatility of inflation through generating higher inflation expectations.

¹⁸Similar loss functions can be derived from micro-founded welfare in a simple RANK economy, where the deep parameters underlying the model map into a particular values for λ_y (see Galí (2015) and Woodford (2003) for text-book derivations).

Figure 2: Efficient tradeoff frontier across alternative policy frameworks



An AIT framework engenders history dependence only partially by restricting policy action to undo the effects of a truncated window of past inflation (in our case, for three years). Correspondingly, forward-looking price setting agents take this future action into account and through the expectation channel, aggregate volatility in inflation and output fall somewhere between IT and PLT.¹⁹

The important finding for our paper is that history-dependent frameworks continue to outperform in the HANK model. Recall that household heterogeneity and uninsurable risk amplify monetary policy and macroeconomic shock transmission to the real economy only through the modified Euler equation (first equation in the system (14)). The relationship between inflation and output gap dictated by the NKPC in the HANK model is unaffected by these features. Any amplified response in the output gap in the HANK model is therefore transmitted one-to-one to inflation. As a result, the efficient trade-off between output gap and inflation volatility is unaffected by the introduction of household heterogeneity and incomplete financial markets. In other words, the additional transmission channels afforded through household heterogeneity does not alter the inflation-output trade-off facing a central bank.

In the ad hoc loss function involving inflation and output gap only, the central bank implicitly restricts its considerations to welfare distortions generated by nominal rigidities, and ignores welfare distortions generated by incomplete financial markets. Our result above shows that household heterogeneity does not provide any meaningful alteration to the trade-offs facing the central bank when we ignore the welfare costs of incomplete financial markets. We now turn to explicitly accounting for the latter as captured through consumption inequality.

Loss function with inflation, output gap, and inequality. Household heterogeneity and incomplete financial markets generate consumption inequality, which is undesirable from a social perspective. The cost of this market failure can be captured by adding consumption inequality to the simple ad hoc loss function introduced in equation (30) as follows:

$$L_t = (\pi_t^{yy} - \bar{\pi}_t)^2 + \lambda_y (\tilde{y}_t)^2 + \lambda_\sigma^c (\sigma_t^c)^2 \quad (31)$$

where the additional argument σ_t^c denotes consumption inequality across households, and λ_σ^c denotes the central bank's preference in attenuating the cyclical variation in inequality.


Figure 3 plots the expected loss for the IT, AIT, and PLT given this new ad hoc loss with inequality. For simplicity, we assume equal weight of inflation and output gap stabilization ($\lambda_y = 1$) and vary $\lambda_\sigma^c \in [0, 2]$ in making comparison.

When the weight on inequality is relatively low, PLT continues dominating AIT and IT in generating the lowest loss.²⁰ As the weight of consumption inequality

¹⁹Note, however, that this result is overturned when a sufficient portion of price setters departs from forward-looking behaviour. Nessen and Vestin (2005) show that in a model with both forward-looking and rule-of-thumb price setters, AIT can dominate PLT.

²⁰The left-most point in Figure 3, with $\lambda_\sigma^c = 0, \lambda_y = 1$, corresponds to the right-most point in Figure 2.

Figure 3: Ad hoc loss with consumption inequality across alternative policy frameworks



increases, the relative advantage of PLT starts to diminish. There exists an indifferent point where IT starts to outperform history-dependent frameworks once CB assigns more weight on inequality. At the upper bound value considered here (λ_σ^c at 2), IT dominates because it yields better stabilization in consumption inequality than other history-dependent rules.

To illustrate the mechanism behind this result, we report volatilities of key variables corresponding to two different sample points in Figure 3—one where $\lambda_\sigma^c = 0.1$ and PLT dominates IT; and another where $\lambda_\sigma^c = 1$ and IT dominates PLT. Table 4 collects these numbers.

In the HANK model, consumption inequality depends on (i) the level of output, which influences income risk directly through the *income risk channel*, and (ii) the pass-through of idiosyncratic income risk to consumption risk, μ_t , which allows consumers to partially self-insure against their idiosyncratic risk through the *self-insurance channel*. The first channel implies that, all else equal, policy frameworks with more volatile output gaps would also see a higher cost of consumption inequality. The second channel implies that, all else equal, policy frameworks with volatile real interest rates would result in unstable pass-through of income risk to consumption risk, which also increases the cost of consumption inequality. Since the central bank cannot stabilize the level of output and the income risk pass-through independently with only one instrument, there is a trade-off.

Table 4: Unconditional moments of the ad-hoc loss function: high vs low weight on inequality

Framework	$std(\pi^{yy})$	$std(\tilde{y})$	$std(\Delta i)$	$std(\sigma^c)$	Ad hoc Loss	Welfare Loss
$\lambda_\sigma^c = 0.1$						
IT	0.77	0.25	0.99	0.97	0.75	0.71
PLT	0.62	0.24	1.07	1.15	0.58	0.46
$\lambda_\sigma^c = 1$						
IT	0.83	0.42	0.61	0.53	1.15	0.83
PLT	0.70	0.27	0.79	0.83	1.25	0.59

Table 4 shows that under the loss function considered, PLT stabilizes both inflation and output better than IT, but at a cost of higher interest rate variability (vis-à-vis IT). The loss stemming from consumption inequality, however, is higher for PLT than in IT regardless of CB’s weight on inequality stabilization. This can be attributed to the unstable interest rate environment that leads to an unstable pass-through of income risk through the self-insurance channel.

The weight CB assigns to inequality in the loss does, however, affect the overall ranking of the regime. For instance, when the central bank puts a small weight on the cost of inequality ($\lambda_\sigma^c = 0.1$), PLT yields smaller loss than IT. In contrast, when the weight on the cost of inequality is larger ($\lambda_\sigma^c = 1$), the disadvantage of PLT in stabilizing inequality becomes more relevant and IT starts to outperform. Importantly, the

relative ranking of regime depends on the preference toward consumption inequality stabilization, relative to its ability in stabilizing price and real activities in a model with potential trade-off in these actions.

6.2.3 Results with framework-specific loss functions

So far our analysis has abstracted from discussing the role of interest rate volatility. In practice, central banks also conduct policy trying to avoid excessive policy rate movements. In this section, we characterize each framework using a regime-specific loss function delegated to the central bank, following Svensson (2020). The literature on delegated loss functions often takes as given that central banks can act only under discretion, and asks whether providing the central bank a loss function different from society's true loss function could improve performance when the central bank lacks a credible commitment device (c.f. Rogoff (1985)).

We explore this question by characterizing alternative frameworks using the delegated loss functions given in Table 5, which follows the specifications suggested in Svensson (2020) and Dorich et al. (2020). Each delegated loss specifies a different target for nominal values. For example, flexible inflation targeting requires that the central bank minimize the squared deviation in year-over-year annualized inflation π_t^{yy} from its target, while price-level targeting requires that the central bank minimize the squared deviation of the price level p_t from its targeted price-level path. Note that all delegated loss functions incorporate the consideration of stabilizing real activities by including a quadratic term on the output gap, except nominal GDP level targeting, which explicitly targets the nominal output level. Moreover, each loss function imposes inertia in policy-rate movements by imposing quadratic losses on interest-rate movements. The latter follows Woodford (2003)'s argument that imposing inertial policy rate movements can introduce elements of history-dependence in discretionary policy.

Table 5: Framework-specific delegated loss functions

Framework	Loss specification
Inflation targeting	$L^{IT} = (\pi_t^{yy} - \bar{\pi})^2 + (\tilde{y}_t)^2 + 0.5 (\Delta i_t)^2$
Average inflation targeting	$L^{AIT} = (\pi_t^{3y} - \bar{\pi})^2 + (\tilde{y}_t)^2 + 0.5 (\Delta i_t)^2$
Price-level targeting	$L^{PLT} = (p_t - \bar{p})^2 + (\tilde{y}_t)^2 + 0.5 (\Delta i_t)^2$
Nominal-GDP level targeting	$L^{NGDPL} = [(p_t + y_t) - (\bar{p}_t + \bar{y}_t)]^2 + 0.5 (\Delta i_t)^2$

We evaluate the performance of each framework based on unconditional moments, as well as the welfare-based loss function. These results are reported in Table 6. As in the welfare loss case, there is no single framework that outperforms others across all the variables considered. For example, PLT is the framework that best stabilizes inflation, which, however, comes at a price of higher volatility in output gap and

consumption inequality. AIT best stabilizes the output gap, while IT best stabilizes consumption inequality.

Table 6: Moments in HANK across different policy frameworks (Regime-specific loss)

Frameworks	$std(\pi^{yy})$	$std(\tilde{y})$	$std(\Delta i)$	$std(\sigma^c)$	Welfare loss	Relative loss (vs IT)
IT	0.731	0.461	0.648	0.900	0.64	100%
AIT (3y)	0.678	0.438	0.644	0.999	0.56	86%
PLT	0.388	0.708	0.634	2.084	0.21	33%
NGDPL	0.436	1.122	0.284	2.134	0.30	47%

6.3 Horse race of frameworks with ELB

The analysis so far has focused on linearized solutions of the tractable HANK model. We now examine how the presence of the ELB changes the relative performance of policy frameworks. Simulations are performed assuming an annualized neutral rate of 2.25 percent and an effective lower bound of 25 bps. The simulations embed an important underlying assumption that the central bank’s unconventional policy tools are completely ineffective in overcoming the constraints of the ELB.

6.3.1 Welfare-based loss function

Table 7 reports unconditional moments of key variables and the welfare-based loss (as a percentage of steady-state output) in the HANK model with an occasional binding ELB. Similar to the case without the ELB, none of the frameworks strictly dominates the others. PLT generates the unconditional mean of inflation nearly at target, with the lowest inflation volatility and low skewness. In contrast, there exists disinflationary bias under IT, and the distribution of inflation is negatively skewed.

The last two columns report the welfare loss (as a percentage of steady-state output), and the relative loss (vis-à-vis IT) across different frameworks. We see again that PLT dominates the ranking, followed by AIT, IT, and finally NGDPL. Comparing these results with Table 3, we see that the ELB increases welfare losses by 0.01 to 0.02 percent of steady-state output across all frameworks except nominal GDP targeting, which has a very low frequency of hitting the ELB. While the ranking of policy frameworks remain the same as the linear case, the benefits of PLT and AIT are larger under the ELB. In particular, moving from IT to PLT reduces welfare losses by 0.06 percent of steady-state output under ELB, compared to 0.04 percent if the ELB was not a binding constraint.

Table 7: Moments in HANK across different policy frameworks

Frameworks	$E(\pi^{yy})$	$std(\pi^{yy})$	$sk(\pi^{yy})$	$E(\tilde{y})$	$std(\tilde{y})$	$sk(\tilde{y})$	$std(\Delta i)$	$std(\sigma^c)$	Freq	Dur	Welfare Loss	Relative Loss (vs IT)
IT	-0.02	0.185	-1.65	-0.06	1.565	-0.017	0.653	3.784	11.2	3	0.182	100%
AIT (3y)	-0.01	0.184	-0.20	-0.04	1.258	-0.026	0.915	3.396	16.5	3	0.133	73%
PLT	0.00	0.172	-0.06	-0.02	1.240	-0.037	0.760	3.420	12.4	3	0.124	68%
NGDPL	0.00	0.435	-0.01	0.02	1.124	-0.063	0.281	2.105	1.3	2.4	0.300	165%

6.3.2 Ad hoc loss functions

We first present the efficient frontier in Figure 4. The solid lines represent the efficient frontier without the ELB, while the dotted lines represent efficient frontiers with the ELB. The presence of the ELB increases the volatility of both inflation and the output gap, shifting the efficient frontiers away from the origin. The relative positioning of the curves remain the same, with history-dependent policy frameworks outperforming even under the ELB.

Figure 5 plots the expected loss for IT, AIT, and PLT using the loss function specified in equation (31). Alongside inflation and the output gap, the variance of consumption inequality in the loss function varies with the specified weight $\lambda_\sigma^c \in [0, 2]$.

The crossing over of loss functions now occurs at a higher value of λ_σ^c under the ELB. When the weight on the volatility of consumption inequality is small, we see that PLT produces the minimum loss, followed by AIT and then IT.²¹ However, as we increase the weight on inequality, very soon the loss under PLT becomes higher than under IT or AIT.

6.3.3 Framework-specific loss functions

Table 8 presents results under the ELB where the central bank has been delegated the framework-specific loss functions described in Section 6.2.3. As before, there is no single framework that dominates others in terms of stabilizing all variables. When CB assigns equal weights to inflation and output gap stabilization while accounting for interest rate volatility, IT stabilizes consumption inequality best, but involves more volatile inflation and interest rate movements. PLT comes the closest to matching the welfare-based loss from Section 6.3.1. As shown in Table 9, the main advantage of PLT lies in its ability to stabilize both inflation and the output during the periods of binding ELB. However, this benefit of PLT comes at the cost of elevated volatility in output gap and, importantly in our HANK model, increased consumption inequality.

²¹As before, the left-most points in Figure 5, with $\lambda_\sigma^c = 0, \lambda_y = 1$, corresponds to the right-most point in Figure 4. In this case, both the loss function in Figure 5 and the efficient frontier in Figure 4 suggest that PLT is preferred to AIT, which in turn is preferred to IT.

Figure 4: Efficient frontier: with and without the ELB




Figure 5: Ad hoc loss with varying weight on consumption inequality under the ELB




Table 8: Moments in HANK across different policy frameworks (Regime-specific loss)

Frameworks	$E(\pi^{yy})$	$std(\pi^{yy})$	$E(\tilde{y})$	$std(\tilde{y})$	$std(\Delta i)$	$std(\sigma^c)$	Freq (ELB)	Dur (ELB)	Welfare Loss	Relative Loss (vs IT)
IT	-0.03	0.720	-0.07	0.590	0.630	0.910	10.1	2.9	0.65	100%
AIT (3y)	-0.02	0.650	-0.05	0.560	0.610	1.000	9.5	2.9	0.53	82%
PLT	0.00	0.390	-0.02	0.770	0.590	2.040	8.6	2.8	0.22	33%
NGDPL	0.00	0.435	0.02	1.124	0.281	2.105	1.3	2.4	0.30	46%

Table 9: Conditional moments in HANK across different policy frameworks (Regime-specific loss, conditional on IT)

Frameworks	$E(\pi^{yy})$	$std(\pi^{yy})$	$E(\tilde{y})$	$std(\tilde{y})$	$std(\Delta i)$	$std(\sigma^c)$	Dur (ELB)
IT	-0.620	0.720	-0.920	0.900	0.350	1.240	2.9
AIT (3y)	-0.520	0.610	-0.750	0.870	0.360	1.210	2.5
PLT	-0.230	0.370	-0.400	1.070	0.380	2.240	2.1
NGDPL	-0.307	0.401	-1.000	1.037	0.241	2.118	0.4

7 Conclusion

In this paper, we analyze the implications of inflation targeting (IT) and several alternative history-dependent monetary policy frameworks (AIT, PLT, and NGDPL) using the tractable HANK model of Acharya et al. (2021) where agents face idiosyncratic earnings risk and financial markets are incomplete. The failure in financial markets generate two distinct features in the model. First, uninsurable risk leads to precautionary savings, which amplify business cycle fluctuations. Second, incomplete markets lead to cyclical variations in consumption inequality across households, which is undesirable from a societal perspective. The social welfare-based loss function in this model captures welfare distortions from missing financial markets, in addition to the usual New Keynesian distortions related to nominal price rigidities.

We find that history-dependent policies are preferred in the HANK economy when the central bank considers the social welfare-based loss function. Although household heterogeneity results in amplification of shocks in this economy, it does not affect the pricing decision of firms. This keeps the NKPC and the inflation-output trade-off faced by the central bank unaltered compared to the RANK case. In this setting, history-dependent policies result in lower inflation volatility through the expectation channel. The benchmark calibration implies that the welfare cost of missing financial markets is low compared to the cost of inflation. Therefore, the better inflation stabilization under AIT and PLT leads to lower societal loss compared to loss under IT. A binding ELB constraint increases the importance of history-dependent policies vis-à-vis IT.

Our analysis contributes to the literature in providing quantitative evidence around the uncertainty of monetary policy reaction function, as well as central bank preferences. If the central bank were to consider an ad hoc loss function that includes a sufficiently large weight on consumption inequality, IT can become more favourable than history-dependent frameworks. AIT outperforms PLT in that PLT involves additional volatility in interest rates in making up all past inflation misses. This diminishes agents' capacity to partially self-insure against idiosyncratic income shocks, and increases the pass-through of income risk to consumption risk. As a result, all else being equal, PLT generates larger consumption inequality than IT or AIT.

In interpreting the results of this paper, a few caveats should be kept in mind. First, our HANK setup is different from the setup in Kaplan et al. (2018) in that agents in this model do not differ in their marginal propensities to consume (MPC). In other words, we do not have a setup where the proportion of hand-to-mouth households vary across the business cycle. However, in so far as MPC heterogeneity implies cyclical amplification but does not affect firms' pricing decisions, the overall quantitative results in this model remain applicable even when considering the endogenous MPC feature. On the other hand, the relative performance and ranking of alternative policy frameworks would differ if we were to consider rule-of-thumb behaviour in price and/or wage setting, or if the NKPC departs from rational expectation.

We also abstract from fiscal instruments targeted at reducing inequality in the economy. Le Grand et al. (2020) suggest that when fiscal instruments are available to tackle inequality, the redistributive role for monetary policy is limited. In that case, the trade-off faced by the central bank is limited between stabilizing inflation and output, whereby history-dependent policies tend to outperform.

Our framework highlights the importance of imperfections in the financial markets, stemming from the limited ability of households to insure idiosyncratic income risk. In addition, we abstract from redistribution channels offered by inflation when debt contracts are nominal, as in Auclert (2019). There remains very limited scope for analyzing the financial stability implications of a monetary policy framework in our model. For policies that provide additional accommodation to mitigate inequality, we are limited to relating frequency and length of periods with very low interest rates with potential building of risk-taking behaviours and financial vulnerability. In future research, it would be useful to account for alternative financial frictions to study the interacting implications of household heterogeneity and financial stability considerations in assessing relative performance of alternative monetary policy frameworks.

Another important avenue to explore would be to scrutinize how central bank communication interacts with distributional concerns. In fact, recent survey evidence suggests that modelling how market participants understand and use economic information in their decision-making might affect the efficacy of monetary policy frameworks. We leave this for future research.

References

- Acharya, Sushant and Keshav Dogra**, “Understanding HANK: Insights from a PRANK,” *Econometrica*, 2020, 88 (3), 1113–1158.
- , **Edouard Challe, and Keshav Dogra**, “Optimal monetary policy according to HANK,” *CEPR Discussion Paper No. DP14429*, 2021.
- Amano, Robert, Stefano Gnocchi, Sylvain Leduc, and Joel Wagner**, “Average is good enough: Average-inflation targeting and the ELB,” Staff Working Papers 20-31, Bank of Canada July 2020.
- Auclert, Adrien**, “Monetary policy and the redistribution channel,” *American Economic Review*, June 2019, 109 (6), 2333–2367.
- Benchimol, Jonathan and Lahcen Bounader**, “Optimal monetary policy under bounded rationality,” IMF Working Papers, International Monetary Fund August 2019.
- Bernanke, Ben S., Michael T. Kiley, and John M. Roberts**, “Monetary Policy Strategies for a Low-Rate Environment,” Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System (U.S.) February 2019.
- Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J. Sargent**, “Inequality, business cycles, and monetary-fiscal policy,” NBER Working Papers, National Bureau of Economic Research, Inc June 2018.
- Bilbiie, Florin Ovidiu**, “Monetary policy and heterogeneity: An analytical framework,” *CEPR Discussion Paper No. DP12601*, 2019.
- Blanchard, Olivier J**, “Debt, deficits, and finite horizons,” *Journal of Political Economy*, 1985, 93 (2), 223–247.
- Bowlus, Audra, Emilien Gouin-Bonenfant, Huju Liu, Lance Lochner, and Youngmin Park**, “Four decades of Canadian earnings dynamics across workers and firms,” Technical Report 2020.
- Brouillette, Dany, Guyllaume Faucher, Martin Küncl, Austin McWhirter, and Youngmin Park**, “Potential output and the neutral rate in Canada: 2021 update,” *Bank of Canada Staff Analytical Note No. 2021-6*.
- Browning, Martin and Thomas F. Crossley**, “The long-run cost of job loss as measured by consumption changes,” *Journal of Econometrics*, July 2008, 145 (1-2), 109–120.
- Brzozowski, Matthew, Martin Gervais, Paul Klein, and Michio Suzuki**, “Consumption, income, and wealth inequality in Canada,” *Review of Economic Dynamics*, January 2010, 13 (1), 52–75.

- Cecioni, Martina, Günter Coenen, Roberto Motto, Hervé Le Bihan, Viktors Ajevskis, Ugo Albertazzi, Niels Gilbert, Alexander Al-Haschimi, Sandra Gomes, Friederike Bornemann et al.,** “The Ecb’s Price Stability Framework: Past Experience, and Current and Future Challenges,” *ECB Occasional Paper No. 2021269*, 2021.
- Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia,** “Innocent bystanders? Monetary policy and inequality,” *Journal of Monetary Economics*, 2017, 88 (C), 70–89.
- Corrigan, Paul, Hélène Desgagnés, José Dorich, Claudia Godbout, Vadym Lepetyuk, Wataru Miyamoto, and Yang Zhang,** “ToTEM III,” mimeo, Bank of Canada May 2020.
- Deaton, Angus,** “Saving and liquidity constraints,” *Econometrica*, 1991, pp. 1221–1248.
- Dorich, José, Rhys Mendes, and Yang Zhang,** “The Bank of Canada’s “Horse Race” of Alternative Monetary Policy Frameworks: Some Interim Results,” *mimeo*, 2020.
- Feiveson, Laura, Nils Gornemann, Julie L. Hotchkiss, Karel Mertens, and Jae W. Sim,** “Distributional considerations for monetary policy strategy,” Finance and Economics Discussion Series 2020-073, Board of Governors of the Federal Reserve System (U.S.) August 2020.
- Furceri, Davide, Prakash Loungani, and Aleksandra Zdzienicka,** “The effects of monetary policy shocks on inequality,” *Journal of International Money and Finance*, 2018, 85 (C), 168–186.
- Gali, Jordi,** *Monetary Policy, Inflation, and the Business Cycle*, Princeton University Press, 2015.
- Gervais, Olivier and Marc-André Gosselin,** “Analyzing and forecasting the Canadian economy through the LENS model,” Technical Report, Bank of Canada 2014.
- Grand, François Le, Alaïs Martin-Baillon, and Xavier Ragot,** “Should monetary policy care about redistribution? Optimal fiscal and monetary policy with heterogeneous agents,” Technical Report, mimeo 2020.
- Güvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song,** “What do data on millions of U.S. workers reveal about life-cycle earnings risk?,” NBER Working Papers 20913, National Bureau of Economic Research, Inc January 2015.
- Holm, Blomhoff, Pascal Paul, and Andreas Tischbirek,** “The Transmission of Monetary Policy under the Microscope,” Working Papers, Federal Reserve Bank of San Francisco 2020.
- Kaplan, Greg and Giovanni L. Violante,** “Microeconomic heterogeneity and macroeconomic shocks,” *Journal of Economic Perspectives*, Summer 2018, 32 (3), 167–194.

- , **Benjamin Moll**, and **Giovanni L. Violante**, “Monetary Policy According to HANK,” *American Economic Review*, March 2018, 108 (3), 697–743.
- Karibzhanov, Iskander**, “Towards a HANK model for Canada: Estimating a Canadian income process,” Staff Discussion Paper 2020-13, Bank of Canada 2020.
- Kryvtsov, Oleksiy, Malik Shukayev, and Alexander Ueberfeldt**, “Adopting price-level targeting under imperfect credibility,” Technical Report, Bank of Canada Working Paper 2008.
- Nakajima, Makoto and Vladimir Smirnyagin**, “Cyclical labor income risk,” Working Papers 19-34, Federal Reserve Bank of Philadelphia September 2019.
- Nakata, Taisuke, Sebastian Schmidt, and Flora Budianto**, “Average inflation targeting and the interest rate lower bound,” Working Paper Series, European Central Bank April 2020.
- Nessen, Marianne and David Vestin**, “Average inflation targeting,” *Journal of Money, Credit and Banking*, October 2005, 37 (5), 837–863.
- Ravn, Morten O. and Vincent Sterk**, “Macroeconomic fluctuations with HANK & SAM: An analytical approach,” Discussion Papers 1633, Centre for Macroeconomics (CFM) October 2016.
- Rogoff, Kenneth**, “The optimal degree of commitment to an intermediate monetary target,” *The Quarterly Journal of Economics*, 1985, 100 (4), 1169–1189.
- Rotemberg, Julio J**, “Monopolistic price adjustment and aggregate output,” *The Review of Economic Studies*, 1982, 49 (4), 517–531.
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron**, “Cyclical dynamics in idiosyncratic labor market risk,” *Journal of Political Economy*, June 2004, 112 (3), 695–717.
- Svensson, Lars E.O.**, “Monetary policy strategies for the Federal Reserve,” *International Journal of Central Banking*, February 2020, 16 (1), 133–193.
- Vestin, David**, “Price-level versus inflation targeting,” *Journal of Monetary Economics*, October 2006, 53 (7), 1361–1376.
- Woodford, Michael**, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, 2003.
- Yaari, Menahem E**, “Uncertain lifetime, life insurance, and the theory of the consumer,” *The Review of Economic Studies*, 1965, 32 (2), 137–150.

A Details of Acharya, Dogra, and challe (2020) model

A.1 Households optimization problem

The individual household's choice problem can be written as:

$$\begin{aligned} \max_{c_{s,t}^i, l_{s,t}^i} \quad & \mathbb{E}_s \sum_{t=s}^{\infty} (\beta\vartheta)^{t-s} \left[v_t \left(-\frac{1}{\gamma} e^{-\gamma c_{s,t}^i} - \rho e^{-\frac{1}{\rho}(\xi_{s,t}^i - l_{s,t}^i)} \right) \right] \\ \text{s.t.} \quad & c_{s,t}^i + q_t a_{s,t+1}^i = a_{s,t}^i + w_t l_{s,t}^i + T_t, \\ & \xi_{s,t}^i \sim \mathcal{N}(\bar{\xi}, \sigma_t^2) \\ & a_{s,s}^i = 0 \end{aligned}$$

Letting $\lambda_{s,t}^i$ be the Lagrange multiplier on the budget constraint, the Langrangian is given by

$$\begin{aligned} \mathcal{L} = \mathbb{E}_s \sum_{t=s}^{\infty} (\beta\vartheta)^{t-s} \left\{ v_t \left[-\frac{1}{\gamma} e^{-\gamma c_{s,t}^i} - \rho e^{-\frac{1}{\rho}(\xi_{s,t}^i - l_{s,t}^i)} \right] \right. \\ \left. + \lambda_{s,t}^i \left[a_{s,t}^i + w_t l_{s,t}^i + T_t - c_{s,t}^i - q_t a_{s,t+1}^i \right] \right\} \end{aligned}$$

First order conditions. Taking the first order conditions of this problem leads to

$$\begin{aligned} [c_{s,t}^i] : v_t e^{-\gamma c_{s,t}^i} &= \lambda_{s,t}^i \\ [l_{s,t}^i] : v_t e^{-\frac{1}{\rho}(\xi_{s,t}^i - l_{s,t}^i)} &= w_t \lambda_{s,t}^i \\ [a_{s,t+1}^i] : \lambda_{s,t}^i &= \beta\vartheta q_t^{-1} \mathbb{E}_t \lambda_{s,t+1}^i \\ [\lambda_{s,t}^i] : c_{s,t}^i + q_t a_{s,t+1}^i &= a_{s,t}^i + w_t l_{s,t}^i + T_t \end{aligned}$$

Rearranging the first order conditions gives three equations: (i) the standard intertemporal Euler equation, (ii) the intratemporal equation connecting consumption to labour supply, and (iii) the budget constraints. These equations fully characterize the individual household's decision rules:

$$\begin{aligned} 1 &= \mathbb{E}_t \left(\beta\vartheta q_t^{-1} \frac{v_{t+1}}{v_t} \frac{e^{-\gamma c_{s,t+1}^i}}{e^{-\gamma c_{s,t}^i}} \right), \\ w_t &= \frac{e^{-\frac{1}{\rho}(\xi_{s,t}^i - l_{s,t}^i)}}{e^{-\gamma c_{s,t}^i}}, \\ c_{s,t}^i + q_t a_{s,t+1}^i &= a_{s,t}^i + w_t l_{s,t}^i + T_t. \end{aligned}$$

Table 10: Model summary

<i>Demand Block</i>	
Euler Equation	$1 = \mathbb{E}_t \left(\beta R_t \frac{v_{t+1}}{v_t} \frac{e^{-\gamma c_{s,t+1}^i}}{e^{-\gamma c_{s,t}^i}} \right)$
Labour supply	$w_t = \frac{e^{-\frac{1}{\rho}(\xi_{s,t}^i - l_{s,t}^i)}}{e^{-\gamma c_{s,t}^i}}$
Budget constraint	$c_{s,t}^i + \frac{\theta}{R_t} a_{s,t+1}^i = a_{s,t}^i + w_t l_{s,t}^i + T_t$
Uninsurable shock	$\xi_{s,t}^i \sim \mathcal{N}(\bar{\xi}, \sigma_\xi^2)$
<i>Supply Block</i>	
Phillip's curve	$(\Pi_t - 1)\Pi_t = \frac{\varepsilon_t}{\Phi} \left(1 - \frac{\varepsilon_t - 1}{\varepsilon_t} \frac{z_t}{(1-\tau)w_t} \right) + \mathbb{E}_t \left(\frac{1}{R_t} \frac{y_{t+1} z_{t+1} w_{t+1}}{y_t z_{t+1} w_t} (\Pi_{t+1} - 1) \Pi_{t+1} \right)$
Aggregate output	$y_t = z_t n_t - \frac{\Phi}{2} (\Pi_t - 1)^2 y_t$
<i>Aggregation, market clearing</i>	
Good market	$c_t = y_t$
Labour market	$n_t = \rho \log w_t - \gamma \rho c_t + \bar{\xi}$
Bond market	$a_t = 0$
Gov Bond market	$B_t = 0$
Agg. Budget constraint	$c_t = w_t (\rho \log w_t - \gamma \rho c_t + \bar{\xi}) + T_t$
<i>Monetary policy rule</i>	
Real interest rate	$R_t = \frac{1+i_t}{\Pi_{t+1}}$
Taylor rule	$1 + i_t = (1 + i^*)^{1-\rho_i} (1 + i_{t-1})^{\rho_i} \left(\frac{\Pi_t}{\bar{\Pi}} \right)^{\theta_\pi (1-\rho_i)}$
<i>Shocks</i>	
Demand shock	$\log v_t = \rho_v \log v_{t-1} + u_{v,t}$
Cost-push shock	$\log \varepsilon_t = \rho_\varepsilon \log \varepsilon_{t-1} + u_{\varepsilon,t}$
Tech shock	$\log z_t = \rho_z \log z_{t-1} + u_{z,t}$
Income risk	$w_t^2 \sigma_t^2 = w^2 \sigma^2 e^{2\phi(y_t - y)}$
An aggregate variable x_t is defined as follow: $x_t = \sum_{s=-\infty}^t (1 - \theta) \theta^{s-t} \int_i x_{s,t}^i di$	

A.2 Solving the household's problem

Step 1: Guess a linear consumption policy function The first step in solving the household problem consists of guessing a linear expression of consumption, where linearity is defined with respect to cash-on-hand, $x_{s,t}^i$. That is:

$$c_{s,t}^i = c_t + \mu_t x_{s,t}^i \quad (32)$$

where c_t and μ_t are determined endogenously in equilibrium. This guess allows us to derive a dynamics equation of cash-on-hand summarized in the following lemma.

Lemma 3 *The cash-on-hand of a household i in a cohort s at time t evolves according to the following equation:*

$$x_{s,t+1}^i = \frac{R_t}{\vartheta} \left(1 - (1 + \gamma \rho w_t) \mu_t \right) x_{s,t}^i + w_{t+1} (\xi_{s,t+1}^i - \bar{\xi}) \quad (33)$$

Proof. Using the definition of cash-on-hand and the household optimal labour supply equation, we rewrite $x_{s,t+1}^i$ as:

$$x_{s,t+1}^i = \frac{R_t}{\vartheta} \left(x_{s,t}^i + w_t \bar{\xi} + \rho w_t \log w_t + T_t - (1 + \gamma \rho w_t) c_{s,t}^i \right) + w_{t+1} (\xi_{s,t+1}^i - \bar{\xi}) \quad (34)$$

We next aggregate individual household budget constraint and impose real bond market clearing condition to obtain the following identity:

$$c_t (1 + \gamma \rho w_t) = w_t (\rho \log w_t + \bar{\xi}) + T_t. \quad (35)$$

The next step consists of plugging this last identity back into $x_{s,t+1}^i$:

$$x_{s,t+1}^i = \frac{R_t}{\vartheta} \left(x_{s,t}^i + (1 + \gamma \rho w_t) (c_t - c_{s,t}^i) \right) + w_{t+1} (\xi_{s,t+1}^i - \bar{\xi}). \quad (36)$$

■

Step 2: Rewrite the Euler equation The Euler equation is given by

$$1 = \mathbb{E}_t \left(\beta R_t \frac{v_{t+1}}{v_t} \frac{e^{-\gamma c_{s,t+1}^i}}{e^{-\gamma c_{s,t}^i}} \right). \quad (37)$$

For ease of derivation we rewrite this Euler equation to explicitly account for the state variable and the uncertainty associated with uninsurable income shocks. We write $c_t(x)$ and $c_{t+1}(x')$ to represent respectively the consumption policy function at time t with a cash-on-hand x ; and consumption policy function at time $t + 1$ with cash-on-hand x' .

From (33) we find an expression of future uninsurable risk given by:

$$\frac{\xi_{s,t+1}^i - \bar{\xi}}{\sigma_{t+1}} = \frac{x'}{w_{t+1} \sigma_{t+1}} - \frac{R_t}{\vartheta} \left(1 - (1 + \gamma \rho w_t) \mu_t \right) \frac{x}{w_{t+1} \sigma_{t+1}} \quad (38)$$

and rewrite the Euler equation as:

$$1 = \mathbb{E}_t \left[\beta R_t \frac{v_{t+1}}{v_t} \frac{1}{\sigma_{y,t+1}} \int \frac{e^{-\gamma c_{t+1}(x')}}{e^{-\gamma c_t(x)}} f \left(\frac{x'}{\sigma_{y,t+1}} - \frac{R_t}{\vartheta} (1 - (1 + \gamma \rho w_t) \mu_t) \frac{x}{\sigma_{y,t+1}} \right) dx' \right]. \quad (39)$$

where $f(\cdot)$ denotes the probability density function of a standard normal distribution and $\sigma_{y,t} = \sigma_t w_t$.

Step 3: Steady state Euler equation

Lemma 4 *In steady state, the Euler equation is given by:*

$$\begin{aligned} 1 &= \frac{\beta R}{\sigma_y} \int e^{-\gamma \mu(x'-x)} f \left(\frac{x'-x}{\sigma_y} \right) dx' \\ 1 &= \frac{R}{\vartheta} (1 - (1 + \gamma \rho w) \mu) \iff \mu = \frac{1 - \frac{\vartheta}{R}}{1 + \gamma \rho w} \end{aligned} \quad (40)$$

which leads to the following two additional identities:

$$\begin{aligned} \gamma \mu \sigma_y &= \frac{\beta R}{\sigma_y} \int e^{-\gamma \mu(x'-x)} f' \left(\frac{x'-x}{\sigma_y} \right) dx' \\ (\gamma \mu \sigma_y)^2 &= \frac{\beta R}{\sigma_y} \int e^{-\gamma \mu(x'-x)} f'' \left(\frac{x'-x}{\sigma_y} \right) dx' \end{aligned} \quad (41)$$

where $f(\cdot)$ denotes the probability density function of a standard normal distribution and $\sigma_y = \sigma w$.

Proof. Using the Euler equation derived from the household's problem and setting all time-varying quantities to a common value delivers the Euler equation in steady state. We differentiate both sides of the Euler equation with respect to x and use the fact that $f'(x) = -x f(x)$ and $f''(x) = -f(x) - x f'(x)$ to obtain the two additional identities. ■

Step 4: Linearization of the Euler equation The inclusion of two additional aggregate shocks (preference shocks) and cost-push shocks renders the derivation of equilibrium equations characterizing the demand block of the model more challenging. However we can retain the tractability of the model by linearizing the model around the equilibrium without aggregate shocks. Following (Acharya et al. 2021), we linearize y_t and w_t and log-linearize the remaining variables in (39)

$$\begin{aligned} 1 &= \frac{\beta R}{\sigma_y} \mathbb{E}_t \left\{ (1 + \hat{R}_t) (1 + \hat{v}_{t+1} - \hat{v}_t) \left(1 - \frac{d \log \sigma_y}{dy} \hat{y}_{t+1} \right) \int e^{-\gamma \mu(x'-x)} \left(1 - \gamma (\hat{c}_{t+1}(x') - \hat{c}_t(x)) \right) \right. \\ &\quad f \left(\frac{x'-x}{\sigma_y} \right) \left(1 + \frac{d \log \sigma_y}{dy} \left(\frac{x'-x}{\sigma_y} \right)^2 \hat{y}_{t+1} + \frac{x'-x}{\sigma_y} \frac{x}{\sigma_y} \hat{R}_t + \frac{x'-x}{\sigma_y} \frac{R}{\vartheta} \gamma \rho \mu \frac{x}{\sigma_y} \hat{w}_t \right. \\ &\quad \left. \left. + \frac{x'-x}{\sigma_y} \left(1 - \frac{R}{\vartheta} \right) \frac{x}{\sigma_y} \hat{\mu}_t \right) dx' \right\} \end{aligned} \quad (42)$$

We simplify this approximation by ignoring higher order terms and using the expression of the Euler equation in steady state to have:

$$\begin{aligned}
0 = & \left[\frac{\beta R}{\sigma_y} \int e^{-\gamma \mu(x'-x)} f\left(\frac{x'-x}{\sigma_y}\right) dx' \right] \left[\mathbb{E}_t \hat{R}_t + \mathbb{E}_t \hat{v}_{t+1} - \hat{v}_t \right] \\
& - \left[\frac{\beta R}{\sigma_y} \int e^{-\gamma \mu(x'-x)} f\left(\frac{x'-x}{\sigma_y}\right) dx' \right] \frac{d \log \sigma_y}{dy} \mathbb{E}_t \hat{y}_{t+1} \\
& - \gamma \left[\frac{\beta R}{\sigma_y} \int e^{-\gamma \mu(x'-x)} f\left(\frac{x'-x}{\sigma_y}\right) \mathbb{E}_t \hat{c}_{t+1}(x') dx' \right] \\
& + \gamma \left[\frac{\beta R}{\sigma_y} \int e^{-\gamma \mu(x'-x)} f\left(\frac{x'-x}{\sigma_y}\right) dx' \right] \hat{c}_t(x) \\
& + \left[\frac{\beta R}{\sigma_y} \int e^{-\gamma \mu(x'-x)} \left(\frac{x'-x}{\sigma_y}\right)^2 f\left(\frac{x'-x}{\sigma_y}\right) dx' \right] \frac{d \log \sigma_y}{dy} \mathbb{E}_t \hat{y}_{t+1} \\
& + \left[\frac{\beta R}{\sigma_y} \int e^{-\gamma \mu(x'-x)} \left(\frac{x'-x}{\sigma_y}\right) f\left(\frac{x'-x}{\sigma_y}\right) dx' \right] \frac{x}{\sigma_y} \mathbb{E}_t \hat{R}_t \\
& + \left[\frac{\beta R}{\sigma_y} \int e^{-\gamma \mu(x'-x)} \left(\frac{x'-x}{\sigma_y}\right) f\left(\frac{x'-x}{\sigma_y}\right) dx' \right] \frac{R}{\vartheta} \gamma \rho \mu \frac{x}{\sigma_y} \hat{w}_t \\
& + \left[\frac{\beta R}{\sigma_y} \int e^{-\gamma \mu(x'-x)} \left(\frac{x'-x}{\sigma_y}\right) f\left(\frac{x'-x}{\sigma_y}\right) dx' \right] \left(1 - \frac{R}{\vartheta}\right) \frac{x}{\sigma_y} \hat{\mu}_t
\end{aligned} \tag{43}$$

We next rearrange this expression using the identities in lemma (4) and the linearized policy function $\hat{c}_t(x) = \hat{c}_t + \hat{\mu}_t x = \hat{y}_t + \hat{\mu}_t x$ and $\hat{c}_{t+1}(x) = \hat{c}_{t+1} + \hat{\mu}_{t+1} x = \hat{y}_{t+1} + \hat{\mu}_{t+1} x$ to have:

$$\begin{aligned}
0 = & \left(\mathbb{E}_t \hat{R}_t + (\rho_v - 1) \hat{v}_t + (-\gamma + (\gamma \mu \sigma_y)^2 \frac{d \log \sigma_y}{dy}) \mathbb{E}_t \hat{y}_{t+1} + (\gamma \mu \sigma_y)^2 \mathbb{E}_t \hat{\mu}_{t+1} + \gamma \hat{y}_t \right) \\
& + \gamma \mu x \left(-\mathbb{E}_t \hat{\mu}_{t+1} + \hat{\mu}_t - \mathbb{E}_t \hat{R}_t + \frac{R}{\vartheta} \gamma \rho \mu \hat{w}_t - \left(1 - \frac{R}{\vartheta}\right) \hat{\mu}_t \right)
\end{aligned} \tag{44}$$

Setting the terms in brackets to zero yields:

$$\begin{aligned}
\hat{y}_t = & \left(1 - \gamma (\mu \sigma_y)^2 \frac{d \log \sigma_y}{dy} \right) \mathbb{E}_t \hat{y}_{t+1} - \frac{1}{\gamma} \mathbb{E}_t \hat{R}_t - \gamma (\mu \sigma_y)^2 \mathbb{E}_t \hat{\mu}_{t+1} + \frac{1}{\gamma} (1 - \rho_v) \hat{v}_t \\
\hat{\mu}_t = & -\gamma \rho \mu \hat{w}_t + \frac{\vartheta}{R} \mathbb{E}_t (\hat{R}_t + \hat{\mu}_{t+1})
\end{aligned} \tag{45}$$

Using the steady state income risk pass-through and the definition of income risk, we have:

$$\begin{aligned}
\hat{y}_t = & \left(1 - \frac{\Lambda \phi}{\gamma} \right) \mathbb{E}_t \hat{y}_{t+1} - \frac{1}{\gamma} \mathbb{E}_t \hat{R}_t - \frac{\Lambda}{\gamma} \mathbb{E}_t \hat{\mu}_{t+1} + \frac{1}{\gamma} (1 - \rho_v) \hat{v}_t \\
\hat{\mu}_t = & -\gamma \rho \mu \hat{w}_t + \tilde{\beta} \mathbb{E}_t (\hat{R}_t + \hat{\mu}_{t+1})
\end{aligned} \tag{46}$$

where $\Lambda = (\gamma \mu \sigma_y)^2 = (\gamma \mu \sigma w)^2$, $\tilde{\beta} = \frac{\vartheta}{R}$ and $\phi = \frac{d \log \sigma_y}{dy}$.

Table 11: Linearized model summary

<i>Demand Block</i>	
Aggregate Euler Equation	$\hat{y}_t = \left(1 - \frac{\Lambda\phi}{\gamma}\right) \mathbb{E}_t \hat{y}_{t+1} - \frac{1}{\gamma} \mathbb{E}_t \hat{R}_t - \frac{\Lambda}{\gamma} \mathbb{E}_t \hat{\mu}_{t+1} + \frac{1}{\gamma} (1 - \rho_v) \hat{v}_t$
Income risk pass-through	$\hat{\mu}_t = -\frac{(1-\tilde{\beta})\gamma\rho}{1+\gamma\rho w} \hat{w}_t + \tilde{\beta} \mathbb{E}_t (\hat{R}_t + \hat{\mu}_{t+1})$
Aggregate Labour supply	$\hat{l}_t = \rho \frac{\hat{w}_t}{w} - \gamma \rho \hat{c}_t$
<i>Supply Block</i>	
Phillip's curve	$\pi_t = \frac{\tilde{\beta}}{\theta} \mathbb{E}_t \pi_{t+1} + \kappa (\hat{y}_t - \hat{y}_t^n) + \sigma_\varepsilon \hat{\varepsilon}_t$
Aggregate output	$\hat{y}_t = y \hat{z}_t + z \hat{n}_t$
<i>Market clearing</i>	
Good market	$\hat{c}_t = \hat{y}_t$
Labour market	$\hat{n}_t = \hat{l}_t$
<i>Monetary policy rule</i>	
Real interest rate	$\hat{R}_t = i_t - \pi_{t+1}$
Taylor rule	$i_t = \rho_i i_{t-1} + (1 - \rho_i) (i^* + \theta_\pi (\pi_t - \bar{\pi}))$
<i>Shocks</i>	
Demand shock	$\hat{v}_t = \rho_v \hat{v}_{t-1} + \hat{u}_{v,t}$
Cost-push shock	$\hat{\varepsilon}_t = \rho_\varepsilon \hat{\varepsilon}_{t-1} + \hat{u}_{\varepsilon,t}$
Tech shock	$\hat{z}_t = \rho_z \hat{z}_{t-1} + \hat{u}_{z,t}$
Income risk	$\hat{\sigma}_t + \frac{\hat{w}_t}{w} = \phi \hat{y}_t$

We further simplify the income risk pass-through equation by noticing that in steady state, $\mu = \frac{1-\frac{\theta}{R}}{1+\gamma\rho w} \Rightarrow \gamma\rho\mu = \frac{(1-\tilde{\beta})\gamma\rho}{1+\gamma\rho w}$.

$$\begin{aligned}
\hat{y}_t &= \left(1 - \frac{\Lambda\phi}{\gamma}\right) \mathbb{E}_t \hat{y}_{t+1} - \frac{1}{\gamma} \mathbb{E}_t \hat{R}_t - \frac{\Lambda}{\gamma} \mathbb{E}_t \hat{\mu}_{t+1} + \frac{1}{\gamma} (1 - \rho_v) \hat{v}_t \\
\hat{\mu}_t &= -\frac{(1-\tilde{\beta})\gamma\rho}{1+\gamma\rho w} \hat{w}_t + \tilde{\beta} \mathbb{E}_t (\hat{R}_t + \hat{\mu}_{t+1})
\end{aligned} \tag{47}$$

A.3 Supply block

Price setting. The price-setting problem can be written as:

$$\begin{aligned}
&\max_{p_t(j), n_t(j), y_t(j)} \mathbb{E}_t \sum_{t=0}^{\infty} Q_{t|0} \left\{ \frac{p_t(j)}{P_t} y_t(j) - (1 - \tau) w_t n_t(j) \right\} \\
&\text{s.t. } y_t(j) = \left(\frac{p_t(j)}{P_t} \right)^{-\varepsilon_t} y_t, \\
&y_t(j) = z_t n_t(j) - \frac{\Phi}{2} \left(\frac{p_t(j)}{p_{t-1}(j)} - 1 \right)^2 y_t.
\end{aligned}$$

We start by finding an expression of $n_t(j)$ from the aggregate resource constraint and plugging it back into the objective function to have:

$$\max_{p_t(j)} \mathbb{E}_t \sum_{t=0}^{\infty} Q_{t|0} \left\{ \left(\frac{p_t(j)}{P_t} \right)^{1-\varepsilon_t} y_t - (1-\tau)w_t \left(\left(\frac{p_t(j)}{P_t} \right)^{-\varepsilon_t} \frac{y_t}{z_t} + \frac{\Phi}{2} \left(\frac{p_t(j)}{p_{t-1}(j)} - 1 \right)^2 \frac{y_t}{z_t} \right) \right\}$$

We next take the first order conditions with respect to $p_t(j)$, impose the symmetric equilibrium $p_t(j) = P_t$, and rearrange to obtain:

$$(\Pi_t - 1)\Pi_t = \frac{\varepsilon_t}{\Phi} \left(1 - \frac{\varepsilon_t - 1}{\varepsilon_t} \frac{z_t}{(1-\tau)w_t} \right) + \mathbb{E}_t \left(\frac{1}{R_t} \frac{y_{t+1}z_{t+1}w_{t+1}}{y_t z_{t+1} w_t} (\Pi_{t+1} - 1)\Pi_{t+1} \right) \quad (48)$$

Log-linearization. we start by rewriting the NKPC as:

$$(\Pi e^{\pi_t} - 1)\Pi e^{\pi_t} = \frac{\varepsilon}{\Phi} e^{\hat{\varepsilon}_t} \left(1 - \frac{\varepsilon e^{\hat{\varepsilon}_t} - 1}{\varepsilon} \frac{z e^{\hat{z}_t}}{(1-\tau)w e^{\hat{\varepsilon}_t + \hat{w}_t}} \right) + \mathbb{E}_t \left(\frac{1}{R} \frac{e^{\hat{y}_{t+1} + \hat{z}_t + \hat{w}_{t+1}}}{e^{\hat{y}_t + \hat{z}_{t+1} + \hat{w}_t}} (\Pi e^{\pi_{t+1}} - 1)\Pi e^{\pi_{t+1}} \right) \quad (49)$$

Taking the first order log-linearization with $\Pi = 1$ and ignoring higher order terms yields:

$$\pi_t = \frac{\varepsilon}{\Phi} \frac{z}{(1-\tau)w} \frac{\varepsilon - 1}{\varepsilon} (\hat{w}_t - \hat{z}_t) + \left(1 - \frac{z}{(1-\tau)w} \right) \frac{\varepsilon}{\Phi} \hat{\varepsilon}_t + \frac{\tilde{\beta}}{\vartheta} \mathbb{E}_t \pi_{t+1}. \quad (50)$$

Given that flexible price output is $\hat{y}_t^n = \frac{\rho+y}{1+\gamma\rho} \hat{z}_t$ and the sticky price output is $\hat{y}_t = \frac{\rho}{1+\gamma\rho} \frac{\hat{w}_t}{w} + \frac{y}{1+\gamma\rho} \hat{z}_t$, we deduce that:

$$\hat{w}_t - \hat{z}_t = \frac{1+\gamma\rho}{\rho} (\hat{y}_t - \hat{y}_t^n). \quad (51)$$

Using this last equality with the fact that $\frac{z}{(1-\tau)w} \frac{\varepsilon-1}{\varepsilon} = 1$ in steady state, we can simplify and obtain:

$$\begin{aligned} \pi_t &= \frac{\varepsilon}{\Phi} \frac{1+\gamma\rho}{\rho} (\hat{y}_t - \hat{y}_t^n) + \frac{\tilde{\beta}}{\vartheta} \mathbb{E}_t \pi_{t+1} + \frac{\varepsilon}{1-\varepsilon} \frac{1}{\Phi} \hat{\varepsilon}_t \\ \Rightarrow \pi_t &= \kappa (\hat{y}_t - \hat{y}_t^n) + \frac{\tilde{\beta}}{\vartheta} \mathbb{E}_t \pi_{t+1} + \sigma_\varepsilon \hat{\varepsilon}_t \end{aligned} \quad (52)$$

where $\kappa = \frac{\varepsilon}{\Phi} \frac{1+\gamma\rho}{\rho}$ and $\sigma_\varepsilon = \frac{\varepsilon}{1-\varepsilon} \frac{1}{\Phi}$

B Steady state

Letting productivity be normalized to $z = 1$, the New Keynesian Philipps curve in the steady state will be given by:

$$\left(1 - \frac{1}{R}\right)(\pi - 1)\pi = \frac{\varepsilon}{\Phi} \left(1 - \frac{\varepsilon - 1}{\varepsilon} \frac{1}{(1 - \tau)w}\right). \quad (53)$$

This equation illustrates a long run trade-off between inflation and economic activity. As usual, the social planner chooses the payroll subsidy, τ , to eliminate steady state distortions created by monopolistic competition. However, the presence of counter-cyclical idiosyncratic risk gives an additional motive for the social planner to deviate from price stability. This is because higher output not only reduces the amount of income risk faced by households, but also improves the ability of households to self-insure themselves against idiosyncratic income risk. Following Acharya et al. (2021), the optimal subsidy consistent with zero inflation in the steady state is then given by:

$$\tau = \frac{1}{\varepsilon} + \frac{\varepsilon - 1}{\varepsilon} \frac{\Omega}{1 + \Omega} (1 + \gamma\rho),$$

where Ω satisfies

$$\begin{aligned} w &= \frac{1 + \Omega}{1 - \gamma\rho\Omega}, \\ \Omega &= \frac{\Lambda(1 - \frac{\phi}{\gamma})}{(1 - \Lambda)(1 - \tilde{\beta})}, \\ \Lambda &= \gamma^2 \mu^2 w^2 \sigma^2, \\ \tilde{\beta} &= \frac{\vartheta}{R}. \end{aligned}$$

This presence of counter-cyclical risk leads to a $\Omega > 1$ and induces a higher steady state wage, $w > 1$, compared to the efficient steady state wage, $w = 1$. In the steady state, the time invariant IS curve and MPC are given by:

$$\begin{aligned} \beta R &= e^{-\frac{\Lambda}{2}}, \\ \mu &= \frac{1 - \tilde{\beta}}{1 + \gamma\rho w}. \end{aligned}$$

When the the mean idiosyncratic risk is set to $\bar{\xi} = 1 + \gamma\rho$, output and aggregate employment are given by:

$$\begin{aligned} y &= \frac{\rho \log w + \bar{\xi}}{1 + \gamma\rho} = 1 + \frac{\rho \log w}{1 + \gamma\rho}, \\ n &= \rho \log w - \gamma\rho y + \bar{\xi} = y. \end{aligned}$$

	RANK efficient	Flexible Price HANK	Sticky price HANK
<i>Dynamics</i>			
Frictions		<ul style="list-style-type: none"> • Mpl. competition • Incomplete market 	<ul style="list-style-type: none"> • Mpl. competition • Incomplete market • Nominal rigidity
Output	$y_t^e = z_t \frac{\rho \log z_t + \bar{\xi}}{1 + \gamma \rho z_t}$	$y_t^n = z_t \frac{\rho(\log w + \log z_t) + \bar{\xi}}{1 + \gamma \rho z_t}$	$y_t = z_t \frac{\rho \log w_t + \bar{\xi}}{1 + \gamma \rho z_t}$
<i>Steady state</i>			
Frictions		• Incomplete market	• Incomplete market
Payroll subsidy	$\tau = \frac{1}{\varepsilon}$	$\tau = \frac{1}{\varepsilon} + \frac{1 + \gamma \rho}{\varepsilon} \frac{\Omega}{1 + \Omega}$	$\tau = \frac{1}{\varepsilon} + \frac{1 + \gamma \rho}{\varepsilon} \frac{\Omega}{1 + \Omega}$
Inflation	$\pi^e = 1$	$\pi^n = 1$	$\pi = 1$
Output	$y^e = \frac{\bar{\xi}}{1 + \gamma \rho} = 1$	$y^n = \frac{\rho \log w + \bar{\xi}}{1 + \gamma \rho}$	$y = \frac{\rho \log w + \bar{\xi}}{1 + \gamma \rho}$
Employment	$n^e = 1$	$n^n = y^n = y$	$n = y$
Wage	$w^e = 1$	$w^n = \frac{1 + \Omega}{1 - \gamma \rho \Omega}$	$w = \frac{1 + \Omega}{1 - \gamma \rho \Omega}$

C Proofs

Proof of proposition 1. For a generic variable x_t , let $\hat{x}_t = \log(x_t) - \log(x)$ be the log-linear approximation of a variable around its steady state allocation. The instantaneous social welfare function is given by the following utility:

$$U_t = v_t \left[-\frac{1}{\gamma} e^{-\gamma y_t} - \rho e^{\frac{1}{\rho}(n_t - \bar{\xi})} \right] \Sigma_t. \quad (54)$$

The second-order Taylor expansion of U_t around the steady state (y, n, Σ, v) gives:

$$\begin{aligned}
U_t \approx & U + v e^{-\gamma y} \Sigma y \left(\frac{y_t - y}{y} \right) - v e^{\frac{1}{\rho}(n - \bar{\xi})} \Sigma n \left(\frac{n_t - n}{n} \right) \\
& - v \left(\frac{1}{\gamma} e^{-\gamma y} + \rho e^{\frac{1}{\rho}(n - \bar{\xi})} \right) \Sigma \left(\frac{\Sigma_t - \Sigma}{\Sigma} \right) - v \left(\frac{1}{\gamma} e^{-\gamma y} + \rho e^{\frac{1}{\rho}(n - \bar{\xi})} \right) \Sigma \left(\frac{v_t - v}{v} \right) \\
& - \frac{\gamma}{2} v e^{-\gamma y} \Sigma y^2 \left(\frac{y_t - y}{y} \right)^2 - \frac{1}{2\rho} v e^{\frac{1}{\rho}(n - \bar{\xi})} \Sigma n^2 \left(\frac{n_t - n}{n} \right)^2 \\
& + v e^{-\gamma y} \Sigma y \left(\frac{y_t - y}{y} \right) \left(\frac{\Sigma_t - \Sigma}{\Sigma} \right) + v e^{-\gamma y} \Sigma y \left(\frac{y_t - y}{y} \right) \left(\frac{v_t - v}{v} \right) \\
& - v e^{\frac{1}{\rho}(n - \bar{\xi})} \Sigma n \left(\frac{n_t - n}{n} \right) \left(\frac{\Sigma_t - \Sigma}{\Sigma} \right) - v e^{\frac{1}{\rho}(n - \bar{\xi})} \Sigma n \left(\frac{n_t - n}{n} \right) \left(\frac{v_t - v}{v} \right) \\
& - v \left(\frac{1}{\gamma} e^{-\gamma y} + \rho e^{\frac{1}{\rho}(n - \bar{\xi})} \right) \Sigma \left(\frac{\Sigma_t - \Sigma}{\Sigma} \right) \left(\frac{v_t - v}{v} \right).
\end{aligned} \quad (55)$$

We can rearrange this expression to get

$$\begin{aligned}
\frac{U_t - U}{\nu e^{-\gamma y} \Sigma y} &\approx \left(\frac{y_t - y}{y} \right) - \frac{e^{\frac{1}{\rho}(n - \bar{\xi})} n}{e^{-\gamma y} y} \left(\frac{n_t - n}{n} \right) \\
&\quad - \left(\frac{1}{\gamma y} + \rho \frac{e^{\frac{1}{\rho}(n - \bar{\xi})}}{e^{-\gamma y} y} \right) \left(\frac{\Sigma_t - \Sigma}{\Sigma} \right) - \left(\frac{1}{\gamma y} + \rho \frac{e^{\frac{1}{\rho}(n - \bar{\xi})}}{e^{-\gamma y} y} \right) \left(\frac{\nu_t - \nu}{\nu} \right) \\
&\quad - \frac{\gamma}{2} y \left(\frac{y_t - y}{y} \right)^2 - \frac{1}{2\rho} \frac{e^{\frac{1}{\rho}(n - \bar{\xi})} n^2}{e^{-\gamma y} y} \left(\frac{n_t - n}{n} \right)^2 \\
&\quad + \left(\frac{y_t - y}{y} \right) \left(\frac{\Sigma_t - \Sigma}{\Sigma} \right) + \left(\frac{y_t - y}{y} \right) \left(\frac{\nu_t - \nu}{\nu} \right) \\
&\quad - \frac{e^{\frac{1}{\rho}(n - \bar{\xi})} n}{e^{-\gamma y} y} \left(\frac{n_t - n}{n} \right) \left(\frac{\Sigma_t - \Sigma}{\Sigma} \right) - \frac{e^{\frac{1}{\rho}(n - \bar{\xi})} n}{e^{-\gamma y} y} \left(\frac{n_t - n}{n} \right) \left(\frac{\nu_t - \nu}{\nu} \right) \\
&\quad - \left(\frac{1}{\gamma y} + \rho \frac{e^{\frac{1}{\rho}(n - \bar{\xi})}}{e^{-\gamma y} y} \right) \left(\frac{\Sigma_t - \Sigma}{\Sigma} \right) \left(\frac{\nu_t - \nu}{\nu} \right).
\end{aligned} \tag{56}$$

Using the equilibrium condition $n = \rho \log w - \gamma \rho y + \bar{\xi}$, we have

$$\frac{e^{\frac{1}{\rho}(n - \bar{\xi})}}{e^{-\gamma y}} = w. \tag{57}$$

Imposing this relation together with $y = n$ in the previous approximation gives:

$$\begin{aligned}
\frac{U_t - U}{\nu e^{-\gamma y} \Sigma y} &\approx \left(\frac{y_t - y}{y} \right) - w \left(\frac{n_t - n}{n} \right) \\
&\quad - \left(\frac{1}{\gamma y} + \rho \frac{w}{y} \right) \left(\frac{\Sigma_t - \Sigma}{\Sigma} \right) - \left(\frac{1}{\gamma y} + \rho \frac{w}{y} \right) \left(\frac{\nu_t - \nu}{\nu} \right) \\
&\quad - \frac{\gamma}{2} y \left(\frac{y_t - y}{y} \right)^2 - \frac{1}{2\rho} w y \left(\frac{n_t - n}{n} \right)^2 \\
&\quad + \left(\frac{y_t - y}{y} \right) \left(\frac{\Sigma_t - \Sigma}{\Sigma} \right) + \left(\frac{y_t - y}{y} \right) \left(\frac{\nu_t - \nu}{\nu} \right) \\
&\quad - w \left(\frac{n_t - n}{n} \right) \left(\frac{\Sigma_t - \Sigma}{\Sigma} \right) - w \left(\frac{n_t - n}{n} \right) \left(\frac{\nu_t - \nu}{\nu} \right) \\
&\quad - \left(\frac{1}{\gamma y} + \rho \frac{w}{y} \right) \left(\frac{\Sigma_t - \Sigma}{\Sigma} \right) \left(\frac{\nu_t - \nu}{\nu} \right).
\end{aligned} \tag{58}$$

Next, we exploit the following log-linear approximation:

$$\begin{aligned}
\frac{x_t - x}{x} &= \frac{x_t}{x} - 1 \\
&= e^{\log(\frac{x_t}{x})} - 1 \\
&= e^{\hat{x}_t} - 1 \\
&\approx 1 + \hat{x}_t + \frac{\hat{x}_t^2}{2} - 1 \\
&\approx \hat{x}_t + \frac{\hat{x}_t^2}{2}
\end{aligned} \tag{59}$$

Plugging back these approximations into (58) gives

$$\begin{aligned} \frac{U_t - U}{e^{-\gamma y \Sigma y}} &\approx \left(\hat{y}_t + \frac{1}{2} \hat{y}_t^2 \right) - w \left(\hat{n}_t + \frac{1}{2} \hat{n}_t^2 \right) - \left(\frac{1}{\gamma y} + \frac{\rho w}{y} \right) \left(\hat{\Sigma}_t + \frac{1}{2} \hat{\Sigma}_t^2 \right) \\ &\quad - \left(\frac{1}{\gamma y} + \frac{\rho w}{y} \right) \left(\hat{v}_t + \frac{1}{2} \hat{v}_t^2 \right) - \frac{\gamma}{2} y \hat{y}_t^2 - \frac{1}{2\rho} w y \hat{n}_t^2 + \hat{y}_t \hat{\Sigma}_t + \hat{y}_t \hat{v}_t \\ &\quad - w \hat{n}_t \hat{\Sigma}_t - w \hat{n}_t \hat{v}_t - \left(\frac{1}{\gamma y} + \rho \frac{w}{y} \right) \hat{\Sigma}_t \hat{v}_t. \end{aligned} \quad (60)$$

Applying logarithm to both sides of the aggregate production function $y_t = \frac{z_t n_t}{1 + \frac{\Phi}{2} (\Pi_t - 1)^2}$ and subtracting its steady state counterpart, we get:

$$\begin{aligned} \hat{n}_t &= \hat{y}_t - \hat{z}_t + \frac{\Phi}{2} \pi_t^2 \\ \hat{n}_t^2 &= (\hat{y}_t - \hat{z}_t)^2 \end{aligned} \quad (61)$$

where the second equation ignores higher order terms.

Plugging this back into equation (60) gives:

$$\begin{aligned} \frac{U_t - U}{e^{-\gamma y \Sigma y}} &\approx \left(\hat{y}_t + \frac{1}{2} \hat{y}_t^2 \right) - w \left(\hat{y}_t - \hat{z}_t + \frac{\Phi}{2} \pi_t^2 + \frac{1}{2} (\hat{y}_t - \hat{z}_t)^2 \right) \\ &\quad - \left(\frac{1}{\gamma y} + \frac{\rho w}{y} \right) \left(\hat{\Sigma}_t + \frac{1}{2} \hat{\Sigma}_t^2 \right) - \left(\frac{1}{\gamma y} + \frac{\rho w}{y} \right) \left(\hat{v}_t + \frac{1}{2} \hat{v}_t^2 \right) - \frac{\gamma}{2} y \hat{y}_t^2 - \frac{1}{2\rho} w y (\hat{y}_t - \hat{z}_t)^2 \\ &\quad + \hat{y}_t \hat{\Sigma}_t + \hat{y}_t \hat{v}_t - w \left(\hat{y}_t - \hat{z}_t + \frac{\Phi}{2} \pi_t^2 \right) \hat{\Sigma}_t - w \left(\hat{y}_t - \hat{z}_t + \frac{\Phi}{2} \pi_t^2 \right) \hat{v}_t - \left(\frac{1}{\gamma y} + \rho \frac{w}{y} \right) \hat{\Sigma}_t \hat{v}_t \\ &\approx -\frac{1}{2} (\gamma y - 1) \hat{y}_t^2 - \frac{1}{2} w \Phi \pi_t^2 - \frac{1}{2} \left(\frac{1}{\gamma y} + \frac{\rho w}{y} \right) \hat{\Sigma}_t^2 - \frac{1}{2} w \left(1 + \frac{y}{\rho} \right) \left(\hat{y}_t^2 - 2 \hat{z}_t \hat{y}_t \right) \\ &\quad - (w - 1) \hat{y}_t - \left(\frac{1}{\gamma y} + \frac{\rho w}{y} \right) \hat{\Sigma}_t - (w - 1) \hat{y}_t \hat{\Sigma}_t + w \hat{z}_t \hat{\Sigma}_t \\ &\quad - (w - 1) \hat{y}_t \hat{v}_t - \left(\frac{1}{\gamma y} + \rho \frac{w}{y} \right) \hat{\Sigma}_t \hat{v}_t + t.i.p. \end{aligned} \quad (62)$$

where t.i.p. indicates terms independent of policy.

The natural output level depends entirely on productivity and is given by $\hat{y}_t^n = \frac{y + \rho}{1 + \rho \gamma} \hat{z}_t$. We use this last expression to express \hat{z}_t in terms of \hat{y}_t^n and plug it back in:

$$\begin{aligned} \frac{U_t - U}{e^{-\gamma y \Sigma y}} &\approx -\frac{1}{2} (\gamma y - 1) \hat{y}_t^2 - \frac{1}{2} w \Phi \pi_t^2 - \frac{1}{2} \left(\frac{1}{\gamma y} + \frac{\rho w}{y} \right) \hat{\Sigma}_t^2 - \frac{1}{2} w \left(1 + \frac{y}{\rho} \right) \left(\hat{y}_t^2 - 2 \frac{1 + \gamma \rho}{y + \rho} \hat{y}_t^n \hat{y}_t \right) \\ &\quad - (w - 1) \hat{y}_t - \left(\frac{1}{\gamma y} + \frac{\rho w}{y} \right) \hat{\Sigma}_t - (w - 1) \hat{y}_t \hat{\Sigma}_t + w \hat{z}_t \hat{\Sigma}_t \\ &\quad - (w - 1) \hat{y}_t \hat{v}_t - \left(\frac{1}{\gamma y} + \rho \frac{w}{y} \right) \hat{\Sigma}_t \hat{v}_t + t.i.p. \end{aligned} \quad (63)$$

This approximation can be rewritten in the standard form to isolate the welfare cost of market incompleteness and nominal rigidity. To do so, let $\tilde{y}_t = \hat{y}_t - \hat{y}_t^n$, then the approximation becomes:

$$\begin{aligned}
L_t = \frac{U_t - U}{e^{-\gamma y \Sigma y}} \approx & -\frac{1}{2} w \Phi \pi_t^2 - \frac{1}{2} w \left(1 + \frac{y}{\rho}\right) \frac{1 + \gamma \rho}{y + \rho} \tilde{y}_t^2 \\
& - \frac{1}{2} \left(\frac{1}{\gamma y} + \frac{\rho w}{y}\right) \hat{\Sigma}_t^2 + w \hat{z}_t \hat{\Sigma}_t - \left(\frac{1}{\gamma y} + \rho \frac{w}{y}\right) \hat{\Sigma}_t \hat{v}_t - \left(\frac{1}{\gamma y} + \frac{\rho w}{y}\right) \hat{\Sigma}_t \\
& - (w - 1) \hat{y}_t \hat{v}_t - (w - 1) \hat{y}_t \hat{\Sigma}_t - \frac{1}{2} \left(\gamma y - 1 - w \frac{(1 - y + \rho(\gamma - 1))}{\rho}\right) \hat{y}_t^2 - (w - 1) \hat{y}_t
\end{aligned} \tag{64}$$

The first two terms contain the welfare loss due to price rigidities and coincides exactly with the welfare loss in the RANK model, as given by $L_t = -\frac{1}{2} w \Phi \pi_t^2 - \frac{1}{2} w \left(1 + \frac{y}{\rho}\right) \frac{1 + \gamma \rho}{y + \rho} \tilde{y}_t^2$. This expression says that nominal rigidities create distortions that can be measured in terms of inflation and inefficient deviations of output from its natural rate.

The next four terms, $-\frac{1}{2} \left(\frac{1}{\gamma y} + \frac{\rho w}{y}\right) \hat{\Sigma}_t^2 + w \hat{z}_t \hat{\Sigma}_t - \left(\frac{1}{\gamma y} + \rho \frac{w}{y}\right) \hat{\Sigma}_t \hat{v}_t - \left(\frac{1}{\gamma y} + \frac{\rho w}{y}\right) \hat{\Sigma}_t$, contain the welfare loss from consumption dispersion due to incomplete markets. The first three terms, $-\frac{1}{2} \left(\frac{1}{\gamma y} + \frac{\rho w}{y}\right) \hat{\Sigma}_t^2 + w \hat{z}_t \hat{\Sigma}_t - \left(\frac{1}{\gamma y} + \rho \frac{w}{y}\right) \hat{\Sigma}_t \hat{v}_t$, capture the cost of inequality deviating from its steady-state, net of technology shocks and cost-push shocks. Even if we assume the steady-state of consumption dispersion to be 1, i.e., starting value of no inequality, this term is still active. The fourth term, $-\left(\frac{1}{\gamma y} + \frac{\rho w}{y}\right) \hat{\Sigma}_t$, shows that there is intrinsic value to the direction of consumption dispersion. In other words, the level of inequality following a shock also matters. An increase in the level of inequality directly generates a welfare loss, while a decline in inequality directly lowers welfare loss (the weight here is a negative number). Note, however, that the expected value of this term is zero, as inequality is assumed to be at its steady-state in the long-term.

The remaining terms, $-(w - 1) \hat{y}_t \hat{v}_t - (w - 1) \hat{y}_t \hat{\Sigma}_t - \frac{1}{2} \left(\gamma y - 1 - w \frac{(1 - y + \rho(\gamma - 1))}{\rho}\right) \hat{y}_t^2 - (w - 1) \hat{y}_t$, are active only for the HANK optimal steady-state. In that steady-state, the planner internalizes that the starting value is not a world with zero consumption dispersion, but rather a world where consumption inequality already exists. In that case, the planner realizes that a steady-state output that is higher than the natural rate of output can result in lower consumption dispersion, and hence lower societal cost in the long run. In order to mitigate the cost from incomplete financial markets, the planner accepts to deviate from the efficient allocation.

The social welfare loss expressed in term of steady state is given by:

$$\mathbb{W}_0 = -\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t L_t \tag{65}$$

The average per-period welfare loss is the expectation of the per-period welfare loss, given by $\mathbb{L}_0 = -\mathbb{E}_0(L_t)$. The processes \hat{y}_t , \hat{z}_t , \hat{v}_t , and $\hat{\Sigma}_t$ are mean zero processes given that they are deviations from their steady state values. Using this insight, we have:

$$\begin{aligned} \mathbb{L}_0 = & \frac{1}{2} \left\{ w\Phi \text{Var}(\pi_t) + \left(\frac{1}{\gamma y} + \frac{\rho w}{y} \right) \text{Var}(\hat{\Sigma}_t) + w \left(1 + \frac{y}{\rho} \right) \frac{1 + \gamma \rho}{y + \rho} \text{Var}(\tilde{y}_t) \right. \\ & + \left(\gamma y - 1 - w \frac{(1 - y + \rho(\gamma - 1))}{\rho} \right) \text{Var}(\hat{y}_t) \\ & \left. + 2(w - 1) \text{Cov}(\hat{y}_t, \hat{\Sigma}_t) - 2w \text{Cov}(\hat{z}_t, \hat{\Sigma}_t) + 2(w - 1) \text{Cov}(\hat{y}_t, \hat{v}_t) + 2 \left(\frac{1}{\gamma y} + \rho \frac{w}{y} \right) \text{Cov}(\hat{\Sigma}_t, \hat{v}_t) \right\}. \end{aligned} \quad (66)$$

The normalized welfare is given by:

$$\begin{aligned} \mathbb{L}_0 = & \text{Var}(\pi_t) + \frac{1}{w\Phi} \left(\frac{1}{\gamma y} + \frac{\rho w}{y} \right) \text{Var}(\hat{\Sigma}_t) + \frac{1}{\Phi} \left(1 + \frac{y}{\rho} \right) \frac{1 + \gamma \rho}{y + \rho} \text{Var}(\tilde{y}_t) \\ & + \frac{1}{w\Phi} \left(\gamma y - 1 - w \frac{(1 - y + \rho(\gamma - 1))}{\rho} \right) \text{Var}(\hat{y}_t) \\ & + 2 \frac{w - 1}{w\Phi} \text{Cov}(\hat{y}_t, \hat{\Sigma}_t) - 2 \frac{1}{\Phi} \text{Cov}(\hat{z}_t, \hat{\Sigma}_t) \\ & + 2 \frac{w - 1}{w\Phi} \text{Cov}(\hat{y}_t, \hat{v}_t) + 2 \frac{1}{w\Phi} \left(\frac{1}{\gamma y} + \rho \frac{w}{y} \right) \text{Cov}(\hat{\Sigma}_t, \hat{v}_t). \end{aligned} \quad (67)$$

If we assume the RANK efficient steady-state, we have $y = n = w = 1$ and the expected loss becomes:

$$\begin{aligned} \mathbb{L}_0 = & \text{Var}(\pi_t) + \frac{1}{\Phi} \left(\frac{1}{\gamma} + \rho \right) \text{Var}(\hat{\Sigma}_t) + \frac{1}{\Phi} \frac{1 + \gamma \rho}{\rho} \text{Var}(\tilde{y}_t) - \frac{2}{\Phi} \text{Cov}(\hat{z}_t, \hat{\Sigma}_t) \\ & + 2 \frac{1}{w\Phi} \left(\frac{1}{\gamma y} + \rho \frac{w}{y} \right) \text{Cov}(\hat{\Sigma}_t, \hat{v}_t). \end{aligned} \quad (68)$$

Proof of proposition 2. As shown in appendix B.1 of Acharya et al. (2021), the cross-sectional distribution of consumption is normally distributed with mean y_t and variance:

$$\sigma_{c,s,t}^2 = \mu_t^2 w_t^2 \sigma_t^2 + \mu_t^2 \sigma_{a,s,t}^2. \quad (69)$$

while the asset holdings follows:

$$\sigma_{a,s,t}^2 = \frac{R_{t-1}^2}{\vartheta^2} \left(1 - (1 + \gamma \rho w_{t-1}) \mu_{t-1} \right)^2 \left(\sigma_{a,s,t-1}^2 + w_{t-1}^2 \sigma_{t-1}^2 \right). \quad (70)$$

Combining the previous two equations gives:

$$\begin{aligned}
\sigma_{c,s,t}^2 &= \mu_t^2 w_t^2 \sigma_t^2 + \mu_t^2 \sigma_{a,s,t}^2 \\
&= \mu_t^2 w_t^2 \sigma_t^2 + \mu_t^2 \frac{R_{t-1}^2}{\vartheta^2} \left(1 - (1 + \gamma \rho w_{t-1}) \mu_{t-1} \right)^2 \left(\sigma_{a,s,t-1}^2 + w_{t-1}^2 \sigma_{t-1}^2 \right) \\
&= \mu_t^2 w_t^2 \sigma_t^2 + \frac{\mu_t^2}{\mu_{t-1}^2} \frac{R_{t-1}^2}{\vartheta^2} \left(1 - (1 + \gamma \rho w_{t-1}) \mu_{t-1} \right)^2 \left(\mu_{t-1}^2 \sigma_{a,s,t-1}^2 + \mu_{t-1}^2 w_{t-1}^2 \sigma_{t-1}^2 \right) \\
&= \mu_t^2 w_t^2 \sigma_t^2 + \frac{\mu_t^2}{\mu_{t-1}^2} \frac{R_{t-1}^2}{\vartheta^2} \left(1 - (1 + \gamma \rho w_{t-1}) \mu_{t-1} \right)^2 \sigma_{c,s,t-1}^2 \\
&= \mu_t^2 w_t^2 \sigma_t^2 + \sigma_{c,s,t-1}^2
\end{aligned} \tag{71}$$

The aggregate consumption volatility is:

$$\begin{aligned}
\sigma_{c,t}^2 &= (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \sigma_{c,s,t}^2 \\
&= (1 - \vartheta) \sigma_{c,t,t}^2 + (1 - \vartheta) \sum_{s=-\infty}^{t-1} \vartheta^{t-s} \sigma_{c,s,t}^2 \\
&= (1 - \vartheta) \mu_t^2 w_t^2 \sigma_t^2 + (1 - \vartheta) \sum_{s=-\infty}^{t-1} \vartheta^{t-s} \sigma_{c,s,t}^2 \\
&= (1 - \vartheta) \mu_t^2 w_t^2 \sigma_t^2 + (1 - \vartheta) \vartheta \sum_{s=-\infty}^{t-1} \vartheta^{t-1-s} \sigma_{c,s,t}^2 \\
&= (1 - \vartheta) \mu_t^2 w_t^2 \sigma_t^2 + (1 - \vartheta) \vartheta \sum_{s=-\infty}^{t-1} \vartheta^{t-1-s} \left(\sigma_{c,s,t-1}^2 + \mu_{t-1}^2 w_{t-1}^2 \sigma_{t-1}^2 \right) \\
&= (1 - \vartheta) \mu_t^2 w_t^2 \sigma_t^2 + \vartheta (\sigma_{c,t-1}^2 + \mu_{t-1}^2 w_{t-1}^2 \sigma_{t-1}^2) \\
&= \mu_t^2 w_t^2 \sigma_t^2 + \vartheta \sigma_{c,t-1}^2
\end{aligned} \tag{72}$$

To summarize:

$$\begin{aligned}
\sigma_{c,t}^2 &= \mu_t^2 w_t^2 \sigma_t^2 + \vartheta \sigma_{c,t-1}^2 \\
\sigma_c^2 &= \frac{\mu^2 w^2 \sigma^2}{1 - \vartheta} = \frac{\Lambda}{\gamma^2 (1 - \vartheta)} \\
\hat{\sigma}_{c,t} &= (1 - \vartheta) (\hat{\mu}_t + \hat{w}_t + \hat{\sigma}_t) + \vartheta \hat{\sigma}_{c,t-1} \\
&= (1 - \vartheta) (\hat{\mu}_t + \phi y \hat{y}_t) + \vartheta \hat{\sigma}_{c,t-1}
\end{aligned} \tag{73}$$

D Empirical evidence of the cyclical of earnings risk

Empirical evidence suggests that earnings risk is countercyclical both for the U.S. and for Canada. However, the evidence is divided in terms of which moment captures the countercyclical of earnings risk.

For the U.S., Storesletten et al. (2004) show that the first-moment of idiosyncratic labour earnings is strongly counter-cyclical in survey-based panel data (PSID). They find that the cross-sectional standard deviation of earnings increases by 75 percent as the U.S. economy moves from peak to trough. Guvenen et al. (2015) use administrative tax-filer data in the U.S. to demonstrate that while the variance of earnings is acyclical, the skewness of earnings is strongly counter-cyclical. Nakajima and Smirnyagin (2019) re-examine PSID data and find that both the variance and skewness in PSID exhibit counter-cyclical idiosyncratic risk. In particular, they find that when GNP per-capita growth decreases by 3.7 percent (peak to trough in a business cycle), the standard deviation of labour income increases by 0.1 (from 0.1 to 0.2).

Earnings risk is counter-cyclical for Canada as well. Using a variety of sources, including the Survey of Labour Income Dynamics (SLID), Brzozowski et al. (2010) find that during recessions, wage and income inequality rises substantially in Canada.

In what follows, we estimate possible values for ϕ in Canada using three distinct methods to infer relevant aggregate moments: (a) aggregate moments from Canadian tax-filer data (Longitudinal Administrative Database—LAD), (b) aggregate moments from survey-based panel data (Survey of Labour Income Dynamics—SLID), and (c) the Labour Force Survey (LFS) job-loss dynamics.

D.1 Longitudinal Administrative Database

First, we consider a time-series of unconditional annual moments from Canadian tax-filer data from the Longitudinal Administrative Database (LAD). The LAD database is a panel comprising a 20 percent sample of annual tax filings between 1982 and 2016. Following the influential work of Guvenen et al. (2015), we only consider earnings data for working-age males earning above minimum wage in the previous year to ensure sufficient ties to the labour market. Karibzhanov (2020) show that key moments from the LAD database are comparable to those obtained from tax filer data for the U.S.

We measure earnings risk using the variance of log earnings, variance of earnings growth, and log variance of log earnings. We run regressions on any of three explanatory variables: (i) mean log earnings from the LAD database, (ii) real GDP—both de-trended using an HP filter²², and (iii) GDP growth. The regressions also include a constant and a time trend. The regressions are run on annual data spanning from 1983 to 2016. Table 12 summarizes the findings. One, two, and three asterisks denote statistical significance at the 10 percent, 5 percent, and 1 percent level, respectively. The corresponding values for ϕ range from -0.47 to -2.6.

D.2 Survey of Labour Income Dynamics

Brzozowski et al. (2010) estimate a persistent idiosyncratic income process using the Survey of Longitudinal Income Database (SLID) for 1993-2005. They report the cross-

²²Following Ravn and Sterk (2016), we adopt a value of $\lambda = 6.25$ for annual data.

Table 12: Cyclicalities of earnings risk based on LAD

LHS variable	Mean income(LAD)	Output gap	Output growth
Variance of log earnings	-0.99**	-1.82*	-1.37
R-squared	0.78	0.77	0.73
Variance of earnings growth	-0.47***	-0.63*	-1.16***
R-squared	0.52	0.35	0.46
Log variance of log earnings	-1.42**	-2.62*	-2.12
R-squared	0.78	0.76	0.73

Table 13: Cyclicalities of earnings risk based on SLID

LHS variable:	Output gap	Output growth
Variance of Individual wage income	-0.17	-1.28**
R-squared	0.06	0.46
Log variance of Individual wage income	-3.1	-23.43
R-squared	0.06	0.47

sectional variance of both permanent and transitory components of their income processes, controlling for age, gender, marital status, education, province of residence, immigration status, and mother tongue for individual wage earners. We run regressions on their reported cross-sectional variance on Canadian output (de-trended using an HP filter) and output growth, including a constant and time trend. Table 13 summarizes the findings:

D.3 Inferring cyclicalities of income risk from labour market moments

From an individual worker's perspective, job losses are the most salient risk to earnings. In this sub-section, we lean on the literature for estimates of earnings and consumption-based losses from losing one's job and infer the variance of earnings based on the cyclicalities of job-market dynamics.

Browning and Crossley (2008) measure the cost of job loss in consumption terms using the Canadian Out of Employment Panel (COEP)—a panel survey collected by HRDC. They find that permanent layoffs are associated with a consumption decline of 6.4 percent.

Now, consider a simple Bernoulli experiment with probability of job loss p_t , and a consequent decline in consumption (or earnings) proportional to a . A worker keeps their job next period with probability $(1 - p_t)$ and retains their earnings/consumption normalized to 1. With probability p_t , they lose their job and earns/consumes $(1 - a)$. In this case, their expected earning will be: $E[y_t^i] = p_t(1 - a) + (1 - p_t) = 1 - ap_t$.

Table 14: Cyclicalty of earnings risk based on LFS

LHS variable	Output gap	R-squared
Variance	-0.035***	0.83
Log variance	-7.38***	0.85

The variance of their earnings/consumption in any period is given by: $Var(y_t^i) = p_t[(1 - a) - (1 - ap_t)]^2 + (1 - p_t)[1 - (1 - ap_t)]^2 = a^2 p_t(1 - p_t)$.

If we let $p_t(1 - p_t) = f(\tilde{y}_t) = c\tilde{y}_t$ be a linear function of the output gap, then we can estimate $\phi = \frac{\partial var(y_t^i)}{\partial \tilde{y}_t} = a^2 f'(\tilde{y}_t) = a^2 c$.

We calculate the probability of a permanent layoff, p_t , from LFS data as the ratio of people who are unemployed or out of the labour force in period t due to a permanent layoff within the last year, over the number of employed people in the period $t - 1$. We then infer a series for income risk following Browning and Crossley (2008), by setting $a = 0.064$, to get $var(y_t^i) = (0.064)^2 p_t(1 - p_t)$. We then run regressions for the variance and log-variance of imputed earnings risk on the Canadian output gap (detrended using an HP filter), including a constant and time trend. Table 14 summarizes the results. Again, we see that earnings risk is counter-cyclical, with values of ϕ between -0.04 and -7.4 .

E Sensitivity analysis for the cyclicalty of income risk

The level and cyclicalty of idiosyncratic risk affects the transmission mechanism in the HANK framework. In this appendix, we explain this mechanism by contrasting three cases of the HANK model where idiosyncratic risk is (a) countercyclical, (b) procyclical, or (c) acyclical. We then provide sensitivity analysis for the level and degree of cyclicalty for the empirically relevant case of counter-cyclical risk.

Figure 6 shows impulse responses for demand and cost-push shocks across different cyclicalities of idiosyncratic risk for an Inflation Targeting framework, where the monetary authority follows the policy rule $i_t = 2.5\pi_t + 0.25\tilde{y}_t$. Alongside the benchmark countercyclical case with $\phi = -2.62$, we consider a procyclical case of $\phi = 7$ and the acyclical case of $\phi = 0$, and compare them to the RANK case. We keep all other parameters the same, including $\sigma = 0.74$.

The first column shows impulse responses for a persistent positive demand shock. Under countercyclical risk, the two channels of precautionary savings go in opposite directions. A rise in expected future output due to the positive demand shock directly lowers the cross-sectional dispersion in income, *decreasing* precautionary savings through the *income risk* channel. The accompanying rise in policy rates, however, make it harder for agents in the economy to partially self-insure against idiosyncratic

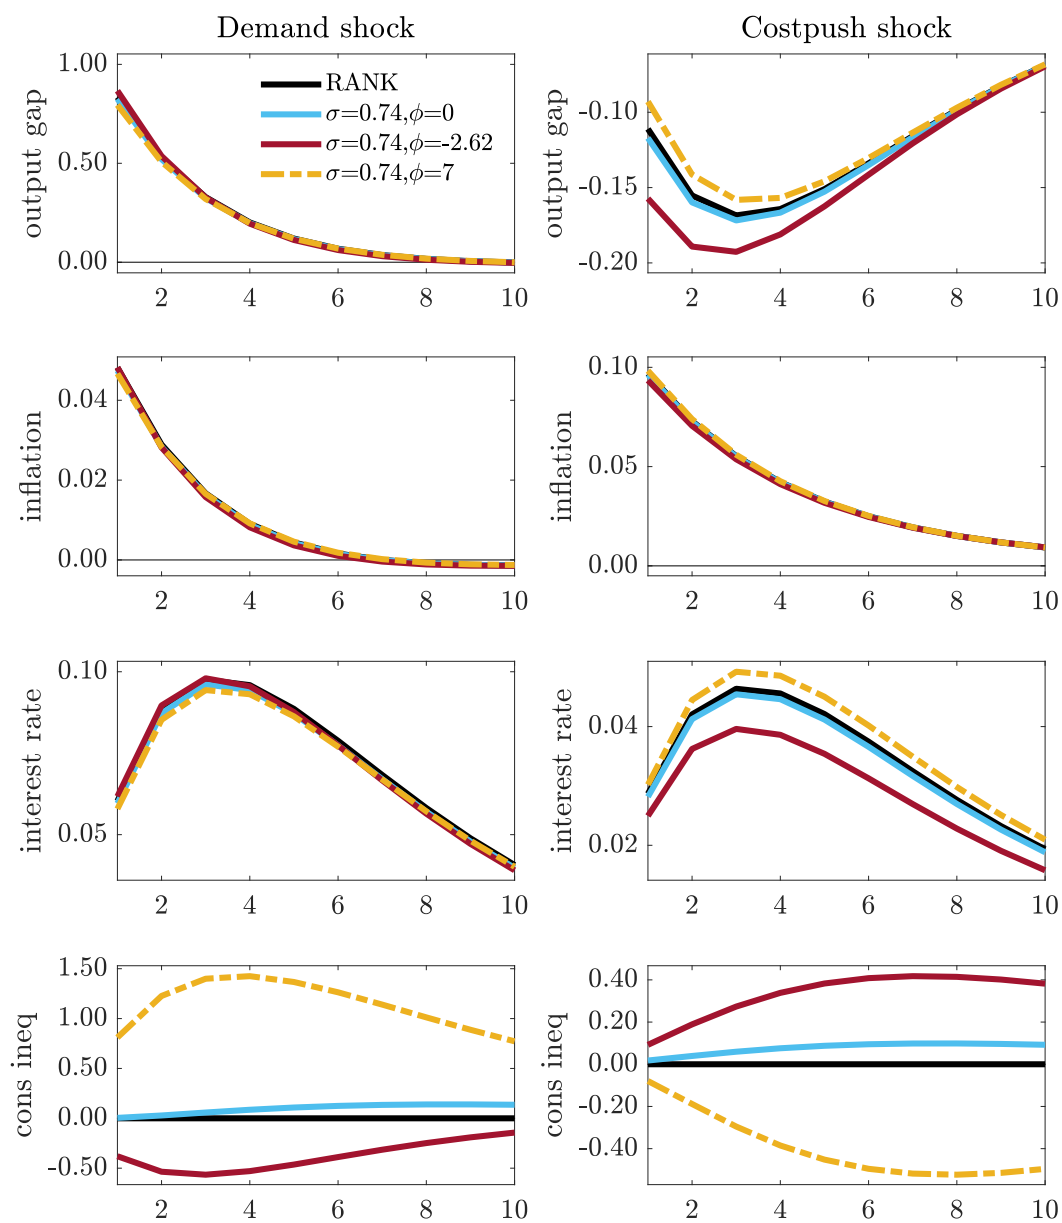


Figure 6: Impulse responses to demand and cost-push shocks across different cyclicalities of risk

risk using the aggregate savings vehicle. As a result, the *self-insurance* channel implies an *increase* in precautionary savings. In the baseline calibration, the income risk channel dominates, and lower precautionary savings lead an amplification of output, and, through the NKPC, higher inflation vis-à-vis the RANK case.

Under procyclical risk, the rise in expected future output from the demand shock *increases* precautionary savings through the *income risk channel* by directly increasing income risk. Both channels now imply an increase in precautionary savings, putting downward pressure on output, resulting in a *dampening* of output and inflation vis-à-vis the RANK case. In the acyclical case, the *income risk channel* is absent, but higher rates still imply an increase in precautionary savings through the *self-insurance channel*, which ends up marginally *dampening* the cycle.

Cross-sectional dispersion in consumption and its welfare cost are also affected by the same channels. In the countercyclical (procyclical) case, a rise in output directly lowers (increases) the dispersion in earnings and consumption. This is mitigated (amplified) by the fact that higher rates imply higher pass-through from income risk to consumption risk.

The second column shows responses for a persistent cost-push shock, modelled as a shock to the NKPC. This time, the two channels of precautionary savings go in the same direction for the countercyclical case, and in opposite directions in the procyclical case. When $\phi < 0$, the decline in output following a cost-push shock leads to higher *income risk* and higher precautionary savings, *amplifying* the fall in output. At the same time, the rise in rates result in higher precautionary savings through the *self-insurance* channel, further amplifying the downturn. The combined effect of lower output weighs down on inflation via the NKPC. As a result, both output and inflation are lower in the HANK model vis-à-vis the RANK case.

In the procyclical case, the decline in output leads to lower *income risk* and lower precautionary savings, dampening the fall. The rise in rates still leads to a rise in precautionary savings through the *self-insurance* channel, partially counteracting the dampening effect. The acyclical case only features the self-insurance channel, which increases precautionary savings and marginally dampens the recession brought forth by central bank action.

Consumption inequality follows the direction of aggregate precautionary savings, rising in the countercyclical case and falling in the procyclical case. It's important to note that in the countercyclical case, central banks now face an additional trade-off. In order to bring inflation back to target following a cost-push shock, central banks now not only have to generate a recession, they also have to accept higher inequality as a consequence of that recession. Even when income risk is acyclical, inflation stabilization leads to an increase in consumption inequality via the self-insurance channel of precautionary savings.

The two key parameters that determine transmission in the HANK model are the steady-state earnings risk, σ , and the cyclical risk, ϕ . Equation 15 suggests that

higher values of σ enhance the transmission through both channels of precautionary savings, while higher (absolute) values of ϕ enhance the income risk channel.

Figure 7 and 8 shows the transmission mechanism of the HANK model is enhanced with larger (absolute) values of σ and ϕ , respectively, when the remaining model parameters at their benchmark values.

Kaplan et al. (2018) measure the degree of amplification in a HANK framework by looking at the interest rate elasticity of consumption for a monetary policy shock. They find that their calibration provides an interest elasticity of consumption in the HANK model that is 50 percent higher than a similarly calibrated RANK case. Figures 9 and 10 show the interest elasticity of consumption in our HANK framework vis-à-vis the embedded RANK case across different values of σ and ϕ , respectively. We calculate the elasticity of initial consumption to the initial change in interest rates $\frac{d \log(c_0)}{dr_0}$ as the ratio of the period 1 impulse response of consumption (which is equal to output in our case) and the period 1 impulse response of the policy rate following a monetary policy shock. In this exercise, we follow Kaplan et al. (2018) and set the monetary policy shock persistence to 0.61.

As the figures show, increasing the (absolute) values of σ and ϕ increases the interest elasticity of consumption in the HANK case vis-à-vis the RANK case. The benchmark choice of $\sigma = 0.74$ and $\phi = -2.62$ generates an elasticity of initial consumption to the initial change in interest rate that is 15.5 percent higher than the embedded RANK model. Increasing the parameter value of σ to 0.785 would increase the interest elasticity to 35.5 percent higher than the RANK case. However, the amplification properties of the tractable HANK model are still less than the reported 50 percent value reported in the HANK model of Kaplan et al. (2018).

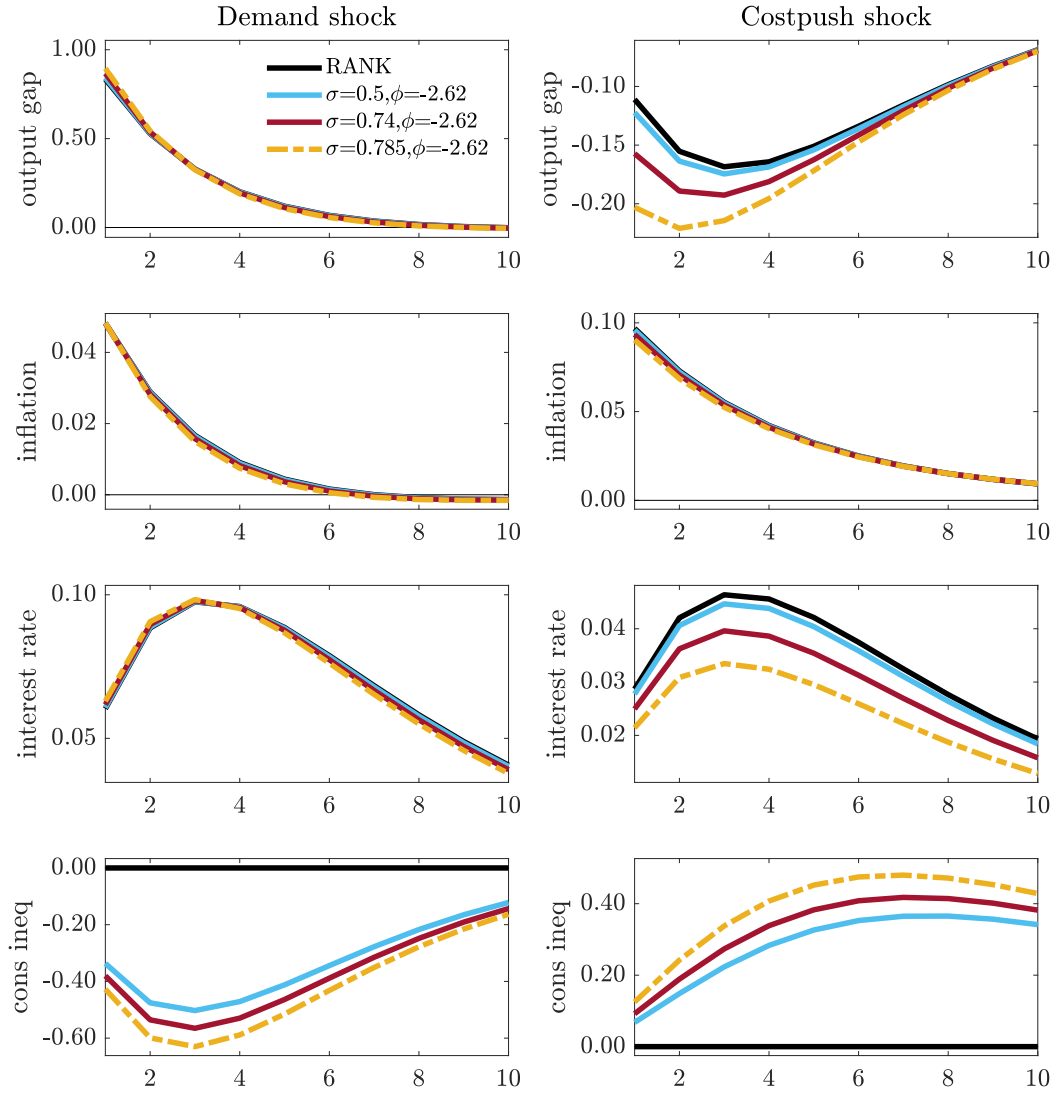


Figure 7: Impulse responses to demand and cost-push shocks across different values of σ

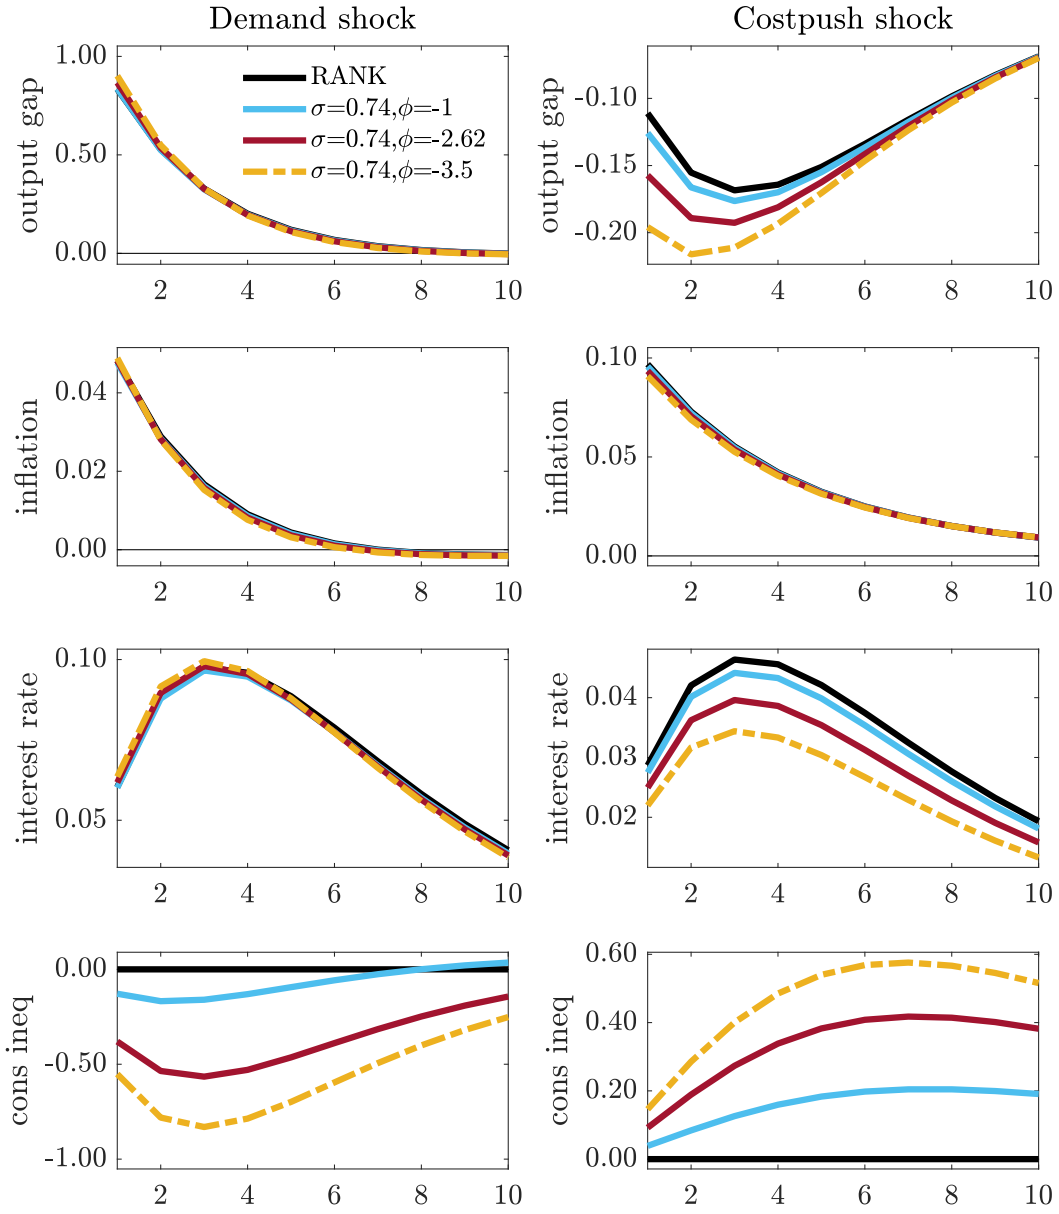


Figure 8: Impulse responses to demand and cost-push shocks across different values of ϕ

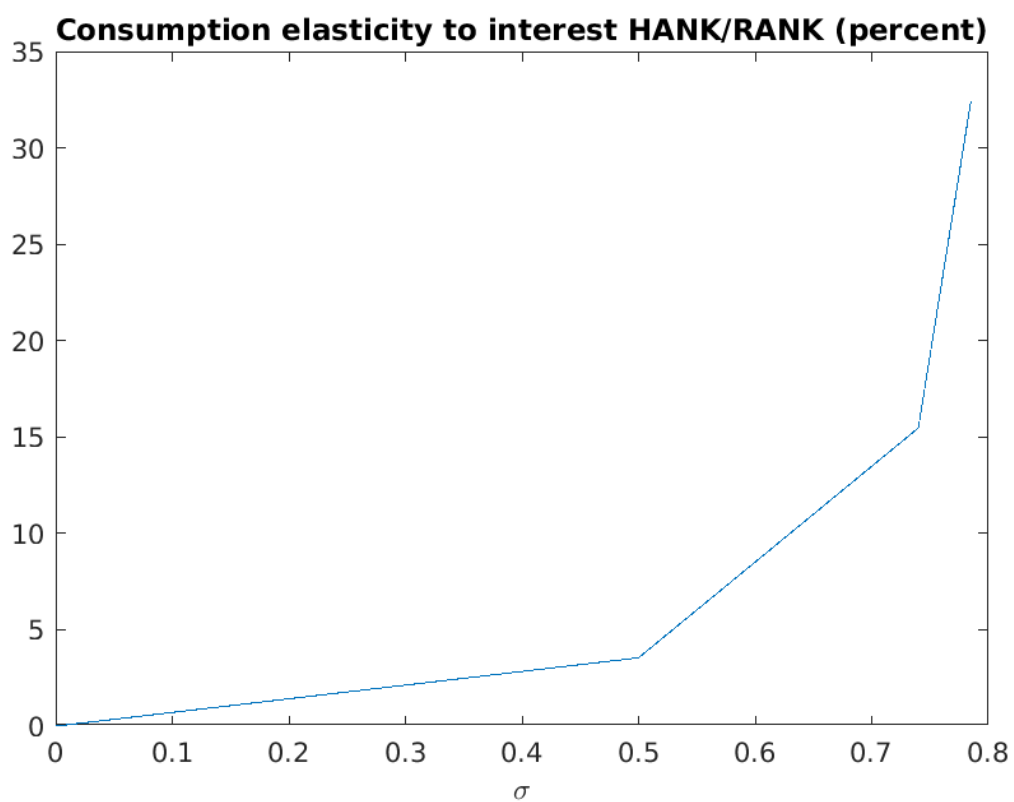


Figure 9: Interest elasticity of output compared to RANK (percent increase) across different values of σ




Figure 10: Interest elasticity of output compared to RANK (percent increase) across different values of ϕ