

Optimal Unconventional Policy in a New Keynesian Preferred Habitat Model*

Rupal Kamdar

Indiana University, Bloomington

Walker Ray

Federal Reserve Bank of Chicago

& CEPR

September 2026 (*latest version*)

Abstract

We study conventional and balance sheet policy in a general equilibrium model with heterogeneous households, nominal rigidities, and limits to arbitrage. Both instruments move risk prices through portfolio rebalancing; because household portfolios differ in duration, the induced term premium fluctuations lower welfare through imperfect risk-sharing. Conservative calibrations imply these welfare costs are comparable to or larger than those of traditional stabilization motives. When divine coincidence holds and both instruments are unconstrained, optimal policy satisfies a separation principle: conventional policy stabilizes output and inflation, while balance sheet policy stabilizes term premia. If the short rate is constrained, balance sheet policy is a necessary but imperfect tool for macroeconomic stabilization; with costly balance sheet tools, a less reactive short rate improves welfare at the cost of greater macroeconomic volatility. Under commitment, the path of the short rate can mitigate households' exposure to aggregate risk. Our results justify a central bank "triple mandate" that includes term premia alongside output and inflation, and challenge the view that balance sheet policy is a close substitute for the short rate.

JEL Classification: E52, E44, E43, D14

*rkamdar@iu.edu, walkerdray@gmail.com. This paper updates and replaces "Monetary Policy and the Limits to Arbitrage: Insights from a New Keynesian Preferred Habitat Model" ([Ray, 2019](#)). We thank Yuriy Gorodnichenko, Pierre-Olivier Gourinchas, David Romer, Michael Bauer, Gadi Barlevy, Markus Brunnermeier, Max Croce, Ian Dew-Becker, Mike Droste, Andrew Foerster, Ben Hébert, Erik Hurst, Anil Kashyap, Rohan Kekre, Peter Kondor, Jennifer La'O, Byoungchan Lee, Andy Levin, Karen Lewis, Dima Mukhin, Nick Sander, Dmitriy Sergeyev, Jón Steinsson, Mauricio Ulate, Dimitri Vayanos, Cynthia Wu, and participants at Berkeley, Bocconi, Chicago Booth, Cleveland Fed, Columbia GSB, CREI, CUHK, Dallas Fed, Dartmouth, Fed Board, HKU, HKUST, INSEAD, LSE, Marquette, Notre Dame, Oxford, Rice, SF Fed, St. Louis Fed, Toulouse, Wharton, UIUC, USC, CEBRA, CEPR AP, CESifo, and SED for helpful comments. The views expressed do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

1 Introduction

Since the global financial crisis, central bank balance sheets have become an essential tool of monetary policy. With the federal funds rate at its lower bound, the Federal Reserve used quantitative easing (QE) to lower long-term rates, expanded its balance sheet to about \$9 trillion after the onset of COVID-19, and then began quantitative tightening (QT) in 2022 while raising its policy rate. Yet the consensus view among policymakers remains that the short rate is the primary tool for adjusting the stance of monetary policy; the balance sheet is an effective substitute when the short rate is constrained, but at other times should only adjust gradually and predictably in the background.¹ This paper studies how substitutable the short rate and balance sheet policies are, whether balance sheet policy has a role beyond substituting for the short rate, and how the two instruments should be designed together.

We argue that the consensus view overlooks the welfare consequences of both instruments when risk prices are endogenous. In particular, QE is an imperfect substitute for the short rate, and is instead better suited to dampening inefficient fluctuations in risk prices. We develop a tractable New Keynesian model in which households and firms borrow and save in a segmented, preferred habitat bond market with imperfect arbitrage. Balance sheet operations move the market price of risk directly, while the short rate moves it indirectly through changes in the bond demand of households and firms. Because households hold portfolios of different maturities, fluctuations in risk prices and term premia create dispersion in borrowing rates and consumption across households; with imperfect risk-sharing, this dispersion lowers welfare. Combining household balance sheet data with estimates of term premia, our model implies this cost is comparable to or larger than the welfare costs of output and inflation volatility. Hence, even when the short rate is unconstrained and “divine coincidence” holds, a central bank which targets only output and inflation, with either instrument, does not achieve the first-best.

Optimal policy assigns each instrument to the objective where it has a comparative advantage. When neither instrument is costly and divine coincidence holds, a *separation principle* applies: the short rate stabilizes output and inflation, while the balance sheet stabilizes term premia. Away from this benchmark, optimal policy trades off macroeconomic stabilization against risk-sharing. When the short rate is constrained, the balance sheet must also stabilize the macroeconomy, but using QE to fully stabilize output worsens risk-sharing, and optimal policy stops short of full stabilization. Conversely, when balance sheet policy is costly, the optimal short rate rule responds less to aggregate shocks in order to dampen term premium fluctuations. Under commitment, the central bank can also reduce the exposure of

¹See, for example, [Federal Open Market Committee \(2022\)](#).

household portfolios to aggregate risk through the path of the short rate.

Our model starts from a conventional New Keynesian model with nominal rigidities. Building on [Vayanos and Vila \(2021\)](#), households and firms have differentiated access to bonds of different maturities, and the positions they take induce “preferred habitat” segmentation across the term structure. Specialized bond arbitrageurs re-integrate markets, but only partially when their risk-bearing capacity is imperfect. Both conventional and unconventional policy therefore move term premia: the policy rate is not transmitted one-for-one to all borrowing rates, pass-through differs across household portfolios, and QE moves borrowing rates even when the policy rate is unchanged.

Through market clearing, the bond positions of heterogeneous households matter for term premia, and movements in term premia in turn shift relative bond positions across the entire distribution of households. Nevertheless, we characterize the dynamics of the model with a small number of sufficient statistics of the household asset distribution. Intuitively, term premia depend on the equilibrium exposure of bond arbitrageurs to aggregate risk. Market clearing requires the entire distribution of bond positions to pin down arbitrageur holdings; however, only the *risk-weighted distribution* across the term structure matters for the pricing of risk. Because there are only a small number of aggregate risk factors, the dynamics of the risk-weighted distribution are summarized by a small number of factors.

The rebalancing mechanism implies that the comovement of term premia with the short rate depends on the horizon and on why the short rate moves. When a recessionary shock lowers firm borrowing, term premia fall with the short rate on impact, but slow-moving household portfolio adjustments reverse this pattern over longer horizons; the comovement on impact is stronger when the central bank also uses its balance sheet, and weaker or even of the opposite sign following monetary policy shocks. These predictions help rationalize findings about the response of Treasury forward rates to short rates (see, e.g., [Hanson and Stein 2015](#), [Hanson et al. 2021](#), [Gürkaynak et al. 2005](#), [Nakamura and Steinsson 2018](#)).

For aggregate demand, conventional and unconventional policies are substitutes: conditional on the paths of the household and firm “effective rates,” which are weighted averages of bond returns across the term structure, the first-order dynamics of output and inflation are those of a representative agent model in which households and firms borrow at these rates. When prices are rigid, or when firms and households face the same effective rate, a central bank whose loss function depends only on output and inflation is therefore indifferent between the two tools, and if divine coincidence holds, either tool alone can fully stabilize output and inflation. Such a loss function omits the welfare cost of term premium fluctuations. Data from the Survey of Consumer Finances (SCF) and estimates of term premia put this cost at 0.4 to 1.7 percent of steady-state consumption over 1990-2019, compared with

0.05 to 0.33 percent for output and inflation volatility.

Beyond simple rules, in which the policy tools respond only to the current state of the economy, a central bank with full commitment can choose its policy tools as functions of the entire history of shocks, which improves welfare. As in standard New Keynesian models, commitment introduces promise-keeping, which makes policy inertial: the central bank stabilizes output and inflation through promises about future policy rather than through large movements in current rates. In our model, when moving the short rate is costly, inertia has an additional benefit: less reactive policy implies *ex ante* smaller expected variation in the policy tools, which reduces bond volatility and the risk compensation arbitrageurs require, and therefore improves risk-sharing directly.

Commitment also gives rise to a *risk-shaping* motive. Because bond prices reflect the expected path of policy over the life of each bond, the exposure of household portfolios to aggregate risk is itself an equilibrium outcome. The central bank can influence this exposure by committing to the shape of its responses to shocks, even when the market price of risk varies. This motive is quantitatively important. In a calibrated version of the model without balance sheet tools, optimal commitment policy removes more than 80 percent of the welfare loss under a historical Taylor rule; a central bank which commits but does not internalize the risk-shaping motive (and so takes portfolio exposures as given) removes only about two thirds, and is left with more than twice the loss. When balance sheet tools are cheap, the central bank can stabilize risk prices directly, and risk-shaping adds little.²

Finally, our results show that optimal policy pursues a “triple mandate” of stable prices, output, and term premia. The statutory mandate of the Federal Reserve is to “promote effectively the goals of maximum employment, stable prices, and moderate long-term interest rates.” Academics, practitioners, and policymakers typically interpret this mandate in terms of its first two goals, reasoning that macroeconomic stability delivers moderate long-term yields. In our model, macroeconomic stabilization is not always sufficient to stabilize term premia, which gives the third goal a welfare rationale of its own.³

Our paper builds on the preferred habitat model of [Vayanos and Vila \(2021\)](#), which formalizes the concept of [Modigliani and Sutch \(1966\)](#). In preferred habitat models, demand

²Our calibration implies that the risk-shaping motive is powerful. However, the central bank is required to know the structure of household portfolio dispersion, and to condition policy on the history of shocks. Simple balance sheet rules which reduce term premium volatility achieve most of the same gains and can be implemented as functions of observables alone. Relatedly, a fiscal authority with a full set of state-contingent, household-specific transfers could remove consumption dispersion directly (as in any model with heterogeneous agents), but such transfers must be individual-specific and require even more knowledge of each household’s portfolio; simple balance sheet tools only require transactions in secondary bond markets.

³See [Mishkin \(2007\)](#) for the standard interpretation. Our argument interprets “moderate long-term interest rates” as stable term premia; our model is silent on fiscal motives for the level of long-term rates.

and supply shocks from clientele investors change the risk exposure of marginal bond investors; when their risk-bearing capacity is imperfect, risk premia fluctuate, and the term structure departs from the expectations hypothesis.⁴ In this tradition the short rate is exogenous. General equilibrium models with segmented bond markets have generated term premia through portfolio adjustment costs or intermediary net worth (see, e.g., [Andrés et al. 2004](#), [Chen et al. 2012](#), [Ellison and Tischbirek 2014](#), [Carlstrom et al. 2017](#)); [Ray et al. \(2024\)](#) embed the preferred habitat framework in a New Keynesian model with a representative household, with a richer risk factor structure and a wider set of assets, to study the positive implications of QE, so that the dynamics induced by household heterogeneity and the normative questions we study are absent. Macro-finance term structure models in the tradition of [Ang and Piazzesi \(2003\)](#) and [Piazzesi \(2005\)](#) link term premia to macroeconomic factors and to the monetary policy rule; our model derives that link from market clearing between clientele investors, firms, and arbitrageurs.

More generally, we contribute to a large literature studying the effects of QE in general equilibrium models, in which financial frictions break the “QE neutrality” result of [Wallace \(1981\)](#) (see, e.g., [Gertler and Karadi 2011](#), [Gertler and Karadi 2013](#), [Cúrdia and Woodford 2011](#), [Sims and Wu 2021](#), [Karadi and Nakov 2021](#)). Closest to our normative results, [Sims et al. \(2023\)](#) argue that optimal policy uses the short rate to offset natural rate shocks and QE to offset credit shocks; [Alexandrov and Brunnermeier \(2026\)](#) assign the interest rate to aggregate efficiency and the balance sheet to the distribution of risk; and [Silva \(2026\)](#) studies how asset purchases reallocate aggregate risk between leveraged intermediaries and savers. Relative to these papers, term premium volatility is itself a welfare cost in our model because of household portfolio heterogeneity, which gives balance sheet policy a stabilization objective even when the only shocks are to the natural rate; and under commitment, the path of the short rate shapes household exposure to risk as well.

Our welfare analysis relates to optimal monetary policy with heterogeneous households ([Bhandari et al. 2021](#), [Acharya et al. 2023](#)) and to the HANK literature ([Kaplan et al. 2018](#), [Auclert 2019](#), [Cui and Sterk 2021](#), [Debortoli and Galí 2017](#)), in which monetary policy redistributes across households through marginal propensities to consume and unhedged interest rate exposures. [Kekre and Lenel \(2022\)](#) study how heterogeneity in risk tolerance makes monetary policy move risk premia. Heterogeneity in our model arises from portfolio

⁴Theoretical contributions include [Greenwood and Vayanos \(2014\)](#), [Greenwood et al. \(2016\)](#), [King \(2019b\)](#), [King \(2019a\)](#), and [Kekre et al. \(2024\)](#). Empirically, bond supply and asset purchases move term premia, consistent with the demand and supply “preferred habitat” frictions considered in this paper, and these frictions help explain the transmission of large-scale asset purchases (see, e.g., [Krishnamurthy and Vissing-Jorgensen 2011](#), [Hamilton and Wu 2012](#), [D’Amico and King 2013](#), [Li and Wei 2013](#), [Krishnamurthy and Vissing-Jorgensen 2012](#), [Cahill et al. 2013](#), [King 2019b](#), [Fieldhouse et al. 2018](#), [Di Maggio et al. 2020](#), [Debortoli et al. 2020](#)).

exposure to *aggregate* risk, and thus the induced premium movements are both a welfare cost and a first-order transmission channel, which balance sheet policies can target directly. Our paper also relates to models in which policy responds to financial conditions or credit spreads (see, e.g., [Angeletos et al. 2023](#), [Campbell and Nemyrev 2026](#)). Our results on both the transmission and the optimal design of conventional policy are closely related to those of [Caballero and Simsek \(2024\)](#) and [Caballero et al. \(2024\)](#): the transmission of monetary policy to the borrowing rates relevant for macroeconomic dynamics is imperfect and, more fundamentally, depends on the systematic policy rule, which determines the endogenous pricing of risk; in our model, stable term premia are in addition a social objective.⁵

Our work also relates to research in an international setting. Most closely related to our optimal policy results is [Itskhoki and Mukhin \(2023\)](#), who show that with segmented currency markets, monetary policy should target the output gap while foreign exchange interventions eliminate currency risk; our separation principle is the analogue for a closed economy across the entire term structure. Closer to our framework, [Gourinchas et al. \(2025\)](#) and [Greenwood et al. \(2023\)](#) study the determination of bond yields and exchange rates in a preferred habitat model, but do not study general equilibrium effects.

2 Model

Time is continuous and denoted by $t \in [0, \infty)$. The model features relatively standard firm and household sectors. The household sector comprises differentiated households making labor, consumption, and savings decisions. Intermediate goods are produced by monopolistically competitive firms using labor; these firms set prices but face nominal rigidities in the form of Rotemberg pricing frictions. Differentiated goods are aggregated by a perfectly competitive final goods retail sector.

There is a continuum of zero-coupon bonds with maturity $\tau \in [0, T]$: a τ -maturity bond has price $P_t^{(\tau)}$ at time t and pays one dollar at maturity date $t + \tau$. Both firms and households face financial frictions when trading bonds. Each household makes borrowing and savings decisions, but must trade with a bond fund which has fixed portfolio weights across bonds. This friction captures a number of realistic features of household portfolio decisions. The different portfolio weights across households represent a type of preferred

⁵Methodologically, we contribute tools for models combining endogenous risk premia, heterogeneous agents, and optimal policy. Our perturbation scales arbitrageur risk-bearing capacity with the size of aggregate risk, a device also used by [Borovička and Hansen \(2014\)](#), [Itskhoki and Mukhin \(2023\)](#), and [Kekre and Lenel \(2024\)](#). Our method places consumption dispersion in social welfare at the same order as macroeconomic volatility, and remains tractable even with a continuum of households. Relative to the recursive characterization of commitment policy ([Backus and Driffill 1986](#); [Marcet and Marimon 2019](#)), our optimal plans remain finite-dimensional even though policy manipulates a continuum of endogenous volatilities.

habitat for households (for instance, due to demographics such as age) which we summarize parsimoniously as heterogeneity in duration exposure. For example, some households own property financed with long-duration mortgage securities, while others save predominantly in short-duration securities (such as money market mutual funds or bank deposits). Our assumption also reflects the fact that households typically do not buy and sell financial assets directly, but instead invest through pension funds and mutual funds, which tend to allocate their portfolios using fixed weights and rebalance sluggishly (see, e.g., [Koijen and Yogo 2019](#), [Koijen and Yogo 2022](#), [Bretscher et al. 2022](#), [Peng and Wang 2022](#)). In the spirit of the “preferred habitat” literature, we capture these frictions with the parsimonious but stark assumption that each household faces heterogeneous but fixed portfolio weights. Appendix B.4 provides a microfoundation. Firms also face a financial friction in the form of a “working capital” constraint: each firm must borrow a portion of its wage bill in advance of production. Analogous to the household portfolio problem, we also assume that the financing costs for firms depend on a portfolio of bonds with fixed portfolio weights.

Because neither households nor firms can freely trade bonds of different maturities, specialized bond arbitrageurs intermediate bond markets. Arbitrageurs trade bonds across the entire term structure, but have limited risk-bearing capacity. These agents are owned by the household sector, but due to financial frictions do not price bonds using the household stochastic discount factor (SDF). We represent the arbitrageur portfolio problem as a function of risk aversion, but view the risk aversion parameter as a proxy for the risk-bearing capacity of financial intermediaries. This is a friction which hinders arbitrageurs’ ability to trade assets perfectly on behalf of households as a whole.

The monetary authority sets the short-term nominal interest rate i_t , and conducts balance sheet policies. A fiscal authority provides optimal production subsidies but is otherwise passive. We study policy implications in a linearized equilibrium according to second-order welfare approximations. Our perturbation is around a first-best steady state, but our method still allows for nonzero bond term premia which affect first-order macroeconomic dynamics.

2.1 Setup

Risk Factors. An N_z -dimensional vector z_t and an associated vector B_t of independent Brownian motions capture all sources of aggregate shocks; the factors evolve according to

$$dz_t = -\kappa_z z_t dt + \sigma_z dB_t. \quad (1)$$

In the baseline the only factor is aggregate technology, $Z_t = Ze^{z_{1,t}}$, ordered first in z_t (when there is a single factor we simply write z_t); more generally z_t collects household discount-

factor, cost-push, firm financing, bond supply, and policy shocks, which we discuss below. Where applicable, we treat z_t generically.

In equilibrium, bond returns can be written as

$$\frac{dP_t^{(\tau)}}{P_t^{(\tau)}} = \mu_t(\tau) dt + \sigma_t(\tau) dB_t, \quad (2)$$

with $\mu_t(\tau), \sigma_t(\tau)$ determined in equilibrium. No-arbitrage restrictions imply that the nominal short rate is $i_t \equiv \lim_{\tau \rightarrow 0} -\log P_t^{(\tau)}/\tau$ (implemented by the central bank via instantaneously riskless reserves in zero net supply).

Households. The household sector is made up of individuals $i \in [0, 1]$. Household i chooses a consumption bundle $C_t(i)$ and supplies labor $N_t(i)$. Flow utility is

$$u(C_t(i), N_t(i)) = \frac{C_t(i)^{1-\varsigma} - 1}{1-\varsigma} - \frac{N_t(i)^{1+\varphi}}{1+\varphi}, \quad (3)$$

where ς is the inverse intertemporal elasticity of substitution and φ is the inverse Frisch labor elasticity.

As discussed above, households face financial frictions and so cannot trade all bonds freely. Instead, they save and borrow through bond funds, which determine the allocation of the portfolio across assets. A household i therefore faces the following budget constraint

$$d\mathcal{B}_t(i) = [\mathcal{W}_t N_t(i) - \mathcal{P}_t C_t(i)] dt + \mathcal{B}_t(i) d\tilde{R}_t(i) + d\mathcal{F}_t, \quad (4)$$

where $\mathcal{W}_t$ is the nominal wage and $\mathcal{P}_t$ is the price index of the consumption bundle (both of which are taken as given and identical across households). Household i has nominal wealth $\mathcal{B}_t(i)$ (with $\mathcal{B}_0(i)$ given), earning the realized instantaneous return

$$d\tilde{R}_t(i) \equiv \int_0^T \eta_{it}(\tau) \frac{dP_t^{(\tau)}}{P_t^{(\tau)}} d\tau = \tilde{\mu}_t(i) dt + \tilde{\sigma}_t(i) dB_t. \quad (5)$$

The portfolio weights (normalized $\int_0^T \eta_{it}(\tau) d\tau = 1$) are determined by the bond fund with which household i invests, and are taken as given by the household.⁶ Finally, the term $d\mathcal{F}_t$ represents nominal flow transfers (described below). We denote real wealth (in units of steady-state consumption C) by $b + b_t(i) \equiv \frac{\mathcal{B}_t(i)}{C\mathcal{P}_t}$, where b is the steady-state value.

⁶The baseline model assumes each household can only access bond markets through a specific fund. Appendix B.4 shows that fixed weights are individually optimal when households face a fixed cost of rebalancing, because portfolio choice is a second-order problem for households.

The resulting value function of household i at time $t = 0$ is

$$V_0(i) \equiv \max_{\{C_t(i), N_t(i)\}_{t=0}^{\infty}} \mathbb{E}_0 \int_0^{\infty} e^{-\rho t + \varrho_t} u(C_t(i), N_t(i)) dt, \quad (6)$$

where ϱ_t is an aggregate discount-factor shock with dynamics given by

$$d\varrho_t = -\kappa_{\varrho} \varrho_t dt + \sigma_{\varrho} dB_t. \quad (7)$$

The household problem is to maximize (6) subject to (4). As shown in Appendix B.1, intratemporal optimality conditions are standard and given by

$$W_t = C_t(i)^{\varsigma} N_t(i)^{\varphi}, \quad (8)$$

which characterizes consumption and labor as a function of the real wage $W_t = \frac{W_t}{P_t}$. The intertemporal condition (in log-deviations from steady state) is also mostly familiar:

$$\mathbb{E}_t d c_t(i) = \varsigma^{-1} [\tilde{\mu}_t(i) - \pi_t - \kappa_{\varrho} \varrho_t + \xi_t(i)] dt, \quad (9)$$

where $\xi_t(i)$ represents higher-order terms identified in Appendix B.1. Consumption is a function of expected real borrowing rates, but unlike the textbook model, borrowing rates are not pinned down by the policy rate in equilibrium: $\tilde{\mu}_t(i) \neq i_t$.

Intermediate Firms. A continuum of intermediate goods producers indexed by $j \in [0, 1]$ produces differentiated goods $Y_t(j)$ and sets prices $\mathcal{P}_t(j)$. Each firm j faces the same production function $Y_t(j) = Z_t L_t(j)$, where Z_t is aggregate technology and $L_t(j)$ the labor input of firm j . Final output Y_t is produced by a competitive retail sector according to $Y_t^{(\epsilon_t-1)/\epsilon_t} \equiv \int_0^1 Y_t(j)^{(\epsilon_t-1)/\epsilon_t} dj$, where exogenous shocks to the elasticity of substitution between goods ϵ_t represent markup shocks (defined below). This implies the demand curve and price index $Y_t(j) = (\mathcal{P}_t(j)/\mathcal{P}_t)^{-\epsilon_t} Y_t$ and $\mathcal{P}_t^{1-\epsilon_t} = \int_0^1 \mathcal{P}_t(j)^{1-\epsilon_t} dj$, respectively. The flow nominal profit of firm j is

$$d\Pi_t(j) = [(1 + \tau^y) \mathcal{P}_t(j) Y_t(j) - \mathcal{W}_t L_t(j) - \mathcal{T}_t^y] dt - d\Theta_t(j), \quad (10)$$

where τ^y is a production subsidy, financed by transfers $\mathcal{T}_t^y = \int_0^1 \tau^y \mathcal{P}_t(j) Y_t(j) dj$. In addition to the wage bill $\mathcal{W}_t L_t(j)$, firms face the following flow costs of production:

$$d\Theta_t(j) = \frac{\vartheta}{2} \pi_t(j)^2 \mathcal{P}_t Y_t dt + \mathcal{W}_t L_t(j) d\tilde{R}_t. \quad (11)$$

The first term captures nominal rigidities. Firms choose $\pi_t(j)$, the rate of change of their prices ($d\mathcal{P}_t(j) = \mathcal{P}_t(j)\pi_t(j)dt$). Whenever $\vartheta > 0$, firms face Rotemberg costs of choosing $\pi_t(j) \neq 0$. The second term captures a “working capital” financing friction. The cost of financing is

$$d\check{R}_t \equiv \int_0^T \beta_t(\tau) \frac{dP_t^{(\tau)}}{P_t^{(\tau)}} d\tau = \hat{\mu}_t dt + \hat{\sigma}_t dB_t, \quad (12)$$

where $\beta_t \equiv \int_0^T \beta_t(\tau) d\tau$ represents the total fraction of the wage bill which must be financed in advance (taken as given by the firm).

Taking as given CES demand and the costs of production, the firm problem is

$$U_0(j) \equiv \max_{\{\pi_t(j)\}_{t=0}^\infty} E_0 \int_0^\infty e^{-\rho t} Q_t^H d\Pi_t(j). \quad (13)$$

Firms transfer profits to households; we assume that profits are discounted by $e^{-\rho t} Q_t^H$, the average (nominal) SDF across households.

We focus on a symmetric equilibrium in which each firm chooses the same prices. As shown in Appendix B.2, this delivers a mostly familiar nonlinear Phillips curve,

$$E_t d(Q_t^H Y_t \pi_t) = Q_t^H Y_t [(\rho - \pi_t)\pi_t + \vartheta^{-1}((\epsilon_t - 1)(1 + \tau^y) - \epsilon_t MC_t)] dt, \quad (14)$$

where $MC_t \equiv (1 + \hat{\mu}_t) \frac{W_t}{Z_t}$. Interpreting the exogenous fluctuations in ϵ_t as markup shocks, the nonlinear Phillips curve is standard conditional on marginal costs; the departure from the textbook case comes entirely through marginal costs. In particular, the working capital channel depends on duration-weighted borrowing costs.

Arbitrageurs. In addition to bond demand from households and firms, bonds are also traded by arbitrageurs. Arbitrageurs actively choose holdings of short- and long-maturity bonds. The representative arbitrageur solves the following mean-variance problem:

$$\max_{\{\mathcal{X}_t(\tau)\}_{\tau=0}^T} E_t d\mathcal{X}_t - \frac{a_t}{2} \text{Var}_t d\mathcal{X}_t, \quad \text{s.t.} \quad d\mathcal{X}_t = \mathcal{X}_t i_t dt + \int_0^T \mathcal{X}_t(\tau) \left(\frac{dP_t^{(\tau)}}{P_t^{(\tau)}} - i_t dt \right) d\tau, \quad (15)$$

choosing nominal bond holdings $\mathcal{X}_t(\tau)$ across all maturities τ . Arbitrageurs engage in carry trades across the term structure, trading off higher expected returns against the volatility of their balance sheet. This tradeoff is governed by the parameter a_t . Formally, a_t is a risk aversion parameter; more generally, a_t proxies for capital constraints, Value-at-Risk constraints, or any frictions which imply that bond returns are not determined as in a model

with a representative household with perfect access to financial securities. We assume that arbitrageurs transfer all gains and losses in their wealth $d\mathcal{X}_t$ to households each period.

The arbitrageur optimality conditions are given by

$$\mu_t(\tau) - i_t = a_t \int_0^T \mathcal{X}_t(\tau') \sigma_t(\tau) \sigma_t(\tau')^\top d\tau'. \quad (16)$$

Hence, arbitrageurs ensure that no risk-free arbitrage opportunities exist. Equation (16) shows that the expected excess return of any τ -maturity bond is a function of the covariance of bond returns across the entire term structure, scaled by arbitrageur risk aversion a_t and weighted by holdings $\mathcal{X}_t(\tau)$ (determined in equilibrium by market clearing). We assume that $a_t = \frac{a}{p_t Y}$; that is, risk aversion is constant relative to prices (scaled by steady-state output), and we define real arbitrageur bond positions as $X_t(\tau) = \frac{\mathcal{X}_t(\tau)}{p_t Y}$.

Government. The fiscal authority sets the production subsidy τ^y , which is self-financed through lump-sum taxes on firms. The central bank sets the nominal interest rate i_t , and the fiscal authority pays this interest i_t on short-term bonds (reserves). Besides the production subsidy, the consolidated government may also issue bonds $G_t(\tau) \equiv \frac{g_t(\tau)}{p_t Y}$. Here, $G_t(\tau)$ represents the issuance of bonds by the fiscal authority net of central bank asset purchases $QE_t(\tau)$; in deviations from steady state, we write $g_t(\tau)$ for gross fiscal issuance and $qe_t(\tau)$ for central bank holdings, so that net supply is $g_t(\tau) - qe_t(\tau)$. All gains or losses on the consolidated government's bond issuance are rebated lump-sum to households.

As a benchmark, we assume the central bank can use balance sheet tools and the short rate freely. More generally, we suppose that the central bank faces two constraints. First, the central bank faces adjustment costs when setting the short-term rate:

$$Y_t \Psi_t^i \equiv Y_t \frac{\psi^i}{2} (i_t - \bar{i}_t)^2, \quad (17)$$

where $\bar{i}_t$ represents a potentially time-varying (exogenous) policy target. We interpret (17) as capturing, in reduced form, constraints such as the effective lower bound; to maintain our linear-quadratic approximation, we assume a symmetric loss function. Thus, whenever $\psi^i > 0$, setting the short-term interest rate away from $\bar{i}_t$ implies deadweight losses.

Second, central bank bond holdings may be subject to real deadweight holding costs (measured in terms of output) which take the form:

$$Y_t \Psi_t^S \equiv Y_t \int_0^T \frac{\psi(\tau)}{2} QE_t(\tau)^2 d\tau. \quad (18)$$

Thus, whenever $\psi(\tau) > 0$, central bank holdings away from steady state imply deadweight losses. These are holding costs which the central bank faces but private bond arbitrageurs do not. They can also be interpreted as capturing, in reduced form, the fiscal costs associated with gains and losses on the central bank balance sheet (see, e.g., [d'Avernas et al. 2024](#)).

Market Clearing. In the symmetric equilibrium, all firms make the same decisions, so that $Y_t(j) = Y_t$, $\mathcal{P}_t(j) = \mathcal{P}_t$, $\pi_t(j) = \pi_t$, $L_t(j) = L_t$. Define aggregate consumption and labor supply as $C_t = \int_0^1 C_t(i) di$ and $N_t = \int_0^1 N_t(i) di$. Market clearing in production, labor markets, and consumption implies $Y_t = Z_t L_t = Z_t N_t$ and $C_t = Y_t (1 - \frac{\vartheta}{2} \pi_t^2 - \Psi_t^i - \Psi_t^S)$, where aggregate consumption differs from aggregate output due to deadweight loss from price changes and policy frictions. Bond market clearing implies

$$\mathcal{X}_t(\tau) - \beta_t(\tau) \mathcal{W}_t N_t + \int_0^1 \eta_{it}(\tau) \mathcal{B}_t(i) di = \mathcal{G}_t(\tau), \quad \forall \tau > 0. \quad (19)$$

2.2 First- and Second-Order Perturbations

We study perturbations of the nonlinear model by considering the family of economies indexed by the scalar $\xi \geq 0$, where the diffusion matrix from (1) is parameterized by $\sigma_z(\xi) \equiv \xi \sigma_z$ and arbitrageur risk aversion is parameterized by $a(\xi) \equiv \xi^{-2} a$, with σ_z, a fixed. Thus we have $a(\xi) \sigma_z(\xi) \sigma_z(\xi)^\top \rightarrow a \sigma_z \sigma_z^\top$ as $\xi \rightarrow 0$. That is, the product of risk aversion and physical risk remains nonzero and bounded even as risk becomes small.⁷ The following proposition characterizes the optimal allocation in the riskless limit.

Proposition 1 (First-best allocation). *Consider the limiting case where $\xi \rightarrow 0$. Assume zero net steady-state wealth ($b = 0$), and that wealth is initially equalized ($\mathcal{B}_0(i) \equiv \mathcal{B}_0$). Then with perfect arbitrage ($a = 0$), the model is equivalent to a representative agent model. With the optimal production subsidy and absent markup and firm-financing shocks, the first-best allocation is obtained with flexible prices ($\vartheta = 0$) when the nominal short rate $i_t = \rho$.*

Proposition 1 shows that, so long as the production subsidy is chosen optimally to remove steady-state distortions from imperfect competition and the working capital frictions, the first-best is achieved in the flexible price, perfect arbitrage, zero net steady-state wealth model.⁸ Further, in the proof of Proposition 1, we show that if steady-state policy satisfies

⁷This assumption can also be formalized in terms of the size of the intermediary sector relative to the size of asset markets (see, e.g., [Itskhoki and Mukhin 2023](#)), and captures explicitly that the arbitrageur “risk aversion” parameter a measures the degree of imperfect arbitrage.

⁸Away from the deterministic limit, the equivalence to a representative agent model under perfect arbitrage requires $b = 0$. The main text maintains this assumption throughout, but we relax this in Appendix D. With $b \neq 0$, heterogeneous portfolio revaluation effects arise at first order even in the “perfect-arbitrage” econ-

$i = \rho$ and $G(\tau) = -\beta(\tau)$, then the steady state features zero inflation ($\pi = 0$) and zero arbitrageur positions ($X(\tau) = 0$). This in turn implies that the steady states of the first-best equilibrium and of the sticky-price, imperfect-arbitrage equilibrium coincide. In particular, steady-state consumption and labor supply across households are equalized ($C(i) = C$, $N(i) = N$), and the steady state is efficient.

Log-Linear Perturbation. We study the dynamics of the model around the “low risk, low risk-bearing capacity” limiting case as $\xi \rightarrow 0$ but $a \geq 0$. This approximation method delivers tractable first- and second-order approximations of the model, while still allowing for first-order variation in risk premia.⁹

The arbitrageur optimality condition (16) becomes

$$\mu_t(\tau) - i_t = a\sigma(\tau)\Lambda_t \equiv \lambda_t(\tau), \quad \text{where } \Lambda_t = \int_0^T \sigma(\tau)^\top x_t(\tau) d\tau. \quad (20)$$

The expected excess return of any τ -maturity bond is thus a function of bond-specific risk loadings $\sigma(\tau)$ and a global set of risk prices Λ_t . The market price of risk depends on arbitrageur risk aversion a as well as holdings $x_t(\tau)$. Under our perturbation, the loadings $\sigma(\tau)$ are time-invariant, but they are still present at first order, and are endogenous.

Denote aggregate output in the first-best equilibrium by Y_t^n , and define the output gap $X_t \equiv \frac{Y_t}{Y_t^n}$. Aggregating (9) across households, the dynamics of the (log) output gap x_t are governed by a modified aggregate Euler equation:

$$E_t dx_t = \varsigma^{-1} [\tilde{\mu}_t - \pi_t - v_t] dt, \quad \text{where } \tilde{\mu}_t = \int_0^T \eta(\tau)\mu_t(\tau) d\tau, \quad \eta(\tau) \equiv \int_0^1 \eta_i(\tau) di. \quad (21)$$

The term v_t is the “natural” rate, which is the real borrowing rate in the first-best allocation from Proposition 1. The “household effective borrowing rate” (in terms of deviation from steady state) is $\tilde{\mu}_t$, and $\eta(\tau)$ represents the average across households of the steady-state portfolio weight in τ -maturity bonds.

Rewriting the firm optimality condition in terms of output gaps gives a New Keynesian Phillips curve:

$$E_t d\pi_t = \left[\rho\pi_t - \kappa \left(x_t + \alpha_{\hat{\mu}}\hat{\mu}_t + \alpha_{\hat{\beta}}\hat{\beta}_t \right) - u_t \right] dt, \quad (22)$$

where the slope of the Phillips curve κ is given by (B2). In addition to the usual drivers of

omy. This introduces a further redistribution channel (in the spirit of the unhedged interest rate exposures of Auclert 2019); assuming $b = 0$ isolates the imperfect arbitrage risk pricing channel.

⁹For notational convenience, we drop dependence on ξ in the main text.

inflation, our model features

$$\hat{\mu}_t = \int_0^T \beta(\tau) \mu_t(\tau) d\tau, \quad \hat{\beta}_t = \int_0^T [\beta_t(\tau) - \beta(\tau)] d\tau, \quad (23)$$

which capture the “firm effective borrowing rate” and firm financing shocks (respectively), and the parameters $\alpha_{\hat{\mu}} = \frac{1}{(1+\rho\beta)(\varsigma+\varphi)}$, $\alpha_{\hat{\beta}} = \frac{\rho}{(1+\rho\beta)(\varsigma+\varphi)}$, where $\beta = \int_0^T \beta(\tau) d\tau$. The markup shock u_t arises from exogenous variation in demand elasticities ϵ_t .

Second-Order Social Welfare Loss. Aggregate dynamics thus depart from the textbook model already at first order. The welfare consequences differ as well. A second-order expansion of social welfare relative to the first-best allocation gives the per-period social loss

$$\mathcal{L}_t \equiv \underbrace{(\varsigma + \varphi)x_t^2 + \vartheta\pi_t^2}_{\text{conventional}} + \underbrace{\frac{\varsigma}{\varphi}(\varphi + \varsigma)\text{Var}_i \check{c}_t(i)}_{\text{consumption dispersion}} + \underbrace{\int_0^T \psi(\tau) qe_t(\tau)^2 d\tau + \psi^i(i_t - \bar{i}_t)^2}_{\text{policy-instrument costs}}. \quad (24)$$

The first two terms in (24) capture the conventional welfare losses associated with the nominal rigidities in the model; these terms arise in representative agent New Keynesian (RANK) models. Relative to RANK, the social loss also depends on consumption dispersion across households, $\text{Var}_i \check{c}_t(i) \equiv \int_0^1 \check{c}_t(i)^2 di$: because risk-sharing is imperfect, dispersion implies welfare losses. The final terms in (24) represent losses associated with the short rate and balance sheet policies of the central bank, from equations (17) and (18) respectively.

Sufficient Statistics for Wealth Heterogeneity. The dynamics of the output gap x_t and inflation π_t in (21) and (22) are similar to the representative agent case considered in Ray et al. (2024). However, even to first order, our model features more complicated dynamics, due to the endogenous time variation in bond demand from heterogeneous households. Moreover, welfare depends on consumption dispersion $\text{Var}_i \check{c}_t(i)$ across households, which itself varies over time. This complexity arises because the continuum of household bond positions enters bond market clearing, and thus the arbitrageur optimality conditions.

We now show that tracking the entire wealth distribution is not necessary. Instead, only a small number of endogenous sufficient statistics are required. Let $b_t(i)$ denote the real wealth deviations from steady state of household i , and $\check{b}_t(i) \equiv b_t(i) - b_t$ denote the difference from aggregate household wealth $b_t \equiv \int b_t(i) di$; similarly let $\check{c}_t(i) \equiv c_t(i) - c_t$. Define the following risk-weighted portfolio exposures

$$\eta_i^\sigma \equiv \int_0^T \eta_i(\tau) \sigma(\tau) d\tau, \quad \eta^\sigma \equiv \int_0^1 \eta_i^\sigma di, \quad \beta^\sigma \equiv \int_0^T \beta(\tau) \sigma(\tau) d\tau, \quad (25)$$

so that integrating (20) against the portfolio weights gives

$$\tilde{\mu}_t(i) = i_t + a\eta_i^\sigma \Lambda_t, \quad \tilde{\mu}_t \equiv \int_0^1 \tilde{\mu}_t(i) di = i_t + a\eta^\sigma \Lambda_t, \quad \hat{\mu}_t = \beta i_t + a\beta^\sigma \Lambda_t. \quad (26)$$

With zero steady-state wealth $b = 0$, aggregate real wealth evolves according to

$$db_t = [\rho b_t + (\varsigma + \varphi)x_t + f_t] dt, \quad (27)$$

where f_t denotes the aggregate net non-labor transfer flow to households (firm dividends and government rebates). We assume that transfers rebate the aggregate net saving flow ($f_t = -(\varsigma + \varphi)x_t$) so $b_t \equiv 0$ (see Appendix D for arbitrary transfers and for $b \neq 0$, which additionally allows for revaluation effects). Then relative wealth of household i evolves as

$$d\check{b}_t(i) = [\rho\check{b}_t(i) - (1 + \varsigma/\varphi)\check{c}_t(i)] dt, \quad (28)$$

using relative labor supply $\varsigma\check{c}_t(i) + \varphi\check{n}_t(i) = 0$ from (8). Differential consumption $\check{c}_t(i)$ evolves according to

$$E_t d\check{c}_t(i) = \varsigma^{-1}\check{\lambda}_t(i) dt = \varsigma^{-1}a(\eta_i^\sigma - \eta^\sigma)\Lambda_t dt, \quad (29)$$

where $\check{\lambda}_t(i) \equiv \tilde{\mu}_t(i) - \tilde{\mu}_t = a(\eta_i^\sigma - \eta^\sigma)\Lambda_t$ denotes the relative portfolio premium. Heterogeneity is therefore driven by the risk exposure differences $\eta_i^\sigma - \eta^\sigma$, which are N_z -dimensional vectors. This insight motivates the following risk-weighted cross-sectional aggregates:

$$b_t^\sigma \equiv \int_0^1 (\eta_i^\sigma)^\top \check{b}_t(i) di, \quad \check{c}_t^\sigma \equiv \int_0^1 (\eta_i^\sigma)^\top \check{c}_t(i) di, \quad (30)$$

and the positive semi-definite matrix $\check{\Sigma} \equiv \int_0^1 (\eta_i^\sigma - \eta^\sigma)^\top (\eta_i^\sigma - \eta^\sigma) di$. The following theorem shows that these objects are sufficient for characterizing the aggregate first-order dynamics of the model, risk prices, and the time-variation in cross-sectional consumption dispersion.

Theorem 1 (Sufficient statistics). *Suppose initial relative wealth is spanned by exposures: $\check{b}_0(i) = (\eta_i^\sigma - \eta^\sigma)\gamma_0^b$ for some $\gamma_0^b \in \mathbb{R}^{N_z}$. Then:*

1. For all t , $\check{b}_t(i) = (\eta_i^\sigma - \eta^\sigma)\gamma_t^b$ and $\check{c}_t(i) = (\eta_i^\sigma - \eta^\sigma)\gamma_t^c$, where

$$d\gamma_t^b = [\rho\gamma_t^b - (1 + \varsigma/\varphi)\gamma_t^c] dt, \quad E_t d\gamma_t^c = a\varsigma^{-1}\Lambda_t dt. \quad (31)$$

2. The objects (30) satisfy $b_t^\sigma = \check{\Sigma}\gamma_t^b$ and $\check{c}_t^\sigma = \check{\Sigma}\gamma_t^c$.

3. Cross-sectional consumption dispersion is given by $\text{Var}_i \check{c}_t(i) = \gamma_t^{c\top} \check{\Sigma} \gamma_t^c = \check{c}_t^{\sigma\top} \check{\Sigma}^+ \check{c}_t^\sigma$.

4. Risk prices are given by

$$\Lambda_t = (\beta^\sigma)^\top \left[(1 + \varsigma + \varphi)x_t + \frac{1 + \varphi}{\varsigma + \varphi} z_{1,t} \right] + g_t^\sigma - qe_t^\sigma + \hat{\beta}_t^\sigma - b_t^\sigma, \quad (32)$$

where the balance sheet and supply terms are defined as $qe_t^\sigma \equiv \int_0^T \sigma(\tau)^\top qe_t(\tau) d\tau$, $g_t^\sigma \equiv \int_0^T \sigma(\tau)^\top g_t(\tau) d\tau$, and $\hat{\beta}_t^\sigma \equiv \int_0^T \sigma(\tau)^\top [\beta_t(\tau) - \beta(\tau)] d\tau$ collects firm-financing shocks.

Thus, since Λ_t depends on the household side only through b_t^σ , the entire cross-section is summarized by $(b_t^\sigma, \check{c}_t^\sigma)$.

The key insight of Theorem 1 is that what matters to arbitrageurs is risk exposure, not the positions taken in specific bonds *per se*. Risk exposure is a function of the sources of aggregate volatility in the model, which are low-dimensional. Differential consumption choices arise precisely because borrowing rates differ across households; and by the same insight, borrowing rates differ only because risk exposures differ. Thus, all relevant first- and second-order properties of the model are summarized by this small set of risk-weighted wealth statistics. Note, however, that because the bond loadings $\sigma(\tau)$ are endogenous, all of the household and firm risk-weighted objects are endogenous as well.

2.3 Equilibrium

Intuitively, general equilibrium in this model works as follows. From the perspective of households, the key factor is how sensitive their borrowing rates are to the short rate (and movements in asset demand, including balance sheet policies). The model reduces to a benchmark New Keynesian model when these rates move one-for-one, but in general $\mu_t(\tau) \neq i_t$. Suppose that long-term borrowing rates are highly responsive to the policy rate. Then household borrowing is also highly sensitive to the policy rate, and therefore the growth rate of consumption also reacts strongly to the policy rate. On the other hand, when long-term rates are insensitive to the policy rate, the pass-through of changes in the policy rate to households is weakened. Through the borrowing decisions of the household, the growth rate of consumption is less responsive to the policy rate.

However, the sensitivity of the effective borrowing rate to the policy rate is an equilibrium object, which also depends on financial markets. Bond prices adjust to clear bond markets, such that arbitrageurs' portfolio allocation satisfies their optimality conditions given the demand from households and firms. Thus, the term structure is determined by the interactions

between these different classes of investors; these interactions pin down the risk-adjusted dynamics of the economy. Moreover, because differentiated households also make consumption and savings decisions as a function of different borrowing rates, macroeconomic dynamics themselves depend on these interactions. General equilibrium is obtained when these two forces balance. Thus, characterizing equilibrium involves two key steps: first, understanding the differences between the actual and risk-adjusted dynamics of the economy; and second, linking household savings and consumption choices with the bond prices determined in imperfect financial markets.

The following proposition characterizes the behavior of asset prices, conditional on an arbitrary set of state variables.

Proposition 2 (Asset prices). *Suppose that the first-order equilibrium admits a stationary linear representation with an N_s -dimensional state variable s_t , which satisfies $i_t = \iota^\top s_t$ and $\Lambda_t = Ls_t$, and evolves according to $ds_t = -\Gamma s_t dt + \sigma_s dB_t$. Then an affine (log) bond price equilibrium with $-\log P_t^{(\tau)} = A(\tau)^\top s_t + C(\tau)$ satisfies*

$$A(\tau) = [I - e^{-M\tau}] M^{-1} \iota, \quad M^\top = \Gamma + a\sigma_s L. \quad (33)$$

The matrix M describes the risk-adjusted dynamics of the state. In the first-best, arbitrageurs are perfectly risk-neutral ($a = 0$), so we have $M = \Gamma^\top$. However, when $a \neq 0$, M additionally depends on risk prices (through L); and risk prices themselves depend on M , so in equilibrium, M appears on both sides of equation (33).

Bond pricing is similar to [Vayanos and Vila \(2021\)](#): conditional on the equilibrium dynamics of the model, the only differences come from the more general dynamics of the short-term interest rate and from the market clearing conditions in bond markets.

Proposition 2 characterizes risk pricing, taking as given aggregate dynamics. We now approach the problem from the other direction.

Proposition 3 (Aggregate dynamics). *Collect all risk factors z_t and the wealth statistics b_t^σ defined in (30) into the vector s_t . Collect the output gap x_t , inflation π_t , and the consumption statistics $\check{c}_t^\sigma$ into the vector ω_t . Assume that the short rate i_t , bond supply $g_t(\tau)$, and balance sheet tools $qe_t(\tau)$ are functions of these objects.*

Then given $\check{\Sigma}$, η^σ , and β^σ defined in (25), and under the saddle-path condition of Appendix B.3, the linear rational expectations equilibrium is given by

$$ds_t = -\Gamma s_t dt + \sigma_s dB_t, \quad \omega_t = \Omega s_t, \quad (34)$$

where the matrices Γ, Ω , and σ_s are given in Appendix B.3.

Together, Propositions 2 and 3 characterize the equilibrium.

Theorem 2 (Existence and uniqueness). *An affine equilibrium is one in which the state and jump variables evolve according to equations (34), and asset prices are determined by the solution to (33). The per-period loss function can be written $\mathcal{L}_t = s_t^\top \Psi s_t$. Suppose that at $a = 0$ the policy rule satisfies standard regularity conditions. Then in a neighborhood of risk neutrality, the equilibrium exists and is (locally) unique.*

Note that the dynamics matrix of the state Γ depends on the affine coefficients $A(\tau)$ through the effective borrowing rates of households, which are themselves functions of the risk-adjusted dynamics matrix M . Thus, equilibrium is determined as a fixed point that produces asset price dynamics consistent with equilibrium dynamics of the macroeconomy and vice versa. In general, an affine equilibrium of this type may not exist, or there may be multiple solutions to this fixed point problem. However, when $a = 0$, the model reduces to a standard New Keynesian model. Theorem 2 shows that this equilibrium persists and is locally unique as we depart from risk neutrality.

3 Aggregate Dynamics: Analytical Results

Although Theorem 2 characterizes the first-order dynamics and second-order welfare properties of the model, in general it provides little guidance on how either differs from textbook models. In this section, we derive clear predictions regarding first-order dynamics in a simplified setup where the central bank follows ad hoc policy rules.

We focus on a version of the model with only one source of aggregate uncertainty: natural rate shocks $v_t \equiv -\varsigma \kappa_z \frac{1+\varphi}{\varsigma+\varphi} z_t$, which are driven by technology z_t only. We simplify further for most of this section and study the case of fully rigid prices ($\vartheta \rightarrow \infty$). Technology follows an exogenous Ornstein-Uhlenbeck process with mean reversion κ_z and diffusion σ_z , and the dynamics of the rest of the economy are summarized by the following system of equations

$$E_t dx_t = \varsigma^{-1} [i_t + a\eta^\sigma \lambda_t - v_t] dt, \quad E_t d\check{c}_t^\sigma = \varsigma^{-1} a \check{\Sigma} \lambda_t dt, \quad (35)$$

$$db_t^\sigma = [-(1 + \varsigma/\varphi) \check{c}_t^\sigma + \rho b_t^\sigma] dt, \quad \lambda_t = -qe_t^\sigma + \beta^\sigma ((1 + \varsigma + \varphi)x_t - v_t/(\varsigma \kappa_z)) - b_t^\sigma. \quad (36)$$

With one aggregate risk factor, the risk price λ_t and the other composite endogenous risk objects are scalars. Because prices are fully rigid, the financing frictions which firms face do not affect their pricing decisions. However, through market clearing, firm demand for bonds still transmits aggregate fluctuations to risk premia.

Throughout, we assume that the central bank follows Taylor-type rules for its policy tools, $i_t = \phi_x x_t + \epsilon_{i,t}$, $qe_t^\sigma = \phi_x^\sigma x_t + \epsilon_{q,t}^\sigma$ (for dynamics, we only need to specify the risk-weighted

path of the central bank's balance sheet). The terms $\epsilon_{i,t}$ and $\epsilon_{q,t}^\sigma$ represent zero-probability monetary policy shocks, which evolve deterministically according to

$$d\epsilon_{i,t} = -\kappa_i \epsilon_{i,t} dt, \quad d\epsilon_{q,t}^\sigma = -\kappa_q \epsilon_{q,t}^\sigma dt.$$

We first consider a benchmark case in which there is no working capital friction, so that firms do not borrow in bond markets, and the central bank does not use balance sheet tools systematically. The following proposition characterizes the dynamics under conventional and unconventional monetary policy.

Proposition 4 (Rigid price dynamics, no firm borrowing). *Suppose that $a > 0$, $\beta(\tau) = 0$, $\phi_x^\sigma = 0$, $\phi_x > 0$, and $\eta^\sigma > 0$. If household wealth is initially equalized at steady state (so $b_0^\sigma = 0$), then the equilibrium dynamics of the model are equivalent to a representative agent model, and $\lambda_t = 0$. We have*

$$\frac{\partial x_t}{\partial v_t} > 0, \quad \frac{\partial \lambda_t}{\partial v_t} = 0, \quad \frac{\partial x_t}{\partial \epsilon_{i,t}} < 0, \quad \frac{\partial \lambda_t}{\partial \epsilon_{i,t}} = 0.$$

However, following an unconventional monetary policy shock, we have

$$\frac{\partial x_t}{\partial \epsilon_{q,t}^\sigma} > 0, \quad \frac{\partial \lambda_t}{\partial \epsilon_{q,t}^\sigma} < 0,$$

and the model is no longer equivalent to a representative agent model; if $\tilde{\Sigma} > 0$, then $b_{t+k}^\sigma \neq 0$ following the shock.

Proposition 4 shows that without firm borrowing and with initially equal asset positions across households, the model with only conventional monetary policy is identical to a representative agent model. If household wealth is initially at steady state, then arbitrageurs initially take zero positions in bond markets and hence do not require any risk compensation. Following movements in the short-term interest rate (either endogenously or due to policy shocks), while the macroeconomy moves away from steady state, household wealth does not change. Without firm borrowing, arbitrageurs therefore continue to maintain zero positions, and risk prices do not move.

However, even without firm borrowing and with an equalized wealth distribution, the predictions of the model differ from those of a representative agent model when considering unconventional policy. Because an increase in asset purchases by the central bank reduces the holdings of arbitrageurs on impact, required risk compensation falls. Holding fixed the short-term interest rate, the “effective” borrowing rate faced by households falls as well, which is stimulative. But the decline in borrowing rates is larger for households who borrow with long-

maturity assets, so the induced increase in consumption is larger for these households than for households with access to short-maturity bonds. Consumption dispersion $\check{c}_t^\sigma$ becomes nonzero, and households' risk-weighted asset positions b_t^σ move away from steady state.

Proposition 4 shows that conventional and unconventional policies may be substitutes in the sense that shocks to either tool can boost output, but the transmission mechanisms and implied dynamics following the shocks differ. The welfare implications therefore differ as well, as we show in Section 5.

The next result shows that when firms borrow as part of the production process, the dynamics of the model differ further from a representative agent model.

Proposition 5 (Rigid price dynamics, general case). *Suppose that $\beta^\sigma > 0$, $\eta^\sigma > 0$, and $\check{\Sigma} > 0$. If $0 < a < \bar{a}$ for some upper bound on arbitrageur risk aversion, and the Taylor rule coefficient satisfies $0 < \phi_x < \bar{\phi}_x \equiv \varsigma\kappa_z(\varsigma + \varphi)$, then following a technology shock which decreases the output gap, both the short-term rate i_t and term premia λ_t decline on impact. However, for all $k > \bar{k}$, term premia λ_{t+k} are positive, while the short rate remains below steady state. With balance sheet tools ($\phi_x^\sigma < 0$), the impact decline in term premia is larger.*

To see the intuition for Proposition 5, first note that under risk neutrality, the output gap is an increasing function of the natural rate: $x_t \propto v_t$. When $\phi_x < \bar{\phi}_x$, firms reduce borrowing in a recession caused by a decline in the natural rate. When $\beta^\sigma > 0$, this implies that λ_t declines. But the decline in term premia induces households with long-duration investments to reduce savings (or increase borrowing). Arbitrageurs must fund this borrowing, and eventually take long positions in the bond carry trade, which are larger for longer maturities.

Thus, our model predicts that unconditionally, decreases in the short-term interest rate are associated with decreases in term premia over short horizons; but over longer horizons, the reaction of term premia becomes small or changes sign. Moreover, when the central bank utilizes balance sheet tools, the reaction of term premia over short horizons is larger. However, the *conditional* reaction of term premia to monetary policy shocks differs, as shown in the following corollary.

Corollary 5.1 (Rigid price dynamics, monetary shocks). *Under the assumptions of Proposition 5 with $\phi_x^\sigma = 0$, following a conventional monetary policy shock $\epsilon_{i,t}$ which decreases the policy rate i_t , the output gap x_t increases. The conditional reaction of λ_t to i_t is negative, and smaller in absolute value than the unconditional reaction in Proposition 5 for similarly sized movements in the short rate (assuming ϕ_x is not too large).*

The difference between the unconditional and conditional reactions of term premia to movements in the short rate arises from the endogenous movements of quantities in asset

markets. In this version of the model, all else equal, firm borrowing is positively correlated with the output gap. Technology-driven declines in the output gap also put upward pressure on firm borrowing; however, when the Taylor rule is not overly reactive, firm borrowing falls on net, which reduces the risk compensation required by arbitrageurs. Simultaneously, the short rate declines due to the endogenous reaction of the policy rule. Because all stochastic dynamics of this one-factor model are due to natural rate (technology) shocks, the unconditional correlation between the policy rate and term premia is therefore positive whenever the Taylor rule is not overly reactive. However, because an unexpected decline in the short-term interest rate is expansionary, the correlation between term premia and the short-term interest rate *conditional on a monetary policy shock* is negative.

Further, the endogenous dynamics of the model induced by slow-moving changes in the relative wealth distribution across households imply that these predictions change over longer horizons. Finally, the quantitative predictions change in versions of the model where the central bank is also pursuing balance sheet policies.

4 Empirical Evidence

We now compare the predictions of Section 3 with the response of forward rates to short rates. Rather than decomposing yields into expectations and term premium components (which requires imposing additional structure on the data), we follow the literature and examine how forward rates across the term structure respond to changes in short-term yields. Market participants presumably expect most movements in short rates to have dissipated after 10 to 20 years, so movements in long-dated forward rates mostly reflect movements in term premia; we interpret our findings through this lens. The predictions of our model are largely already documented in this literature: [Hanson and Stein \(2015\)](#) and [Hanson et al. \(2021\)](#) show that long-dated forward rates are highly sensitive to movements in short rates, and the latter show that this sensitivity is short-lived; [Gürkaynak et al. \(2005\)](#) show that long-dated nominal forwards fall in response to monetary policy surprises; and [Nakamura and Steinsson \(2018\)](#) find that intermediate maturity TIPS forward rates react to high-frequency monetary shocks, but longer maturity forwards show little response.

We estimate regressions of the form

$$f_{t+h}^{(\tau)} - f_{t-1}^{(\tau)} = \alpha(\tau, h) + \beta(\tau, h) \left(y_t^{(1)} - y_{t-1}^{(1)} \right) + \epsilon_t(\tau, h), \quad (37)$$

where $f_t^{(\tau)}$ is the τ -year instantaneous forward rate and $y_t^{(1)}$ the 1-year zero-coupon yield on day t . The coefficient $\beta(\tau, h)$ measures the change in the τ -year forward rate between days

$t - 1$ and $t + h$ associated with a one-percentage-point change in the 1-year yield on day t ; the horizon h is measured in trading days. We use the daily nominal zero-coupon yield curves of [Gürkaynak et al. \(2007\)](#), from 1982 onward.¹⁰

We estimate (37) in two ways. Our *unconditional* specification uses OLS on all trading days, so that $\beta(\tau, h)$ reflects the comovement of forward rates with short rates induced by whatever shocks move short rates. Our *monetary shock* specification estimates the same equation by instrumental variables on FOMC announcement days, instrumenting the 1-year yield change with the high-frequency monetary policy surprises of [Bauer and Swanson \(2023\)](#).

Figure 1 summarizes our results. The estimates match the predictions of Section 3 in four respects, three of which replicate patterns documented by [Hanson et al. \(2021\)](#). First, unconditionally, short-horizon movements in short rates are associated with large movements in forward rates across the entire term structure (Panel A). Changes in the 1-year yield are associated with essentially one-for-one changes in the 2-year forward rate on the same day. The pass-through remains positive and large across the entire term structure: even 20-year forward rates increase by roughly 40 basis points following a 1 percentage point increase in short rates. Consistent with Proposition 5, short rates and term premia comove on impact.

Second, this comovement at long maturities is short-lived (Panel B). The response of the 2-year forward rate persists at longer horizons: over the next 150 trading days, the pass-through is still one-for-one (if anything, the point estimates increase above one). In contrast, the response of the 20-year forward rate is statistically significant only for the first few weeks; it falls to roughly zero within about three months (-0.12 at $h = 60$ trading days), and the point estimates turn negative at longer horizons (although the confidence bands include zero).¹¹ [Hanson et al. \(2021\)](#) document the same dissipation. The full-sample estimate at $h = 60$ averages different periods: in rolling 5-year windows, the 60-day response of the 20-year forward is negative in windows ending before 1991, close to zero in windows ending between 1991 and 2008, and positive in windows ending after 2020 (Figure E2). The evidence is consistent with the prediction that the reaction of term premia to short rates becomes small over longer horizons; but sign reversal depends on the sample.

Third, the transmission of monetary policy *shocks* to long-dated forwards is qualitatively different (Panels A and D). At the short end, conditional and unconditional pass-through coincide. But the response to monetary shocks declines much faster with maturity, turning

¹⁰Our sample begins in 1982, when high-quality estimates of long-maturity (up to 20 years) forward rates become available. We use nominal yield curves because high-quality estimates are available for both short- and long-maturity spot and forward yields since the early 1980s. In contrast, Treasury Inflation-Protected Securities (TIPS) were only introduced in the late 1990s and trade significantly less than nominal Treasury securities; high-quality real yield curve estimates are at best available starting in the mid-2000s, so that most of the estimation sample would fall in the zero lower bound period.

¹¹Appendix Figure E1 reports results for $h = 60$, and Figure E2 reports rolling-window estimates.




Figure 1: Pass-Through of Short Rates to Forward Rates

Notes: estimates of $\beta(\tau, h)$ from (37), the response of the τ -year forward rate between trading days $t - 1$ and $t + h$ to a one-percentage-point change in the 1-year yield on day t . Unconditional: OLS on all trading days, 1982-2026, with Newey-West standard errors (90-day bandwidth). Monetary shocks: IV on FOMC announcement days, 1988-2023, instrumenting the change in the 1-year yield with the [Bauer and Swanson \(2023\)](#) surprise, with heteroskedasticity-robust standard errors. Panel A: maturities $\tau = 2$ to 20 years at $h = 0$. Panels B and D: 2- and 20-year forwards at horizons h measured in trading days. Panel C: 20-year forward at $h = 0$, estimated over rolling 5-year windows and dated by the end of each window. Bands are ± 2 standard errors.

negative beyond 9 years and reaching -0.17 (standard error 0.17) for the 20-year forward, compared with $+0.41$ unconditionally. Over the following days, the response of the 20-year forward becomes significantly negative: -0.57 (standard error 0.19) one trading day after the announcement, and it remains significantly negative through the third trading day. Corollary 5.1 predicts that the conditional reaction is negative: the same-day estimate has the predicted sign (although we cannot reject a zero response); the estimate is significantly negative one to three trading days after the announcement. The negative response of long-dated forwards to monetary surprises is documented by [Gürkaynak et al. \(2005\)](#), and the near-zero response at ten years matches [Nakamura and Steinsson \(2018\)](#).

Finally, the unconditional comovement of long-dated forwards with short rates is state dependent (Panel C). Over rolling 5-year windows, the same-day response of the 20-year forward ranges between 0.2 and 0.7 in windows ending before 2011. It rises sharply in windows covering the zero lower bound and QE period, peaking at 1.2 in the window ending in 2015, and remains elevated in windows covering the asset purchases of 2020-21, before falling back to about 0.25 in windows ending in 2023-24, which are dominated by the 2022-23 tightening cycle. This is consistent with the prediction that the unconditional reaction of term premia to short rates is stronger when the central bank utilizes balance sheet tools (although it is important to note that the 1-year yield was constrained by the lower bound).

4.1 Quantifying the Risk-Sharing Motive

The patterns above are consistent with the model’s predictions for the sign and horizon of the comovement of term premia with short rates. Before studying counterfactual policy rules and the optimal design of conventional and unconventional policy, we quantify the magnitude of the risk-sharing motive and show that the welfare cost of imperfect risk-sharing is at least comparable to that of output and inflation volatility. We use the following approximation of social loss (derived in Lemma A1):

$$\frac{\rho}{2} \mathbb{E}_0 \int_0^\infty e^{-\rho t} \mathcal{L}_t dt \approx \frac{\varsigma + \varphi}{2} \text{Var}(x_t) + \frac{\vartheta}{2} \text{Var}(\pi_t) + \frac{\varsigma}{2\varphi} (\varphi + \varsigma) \frac{s_0^2 \text{Var}(\tilde{\lambda}_t)}{\varsigma^2 (\rho + \kappa_p)^2}, \quad (38)$$

where $\tilde{\lambda}_t \equiv a\eta^\sigma \Lambda_t$ is the risk premium earned on the aggregate household effective portfolio, and κ_p is its approximate mean reversion rate. As described in Lemma A1, the term s_0 represents the dispersion in household portfolios under the assumption that all households hold the same “market” bond fund, and only differ in their positions in the short rate. The final term in (38) is therefore a conservative approximation for the loss arising from consumption dispersion in the actual social loss function (24). We evaluate each term using US data, fixing $\rho = 0.02$ and $\varsigma = \varphi = 1$. Table 1 reports the results.

Panel A reports the conventional terms. The output gap is the log difference between real GDP and CBO potential output, and inflation is calculated using core PCE prices. For inflation losses, we use the mean squared deviation of inflation from 2%; this also penalizes persistent deviations from target and so boosts these terms. Conventional terms are largely driven by inflation and so scale with the price-adjustment cost ϑ : with annualized quarterly inflation, they are 0.06, 0.18, and 0.33 percent of steady-state consumption for $\vartheta = 10, 50$, and 100, respectively. Measuring inflation over 12 months smooths out higher-frequency variation, giving smaller values (0.26 rather than 0.33 percent at $\vartheta = 100$ over 1990-2019).

Table 1: Welfare Losses: Conventional and Risk-Sharing Terms

<i>Panel A: Conventional terms</i>						
Sample		SD(x)	RMSE($\pi - 2\%$)	loss		
				$\vartheta = 10$	$\vartheta = 50$	$\vartheta = 100$
1990–2019	Core PCE (q/q)	1.70	0.77	0.06	0.18	0.33
	Core PCE (y/y)	1.70	0.67	0.05	0.14	0.26
1999–2019	Core PCE (q/q)	1.85	0.65	0.06	0.14	0.25
	Core PCE (y/y)	1.85	0.49	0.05	0.09	0.15
<i>Panel B: Risk-sharing terms</i>						
Sample		SD(TP ⁽⁵⁾)	κ_p	loss		
				simple	κ_p -consistent	$D_{eq} = 4$
1990–2019	ACM	83.8	0.21	1.74	3.30	6.13
	Kim-Wright	54.0	0.28	0.43	0.98	1.51
1999–2019	ACM	73.6	0.32	0.62	1.57	2.19
	Kim-Wright	37.0	0.54	0.06	0.24	0.20

Notes: losses are in percent of steady-state consumption. Panel A: x is the quarterly output gap; π is core PCE inflation, measured as annualized quarterly or as 12-month log changes; the loss is $(\varsigma + \varphi)/2 \cdot \text{Var}(x) + \vartheta/2 \cdot \text{E}[(\pi - 2\%)^2]$. SD(x) and RMSE($\pi - 2\%$) are in percentage points. Panel B: SD(TP⁽⁵⁾) is the standard deviation of the monthly 5-year yield term premium (basis points), and κ_p the annualized mean reversion of a monthly AR(1) in TP⁽⁵⁾. The “simple” loss is $\varsigma(\varsigma + \varphi)/(2\varphi) \cdot \text{SD}_i(\check{c})^2$ with $\text{SD}_i(\check{c}) = \text{SD}_i(D_i) \text{SD}(\ell)/(\varsigma(\rho + \kappa_p))$, where $\ell = \text{TP}^{(5)}/2.5$ is the premium per year of duration and $\text{SD}_i(D_i) = 9.0$ years is the cross-sectional standard deviation of household duration exposure per dollar of income (SCF 2022, equity duration $D_{eq} = 0$). The κ_p -consistent loss replaces the linear term $1/2.5$ with the one implied by mean reversion at rate κ_p ; the $D_{eq} = 4$ loss assigns equity a duration of 4 years ($\text{SD}_i(D_i) = 17.0$) in the simple loss.

Panel B reports the risk-sharing term. The key object is $s_0 \times \text{SD}(\tilde{\lambda}_t)$, the cross-household dispersion in the premium earned on household portfolios. We measure this as $\text{SD}_i(D_i) \times \text{SD}(\ell_t)$: the cross-sectional dispersion in household duration exposure $\text{SD}_i(D_i)$, times the volatility of the term premium per year of duration $\text{SD}(\ell_t)$. For the first factor, we use the 2022 Survey of Consumer Finances, assigning simple durations to the main balance-sheet categories (deposits 0, certificates of deposit 1, bonds and bond funds 5, retirement accounts a 40/60 bond-equity mix, mortgages -5 , installment loans -1) and scaling each household’s net duration exposure by its income; with equities assigned zero duration, this calculation implies $\text{SD}_i(D_i) = 9.0$ years. For the second factor, we use the 5-year term premium estimates of [Adrian et al. \(2013\)](#) and [Kim and Wright \(2005\)](#). Because a yield term premium is the average expected excess return over the life of the bond, the premium per year of current duration is $\ell_t = \text{TP}_t^{(5)}/2.5$. The persistence κ_p is the annualized mean reversion of a monthly

AR(1), implying a half-life of 3.3 years (ACM) or 2.5 years (Kim-Wright).

The implied differences in borrowing rates across households are substantial: a cross-household standard deviation of 303 basis points (ACM) or 195 basis points (Kim-Wright), sustained over an effective horizon of $1/(\rho + \kappa_p)$, or 4.3 years (ACM) and 3.3 years (Kim-Wright). Over 1990-2019, the resulting consumption dispersion $SD_i(\check{c})$ is 13.2% and 6.5%, and the risk-sharing term is 1.74 and 0.43 percent of steady-state consumption. This is at least comparable to the conventional terms even at $\vartheta = 100$, and with ACM premia more than five times as large. Restricting the sample to 1999-2019 reduces all loss terms; with the exception of the simple calculations based on Kim-Wright premia, risk-sharing losses are still comparable to conventional terms even with $\vartheta = 100$.

These estimates are conservative in several respects. First, a consistent treatment of mean reversion in the term premium implies a larger premium per year of duration than $TP^{(5)}/2.5$, raising the risk-sharing term to 3.30 (ACM) and 0.98 (Kim-Wright). Second, assigning equities zero duration minimizes household exposure dispersion; with an equity duration of 4 years, the risk-sharing term rises to 6.13 and 1.51. Third, we consider only duration risk; extending the logic to other risky asset classes would likely imply larger risk-sharing motives. Finally, the results are essentially unchanged using the pre-pandemic 2019 wave of the SCF (Table E1 in Appendix E).

5 Optimal Policy

We now turn to the optimal design of short rate and balance sheet policy. Section 5.1 studies simple policy rules, in which the instruments respond only to the natural state variables of the economy; Section 5.2 studies full commitment.

5.1 Simple Policy Rules

For now, we continue to assume that the only aggregate source of uncertainty is natural rate shocks v_t , and that prices are fully rigid. By “simple,” we mean policy which (in equilibrium) is only a function of the “natural” state variables; under the stylized model of Section 3, these are natural rate shocks v_t (equivalently, technology shocks z_t) and the risk-weighted household wealth distribution b_t^σ . Thus, we study policy rules which implement¹²

$$i_t = \chi_{i,v}v_t + \chi_{i,b}b_t^\sigma, \quad qe_t(\tau) = \chi_{q,v}(\tau)v_t + \chi_{q,b}(\tau)b_t^\sigma \implies qe_t^\sigma = \chi_{q,v}^\sigma v_t + \chi_{q,b}^\sigma b_t^\sigma, \quad (39)$$

¹²We assume such policies are implementable. Determinacy conditions can be satisfied through a Taylor rule which generates the unique “minimum state variable” equilibrium as in McCallum (1983).

for some choice of policy parameters $\chi_{i,v}, \chi_{i,b}$ and $\{\chi_{q,v}(\tau), \chi_{q,b}(\tau)\}_{\tau=0}^T$. With a single priced risk factor, and for the same reasons as discussed in Section 2, the impact of balance sheet rules across the term structure is fully summarized by the risk-weighted balance sheet qe_t^σ , up to the holding costs $\psi(\tau)$; the following lemma makes this precise.

Lemma 1 (Optimal QE maturity profile). *Equilibrium dynamics and welfare depend on central bank holdings $\{qe_t(\tau)\}_{\tau=0}^T$ only through the risk-weighted vector qe_t^σ and the holding cost $\int_0^T \psi(\tau) qe_t(\tau)^2 d\tau$. Given qe_t^σ , the optimal portfolio therefore solves a simple minimum cost problem when $\psi(\tau) > 0$:*

$$qe_t(\tau) = \frac{\sigma(\tau) \Omega_\psi^{-1} qe_t^\sigma}{\psi(\tau)}, \quad \Omega_\psi \equiv \int_0^T \frac{\sigma(\tau)^\top \sigma(\tau)}{\psi(\tau)} d\tau, \quad (40)$$

and the flow costs of the policy are $\int_0^T \psi(\tau) qe_t(\tau)^2 d\tau = qe_t^{\sigma\top} \Omega_\psi^{-1} qe_t^\sigma$.

Lemma 1 formalizes the extent of “localization” in our model which the policymaker may exploit. Despite a continuum of instruments, the policymaker can only effectively control the market price of risk, which is low-dimensional because arbitrageurs impose tight restrictions between risk prices and term premia across maturities. In the stylized setting considered here, balance sheet rules work through only a single risk factor.

The policymaker chooses policy parameters to minimize the unconditional social loss $\mathcal{W} \equiv \frac{1}{2} E \int_0^\infty e^{-\rho t} \mathcal{L}_t dt$, where the per-period social loss $\mathcal{L}_t$ is given by (24). Thus, we consider a policymaker who is able to commit to simple policy rules, which are functions of the natural state variables of the economy. In Section 5.2, we explore optimal policy for a policymaker who conducts (history-dependent) policy under full commitment.

Short Rate Only. First, consider the benchmark case of risk-neutral arbitrageurs: $a = 0$. Then the expectations hypothesis holds, so regardless of the asset positions of households and firms (or central bank balance sheet policies), borrowing rates are equalized: $\mu_t(\tau) = i_t$. Moreover, as long as initial household wealth is equalized, the model collapses to the standard RANK case. In particular, consumption and wealth dispersion disappear: $b_t^\sigma = 0$ and $\text{Var}_i \tilde{c}_t(i) = 0$. Because inflation is zero due to rigid prices (more generally: when divine coincidence holds), the policy rule with $\chi_{i,v} = 1$ which implements $i_t = v_t$ achieves the first-best: $x_t = 0$ for all periods t . We further recover the standard QE neutrality result: balance sheet policies do not affect bond prices (and therefore have no aggregate effects).

Under risk neutrality, so long as the central bank faces no costs to setting the policy rate, the first-best is achievable with only conventional policy. The next proposition shows that

these results fail whenever arbitrage is imperfect. We derive optimal short-rate policy when the central bank does not have access to balance sheet tools.

Proposition 6 (Optimal short rate policy rule). *Assume risk aversion $a > 0$, $\eta^\sigma > 0$, and $\beta^\sigma > 0$. If bond risk dispersion across households $\check{\Sigma} = 0$, then $\chi_{i,v}^n = 1 + \frac{a\eta^\sigma\beta^\sigma}{\varsigma\kappa_z} > 1$ and $\chi_{i,b} = 0$ in (39) guarantee $x_t = 0$ each period. The wedge $\chi_{i,v}^n - 1$ reflects the endogenous reaction of firm borrowing to natural rate shocks.*

With $\check{\Sigma} > 0$, the rule which closes the output gap, $(\chi_{i,v}, \chi_{i,b}) = (\chi_{i,v}^n, a\eta^\sigma)$, is not optimal: at leading order in a , marginally reducing $\chi_{i,v}$ below $\chi_{i,v}^n$ (holding $\chi_{i,b} = a\eta^\sigma$) strictly improves welfare. The optimal short rate policy rule implies:

1. *Bond carry trade returns $\mu_t(\tau) - i_t$ move in the same direction as i_t if and only if firm borrowing declines in response to natural rate shocks.*
2. *Output gaps x_t are not identically zero.*
3. *Consumption dispersion is nonzero: $\text{Var}_i \check{c}_t(i) \neq 0$.*

The intuition follows from Section 3. Consider a fall in the natural rate inducing a cut in the policy rate. If firm borrowing declines, then arbitrageurs reduce their position in the bond carry trade, which requires a smaller bond carry trade return. If instead firm borrowing increases, through market clearing, bond arbitrageurs increase their position in the bond carry trade, which requires a larger bond carry trade return.

Because of this, risk premia vary over time. Thus, a simple short rate policy is unable to equalize all borrowing rates. If the risk-weighted dynamics of household wealth b_t^σ are negligible, then the central bank can choose $i_t = \chi_{i,v}^n v_t$, which implies that $\tilde{\mu}_t = v_t$ (and therefore closes output gaps). However, fluctuations in the natural rate induce volatile fluctuations in the policy rate, and the volatility of short rate changes is increasing in the responsiveness of conventional policy to natural rates. More volatile short rates imply greater variation in term premia. Individual Euler equations then differ, so that consumption choices differ across households: $c_t(i) \neq c_t(i')$. Therefore, consumption dispersion $\text{Var}_i \check{c}_t(i) \neq 0$. Thus, from a risk-sharing perspective, choosing a policy rule which is less reactive to natural rate fluctuations is welfare-improving.

However, a less responsive short rate implies that output gap volatility increases. For very small choices of $\chi_{i,v}$, consumption dispersion becomes negligible, but output gap volatility is large. The optimal choice of $\chi_{i,v}^*$ balances these forces. At the optimum, shocks to the natural rate do not fully pass through to the effective borrowing rate: $\tilde{\mu}_t \neq v_t$. Thus, aggregate borrowing demand changes, and hence the output gap $x_t \neq 0$.

Finally, when the risk-weighted dynamics of household wealth b_t^σ are non-negligible, no rule of the form $i_t = \chi_{i,v} v_t$ implies $\tilde{\mu}_t = v_t$. Even conditional on the level of the natural rate, fluctuations in b_t^σ affect term premia. Thus, the optimal policy in this case generally is responsive to risk-weighted household wealth b_t^σ .

Short Rate and Balance Sheets. The failure of conventional policy to achieve the first-best is driven by three frictions. First, because of imperfect pass-through of the policy rate to household borrowing rates, natural rate shocks are not fully accommodated, implying excessive volatility of output. Second, borrowing rates across households differ, implying suboptimal consumption dispersion. Third, the dispersion in consumption decisions induces additional equilibrium movements in relative bond positions of households, leading to further fluctuations in term premia. Simple conventional policy rules cannot overcome all frictions simultaneously. However, when the central bank has access to frictionless balance sheet policies, we obtain the following result.

Proposition 7 (Optimal policy separation principle). *Assume risk aversion $a > 0$, $\eta^\sigma > 0$, and $\beta^\sigma > 0$. Further, assume that household wealth is initially equalized. Suppose the central bank implements short rate and balance sheet policy according to $i_t = v_t$, $qe_t^\sigma = -\beta^\sigma v_t / (\varsigma \kappa_z)$. If short rate policy is frictionless ($\psi^i = 0$) and the central bank does not face holding costs ($\psi(\tau) = 0$), then the first-best is achieved:*

1. *Macroeconomic stabilization: $x_t = 0 \forall t$.*
2. *Term premium stabilization: $\mu_t(\tau) = \tilde{\mu}_t \forall \tau$.*
3. *Consumption equalization: $c_t(i) = c_t(i') \forall i, i'$.*

The results follow naturally from our findings regarding ad hoc policy. The balance sheet policy in Proposition 7 stabilizes shocks to bond markets by offsetting all firm borrowing shocks. The policy implies that risk premia are zero and thus equalizes borrowing rates across households: $\mu_t(\tau) = \tilde{\mu}_t$. When this is the case, the model collapses to a standard RANK model. The elimination of risk premia implies that policy shocks are transmitted one-for-one to borrowing rates, so $\tilde{\mu}_t = i_t$. Thus, the policy $i_t = v_t$ implies output gaps $x_t = 0$, which is optimal (since inflation is zero, or more generally whenever divine coincidence holds). Thus, optimal policy satisfies a *separation principle*: optimal balance sheet policy *stabilizes term premia* while optimal short rate policy *stabilizes macroeconomic aggregates*.

Balance Sheet Only. We now derive the optimal use of balance sheet tools when the central bank faces constraints on the short rate. To capture the essence of short rate constraints in a simple way without introducing an explicit zero lower bound, we take $\psi^i \rightarrow \infty$

and assume a simple process for the “target rate” $\bar{i}_t$ in (17). The next proposition shows that balance sheet tools alone are not enough to achieve the first-best.

Proposition 8 (Optimal balance sheet rule). *Assume risk aversion $a > 0$, $\eta^\sigma > 0$, $\beta^\sigma > 0$, and $\tilde{\Sigma} > 0$. Further, assume that household wealth is initially equalized and that balance sheet policy faces no holding costs ($\psi(\tau) = 0$). Suppose that $\psi^i \rightarrow \infty$ and that $\bar{i}_t = \tilde{\chi}_{i,v} v_t$ where $0 < \tilde{\chi}_{i,v} < 1$. The balance sheet rule which implements $qe_t^\sigma = \beta^\sigma ((1 + \varsigma + \varphi)x_t - v_t/(\varsigma\kappa_z))$ equalizes borrowing rates ($\mu_t(\tau) = \tilde{\mu}_t \forall \tau$). Consumption and wealth dispersion are zero ($\text{Var}_i \tilde{c}_t(i) = 0$ and $b_t^\sigma = 0$), but output gaps x_t are no longer zero.*

There exist balance sheet policy parameters $\chi_{q,v}^\sigma \neq 0, \chi_{q,b}^\sigma \neq 0$ such that $\tilde{\mu}_t = v_t$, which satisfies macroeconomic stabilization ($x_t = 0 \forall t$). However, long-term bond return stabilization fails ($\exists \tau : \mu_t(\tau) \neq \tilde{\mu}_t$), and therefore consumption and wealth dispersion are nonzero. The optimal balance sheet rule implies:

1. *Output gaps x_t are not identically zero. Output gap volatility is lower than under the balance sheet rule which equalizes borrowing rates and, for $a > 0$ sufficiently small, lower than without balance sheet policies.*
2. *Consumption dispersion is nonzero: $\text{Var}_i \tilde{c}_t(i) \neq 0$. Consumption dispersion is lower than under full macroeconomic stabilization via balance sheet tools.*

Proposition 8 shows that when the short rate is constrained, the optimal balance sheet policy must sacrifice term premium stabilization in order to (partially) stabilize macroeconomic volatility. While balance sheet tools can continue to equalize borrowing rates, suboptimal short rate policy implies that the effective household borrowing rate $\tilde{\mu}_t \neq v_t$. Thus, this policy does not achieve macroeconomic stabilization: $x_t \neq 0$. On the other hand, balance sheet policies alone can close the output gap, but this is also suboptimal because borrowing rates are no longer equalized.

It may be surprising that balance sheet tools are insufficient to simultaneously achieve macroeconomic stabilization and bond market stabilization. Although we have eliminated the short rate as an active policy tool, we have given the policymaker a continuum of balance sheet tools $\{qe_t(\tau)\}_{\tau=0}^T$. Textbook intuition suggests that the policymaker has enough tools to deal with the two objectives.

The reason for the failure of textbook intuition is that balance sheet policy works by affecting term premia through changes in the market price of risk. Although arbitrage is imperfect in this model, arbitrageurs still enforce tight restrictions between the market price of risk and term premia across maturities. Hence, while in principle the central bank has a continuum of policy tools, in practice it can only manipulate λ_t (which we have already imposed by working with the risk-weighted balance sheet object qe_t^σ).

This is related to the localization results in [Vayanos and Vila \(2021\)](#) and [Ray et al. \(2024\)](#). In the one-factor model considered here, the effects of QE are fully global. With several risk factors, the same conclusion holds whenever localization is not “strong enough” in the sense that no direction of the price of risk moves the aggregate household borrowing rate without also moving relative borrowing rates.

Without Divine Coincidence. Without policy frictions, the optimal separation principle of Proposition 7 shows that when divine coincidence holds, the central bank can achieve the first-best. We now study policy when divine coincidence fails. We maintain the stylized assumptions of the model so far, except that we relax the assumption of fully rigid prices, and we allow for nonzero markup shocks $u_t = u_z z_t$ (for simplicity, we maintain the one-factor assumption, so that markup shocks load on the single source of aggregate risk). In this version of the model, the augmented Phillips curve (22) applies, and

$$E_t dx_t = \varsigma^{-1} [i_t + a\eta^\sigma \Lambda_t - \pi_t - v_t] dt, \quad E_t d\pi_t = [\rho\pi_t - \kappa(x_t + \alpha_{\hat{\mu}}\hat{\mu}_t) - u_t] dt,$$

where the firm effective borrowing rate is $\hat{\mu}_t = \beta i_t + a\beta^\sigma \Lambda_t$. All other variables and endogenous coefficients remain the same. Define the firm-household effective-rate wedge $\Delta^\sigma \equiv \beta^\sigma - \beta\eta^\sigma$, which measures the relevant distance between the firm effective borrowing rate and the aggregate household effective borrowing rate. When $\Delta^\sigma \neq 0$, the central bank can in principle use short rate and balance sheet tools to target these borrowing rates separately. We explore the feasibility and the welfare costs of simultaneously stabilizing the output gap and inflation in the next proposition.

Proposition 9 (Exact macro stabilization). *Suppose that $a > 0, \beta > 0, \beta^\sigma > 0, \eta^\sigma > 0$.*

1. *Whenever $\Delta^\sigma \neq 0$, exact macro stabilization $x_t \equiv \pi_t \equiv 0$ is an equilibrium under*

$$\Lambda_t = -\frac{\tilde{u}_t}{a\Delta^\sigma}, \quad i_t = v_t + \frac{\eta^\sigma}{\Delta^\sigma} \tilde{u}_t, \quad qe_t^\sigma = -\frac{\beta^\sigma}{\varsigma\kappa_z} v_t - b_t^\sigma - \Lambda_t, \quad (41)$$

where $\tilde{u}_t = \beta v_t + \frac{\vartheta}{\epsilon} u_t$.

2. *If $\check{\Sigma} > 0$ and $\tilde{u}_t \not\equiv 0$, then, holding the composites $(\eta^\sigma, \beta^\sigma, \check{\Sigma})$ fixed, as $\Delta^\sigma \rightarrow 0$ the social welfare costs of (41) become unbounded, and policy tools diverge.*

The presence of a working capital channel in the Phillips curve endogenously generates the failure of divine coincidence. But when firm and household borrowing costs differ, policymakers may be tempted to use a combination of short rate and balance sheet tools to separately target these effective rates. With two policy tools this is possible: the policymaker

manipulates the short rate and term premia to set both the household and firm effective rates to any desired target. This allows the central bank to offset the endogenous “cost channel” of conventional monetary policy, and in fact allows for a full offset of any shocks to the Phillips curve. But such a policy comes at a price: because the policy necessarily involves manipulating risk prices, it worsens risk-sharing. In the limit as household and firm effective rates become more similar, the required policy leads to unboundedly large fluctuations in term premia, and therefore the welfare costs of such a policy themselves become unbounded.

5.2 Full Commitment

We now study full commitment: the policymaker chooses short rate and balance sheet policies as functions of the entire history of shocks. The separation principle of Proposition 7 shows that simple policy rules can achieve the first-best when the policymaker faces no short rate or balance sheet frictions and divine coincidence holds. Outside of these cases, full commitment policies can improve on simple policy rules.

We consider the problem of a planner with full commitment minimizing unconditional social loss $\mathcal{W}$ (see, e.g., [Damjanovic et al. 2008](#); Appendix C considers the case of the conditional loss function $E_t \int_t^\infty \frac{1}{2} e^{-\rho(s-t)} \mathcal{L}_s ds$). The problem is unrestricted, except that the policymaker is confined to the class of policies with time-invariant feedback on innovations, which deliver time-invariant bond loadings $\sigma(\tau)$. We therefore extend the state with a set of *policy memory states* m_t :

$$i_t = \phi_i^\top s_t + \phi_{im}^\top m_t + \epsilon_{i,t}, \quad qe_t^\sigma = \Phi_q s_t + \Phi_{qm} m_t + \epsilon_{q,t}, \quad (42)$$

$$\text{where } dm_t = -[\Gamma_m m_t + \Gamma_{ms} s_t] dt + K_m dB_t, \quad (43)$$

with every object (including the dimension of m_t) a choice of the planner, along with the deterministic paths $(\epsilon_{i,t}, \epsilon_{q,t})$. Bond prices are then affine in the extended state $s_t^+ \equiv (s_t^\top, m_t^\top)^\top$ plus a deterministic intercept. Crucially, all elements of the extended state s_t^+ are priced risk factors; the deterministic paths move expected returns without affecting loadings.

Two results characterize the optimal plan. First, holding fixed the endogenous risk-weighted composites $(\eta^\sigma, \beta^\sigma, \tilde{\Sigma}, \Omega_\psi)$, the planner problem is a standard linear-quadratic program, and the solution takes the form of “promise-keeping” costates. The costates associated with the output gap x_t , inflation π_t , and the consumption dispersion vector $\check{c}_t^\sigma$ are given by

$$\begin{aligned} d\theta_{x,t} &= [-(\varsigma + \varphi)x_t + \kappa\theta_{\pi,t} - (1 + \varsigma + \varphi)\beta^\sigma\theta_{\Lambda,t}] dt, \\ d\theta_{\pi,t} &= [-\rho\theta_{\pi,t} - \vartheta\pi_t + \varsigma^{-1}\theta_{x,t}] dt, \quad d\theta_{c,t} = \left[-\frac{\varsigma(\varsigma + \varphi)}{\varphi} \tilde{\Sigma}^+ \check{c}_t^\sigma + (1 + \varsigma/\varphi)\theta_{b,t} \right] dt, \end{aligned} \quad (44)$$

where $\theta_{\Lambda,t}$ is the shadow price of risk and $\theta_{b,t}$ is the costate on the wealth statistics b_t^σ (for the formal results, see Lemma A2).

The second result removes the restriction to fixed composites, and takes into account that the fully optimal planner internalizes that the bond loadings $\sigma(\tau)$ (and hence risk transmission, portfolio dispersion, and balance sheet costs) themselves depend on policy.

Theorem 3 (Full commitment optimal policy, necessary conditions). *Assume $\psi^i > 0$ and, if balance sheet policy is used, Ω_ψ nonsingular. Denote by $\nu(\tau) \in \mathbb{R}^{N_z}$ the multipliers on the endogenous loadings $\sigma(\tau)$, and define the shadow factors*

$$v_t^i \equiv \int_0^\infty \nu_i(u)^\top dB_{t-u}, \quad v_t^\Lambda \equiv \int_0^\infty \nu_\Lambda(u)^\top dB_{t-u}, \quad (45)$$

$$\nu_i(u) \equiv \int_u^T \nu(\tau) d\tau, \quad \nu_\Lambda(u) \equiv a \int_u^T \nu(\tau) \sigma(\tau - u) d\tau. \quad (46)$$

Under standard regularity conditions, any optimal plan satisfies the system of Lemma A2 with $\theta_{\Lambda,t}$ replaced by $\theta_{\Lambda,t} - v_t^\Lambda$ and $-v_t^i$ added to the short-rate condition.

Let $c \in \mathbb{R}^K$ collect a minimal set of coordinates for the endogenous risk-weighted composites. The multiplier function is spanned by K basis functions, $\nu(\tau) = \sum_{k=1}^K \zeta_k \nu^{(k)}(\tau)$ with $\zeta_k = \frac{\partial J}{\partial c_k}$, where $J(c)$ denotes the optimal value of the planner's problem (including the shadow-factor terms, with the shadow factors v^i, v^Λ held fixed) when the composites are fixed at c . Hence, given the plan's risk-price impulse responses, the multiplier function is finite-dimensional. When the optimal plan admits a finite-state realization through policy memory states as in (43), any optimal plan solves a finite-dimensional fixed point problem (solutions are ranked by value): the $2K$ coefficients (c, ζ) together with the pricing matrix M of Proposition 2 for the extended state. The optimal instruments take the form

$$\begin{aligned} i_t &= \epsilon_{i,t} + \phi_i^\top s_t + \phi_{i\theta}^\top \theta_t + \int_0^\infty f_i(u)^\top dB_{t-u}, \\ qe_t^\sigma &= \epsilon_{q,t} + \Phi_q s_t + \Phi_{q\theta} \theta_t + \int_0^\infty F_q(u) dB_{t-u}, \end{aligned} \quad (47)$$

with convolution kernels f_i, F_q linear in ζ (construction is in Appendix C).

The policy characterized by (44) (Lemma A2) and the policy of (47) (Theorem 3) correspond to two notions of commitment policy, which differ in whether the planner internalizes that bond loadings depend on policy. We call the textbook plan of Lemma A2, evaluated at the composites which the plan itself induces, the *Nash plan*: the planner optimizes taking the exposures of bond portfolios to aggregate risk as given, and in equilibrium these exposures are the ones generated by the plan. The plan of Theorem 3 is the *Stackelberg plan*:

the planner acts as a leader in the bond market, and chooses its responses to shocks taking into account how they shape the loadings $\sigma(\tau)$. Comparing the Stackelberg with the Nash plan isolates the value of this additional margin. Both plans are straightforward to compute as fixed points in finitely many numbers, with the Nash plan the special case in which the marginal values of the composites ζ are set to zero (Appendix C).

Both plans share the familiar benefit of commitment. As in standard New Keynesian commitment problems, promises induce inertial, “price-level-targeting”-style policy responses, which stabilize output and inflation through expected future policy rather than through large movements in current rates. Here, inertia has an additional benefit: less reactive policy implies *ex ante* smaller expected variation in policy tools, which in equilibrium reduces endogenous bond volatility and hence the risk compensation arbitrageurs require. Thus, inertial policy additionally improves risk-sharing directly.

The Stackelberg plan has a second, distinct *risk-shaping* motive. Households with different portfolios face different returns, and how much these returns differ depends on two objects: the price of risk Λ_t , and the loadings $\sigma(\tau)$, which determine how exposed each portfolio is to aggregate shocks (through the composites η^σ and $\tilde{\Sigma}$). Simple rules act on the loadings only indirectly, since their responses to shocks inherit the persistence of the natural states, and the Nash plan takes the loadings as given. But the loadings are themselves equilibrium outcomes. Bond prices reflect expectations of the entire future path of the short rate and of risk compensation, so the loading of a τ -period bond on a shock is determined by the cumulated response of these two objects (equation (C2) in Appendix C). Every portfolio’s exposure therefore depends on a weighted average of the impulse responses of policy, with weights that depend on the maturity structure of the portfolio. A planner who commits to the “shape” of these responses (how policy moves at short and long horizons) can reduce the exposure of household portfolios to aggregate risk, even when risk prices are time-varying.

Theorem 3 shows how this motive changes the instruments. The instruments under the Stackelberg plan satisfy

$$\psi^i(i_t - \bar{i}_t) = \kappa \alpha_{\hat{\mu}} \beta \theta_{\pi,t} - \varsigma^{-1} \theta_{x,t} + v_t^i, \quad qe_t^\sigma = \Omega_\psi (\theta_{\Lambda,t} - v_t^\Lambda). \quad (48)$$

The Nash plan satisfies the same conditions without the shadow factors v_t^i and v_t^Λ (equation (A10) in Appendix A). The paths of the promise states also differ between the plans, since $\theta_{\Lambda,t} - v_t^\Lambda$ replaces $\theta_{\Lambda,t}$ in their dynamics, but under both plans they have no diffusion. Under the Nash plan the instruments therefore respond to innovations only gradually, through the promise states. Under the Stackelberg plan they also respond through the shadow factors, which load on current innovations: on impact, the short rate responds to dB_t by $\nu_i(0)^\top / \psi^i$

and the balance sheet by $-\Omega_\psi \nu_\Lambda(0)^\top$. Because the impact response is set by the shadow factors and the later path by the promise states, the Stackelberg plan may move policy early in one direction and later in the other.

The shadow factors v_t^i and v_t^Λ in Theorem 3 capture exactly the risk-shaping motive: the planner values the impulse responses of the short rate and of risk prices not only for their direct effects on the loss, but also because they determine the endogenous loadings. The Stackelberg planner minimizes $\frac{1}{2} \mathbb{E}[\mathcal{L}_t] - \mathbb{E}[i_t v_t^i + \Lambda_t^\top v_t^\Lambda]$, and by (C2) the second term is

$$- \mathbb{E}[i_t v_t^i + \Lambda_t^\top v_t^\Lambda] = \int_0^T \sigma(\tau) \nu(\tau) d\tau, \quad (49)$$

the loadings generated by the plan, each weighted by its multiplier $\nu(\tau)$. By (C3), $\nu(\tau)$ is the marginal value of the τ -bond's loading through the composites (such as η^σ , β^σ and $\check{\Sigma}$), net of its effect on the premia of longer bonds. At the optimum, any change in how policy responds to shocks therefore changes this term by exactly the value of the change in portfolio exposures that it causes. Because the response of the short rate at horizon u is weighted by $\nu_i(u) = \int_u^T \nu(\tau) d\tau$, the sum of the multipliers on all bonds still outstanding at that horizon, the planner cares about the shape of the whole path of policy, not only its size. The Nash planner leaves this term out.

Risk-shaping and balance sheet policy are substitutes. Balance sheet policy acts on the price of risk directly (Proposition 7); when balance sheet frictions are small, the planner can remove the persistent component of the price of risk which generates consumption dispersion, even while using the balance sheet to move term premia at shorter horizons for macroeconomic stabilization, and there is little left to gain from reshaping exposures. When balance sheet policy is costly or unavailable, reshaping exposures through the path of the short rate is the planner's main tool against consumption dispersion.

6 Quantitative Analysis

We now quantify the optimal policy mechanisms in a calibrated version of the model. Relative to the stylized setting studied in Sections 3-5, the quantitative model has sticky prices, a working capital channel, and three sources of aggregate risk: household discount-factor (demand) shocks ϱ_t , markup (cost-push) shocks u_t , and bond supply (or equivalently, bond-noise-demand) shocks. We compare the historical Taylor rule, alternative Taylor-type rules, the optimal simple rule of Section 5.1, and the two full-commitment plans of Section 5.2, with and without balance sheet policy.

6.1 Calibration

Table 2 lists the parameters. The macroeconomic block uses standard New Keynesian values; we set the price-adjustment cost to $\vartheta = 100$ (the upper end of the range considered in Section 4.1). Firms finance three quarters of the wage bill in advance, through a borrowing portfolio with a mean maturity of five years. We assume a Taylor rule $i_t = 1.5\pi_t + 0.5x_t$.

For households we use the simplest representation which delivers the rank-one dispersion matrix $\tilde{\Sigma} = s_0^2(\eta^\sigma)^\top \eta^\sigma$ of Lemma A1: half of all households hold $1 - s_0$ and half hold $1 + s_0$ units of a common bond portfolio $\eta(\tau) = \delta_h e^{-\delta_h \tau}$, with the balance in short-term bonds. The mean maturity of the portfolio is $D_h = 1/\delta_h = 5$ years, and we set $s_0 = 1.8$, so that the cross-sectional standard deviation of household duration exposure, $s_0 D_h$, equals the 9 years measured in the SCF in Section 4.1. Arbitrageurs absorb the net positions of households and firms and the exogenous bond supply, with mean maturity 5.4 years.¹³

Five parameters are calibrated internally: the volatilities of the demand and markup shocks, the correlation of demand and bond-supply shocks, the slope of the Phillips curve, and arbitrageur risk aversion a . The targets are the volatility of the output gap and inflation, the correlation of the output gap and inflation, the volatility of the term premium on the household portfolio, and the correlation between 12-month changes in the short rate and in that premium.¹⁴ The premium target is the 84 basis point standard deviation of the ACM five-year term premium used in Table 1. The model’s household premium is an expected excess return, whereas the ACM estimate is a yield term premium, an average of expected excess returns over the life of the bond; the conversion of Section 4.1 would assign the five-year household portfolio a premium about twice as volatile. Targeting the yield premium is therefore conservative by about a factor of two in the premium and four in the dispersion term. Table E2 shows that the model matches the targets, reports untargeted moments, and shows that the risk-sharing motive is sizable but conservative: under the Taylor rule, consumption dispersion accounts for roughly 60 percent of the social loss.

6.2 Welfare under Rules and Plans

Table 3 reports the social loss, by term, under the policies we consider. In every row the central bank faces the same cost of moving the short rate, $\psi^i i_t^2$; we choose $\psi^i = 1.87$ so that

¹³We assume all portfolios have exponential maturity weights and that $T \rightarrow \infty$, in order to deliver the exact finite-state representation from Theorem 3.

¹⁴The macroeconomic moments are computed from the CBO output gap and year-over-year core PCE inflation over 1999-2019, in order to exclude the disinflation of the early 1990s; the premium moments use the full 1990-2019 sample.

Table 2: Calibration

<i>Panel A: External</i>		
ρ	discount rate	0.02
ς	inverse elasticity of intertemporal substitution	1
φ	inverse Frisch elasticity	1
ϑ	price-adjustment cost (loss weight on inflation)	100
β	share of the wage bill financed by borrowing	0.75
D_f	maturity of firm borrowing (years)	5
$D_h = 1/\delta_h$	maturity of the household bond portfolio (years)	5
s_0	dispersion of household duration exposure	1.8
κ_ϱ	mean reversion, demand factor	0.25
κ_u	mean reversion, cost-push factor	0.85
κ_s	mean reversion, bond-supply factor	0.20
$1/\delta_s$	maturity of bond supply (years)	5.4
ζ_s	scale of bond supply	3
ϕ_π	historical rule: response to inflation	1.5
ϕ_x	historical rule: response to the output gap	0.5
<i>Panel B: Internal</i>		
σ_ϱ	volatility, demand innovation	0.0333
σ_u	volatility, cost-push innovation	0.00334
$\rho_{\varrho s}$	correlation, demand and supply innovations	-0.079
κ	slope of the Phillips curve	0.0319
a	arbitrageur risk aversion	6.63

Notes: annual frequency. Panel A lists externally calibrated parameters. The parameters in Panel B are calibrated internally to the targets in Table E2.

the cost of rate movements is moderate.¹⁵ Where balance sheet policy is available, its cost is a fixed quadratic function of the risk-weighted holdings qe_t^σ . By Lemma 1, this is the cost of the cheapest portfolio which delivers qe_t^σ when Ω_ψ is held at a reference value; holding it fixed keeps the cost invariant to the loadings.

Table 3 yields three findings. First, Taylor-type rules which respond more aggressively to inflation or to the output gap reduce macroeconomic volatility but raise term premium volatility and consumption dispersion, and are worse than the historical rule (losses of 0.36 to 0.47 percent of consumption, compared with 0.31). The optimal simple short-rate-only rule improves on the historical rule only modestly (0.27 versus 0.31); it does so mainly by responding much less to markup shocks. Thus, the under-reaction of Proposition 6 carries

¹⁵The baseline value of ψ^i is set so that the rate-only Stackelberg plan generates a 1 percentage point response of the short rate to an aggregate demand shock. Appendix Table E4 reports total losses, and Table E5 their decomposition, for a range of rate and balance sheet costs; Figures E3 to E6 show the corresponding impulse responses, and Table E3 reports the Taylor-type rules.

Table 3: Social Loss under Rules and Plans

	loss (% of consumption)						SD(<i>i</i>)	prem.	<i>u</i> -resp.
	total	<i>x</i>	π	dispersion	rate	QE			
<i>Panel A: Taylor-type rules and the optimal simple rule</i>									
Historical Taylor rule	0.309	0.034	0.071	0.188	0.016	–	1.31	84	1.52
Taylor, $\phi_{\pi} = 3$	0.361	0.028	0.062	0.248	0.022	–	1.55	100	2.54
Gap-stabilizing Taylor	0.357	0.000	0.046	0.283	0.028	–	1.73	109	1.17
Inflation-stabilizing Taylor	0.469	0.028	0.055	0.353	0.033	–	1.88	126	3.82
Natural-rate rule, $\chi_{i,v} = 1$	0.406	0.021	0.051	0.310	0.024	–	1.60	116	1.52
Optimal simple rule	0.266	0.033	0.079	0.141	0.014	–	1.22	70	0.22
<i>Panel B: Full commitment, short rate only</i>									
Nash plan	0.110	0.008	0.023	0.072	0.008	–	0.93	46	–0.23
Stackelberg plan	0.047	0.015	0.023	0.002	0.007	–	0.84	6	–0.10
<i>Panel C: Short rate and balance sheet</i>									
Natural-rate rule, $\Lambda = 0$	0.061	0.000	0.045	0.000	0.016	0	1.30	0	1.52
Historical rule, $\Lambda = 0$	0.098	0.021	0.066	0.000	0.011	0	1.10	0	1.52
Nash plan, moderate QE cost	0.057	0.008	0.023	0.007	0.007	0.013	0.85	31	–0.23
Stackelberg plan, moderate QE cost	0.044	0.012	0.023	0.002	0.006	0.002	0.79	10	–0.13
Nash plan, low QE cost	0.038	0.007	0.023	<0.001	0.006	0.002	0.81	17	–0.13
Stackelberg plan, low QE cost	0.038	0.008	0.023	<0.001	0.006	0.001	0.78	15	–0.11

Notes: the loss $\frac{1}{2} E[\mathcal{L}_t]$ and its terms, in percent of steady-state consumption; “rate” and “QE” are the costs of the two instruments. All rows face the same rate cost $\psi^i i_t^2$ with $\psi^i = 1.87$. The Nash plan is the commitment plan under which the central bank takes household portfolio exposures to aggregate risk as given; the Stackelberg plan internalizes the effect of policy on those exposures (Section 5.2). SD(i) is in percent; “prem.” is the standard deviation of the household term premium, in basis points; “ u -resp.” is the coefficient of the short rate on the markup factor u_t . Panel C: “ $\Lambda = 0$ ” rows combine the short-rate rule with costless balance sheet policy which holds the price of risk at zero; moderate and low QE costs are 1/10 and 1/240 of the reference cost at which balance sheet policy roughly halves consumption dispersion under the Nash plan. Taylor-type rules are $i_t = \phi_\pi \pi_t + \phi_x x_t$ with $(\phi_\pi, \phi_x) = (1.5, 0.5)$ (historical), $(3, 0.5)$, $(1.5, 500)$ (gap-stabilizing) and $(5, 0.5)$ (inflation-stabilizing); the natural-rate rule adds $\chi_{i,v} v_t$ to the historical rule.

over to the quantitative model, but applies mainly to cost-push shocks, where the working capital channel makes an aggressive response most costly in terms of risk-sharing.

Second, commitment is valuable, and the risk-shaping motive of Section 5.2 accounts for about a quarter of this value at our baseline rate cost. The Nash plan removes about two thirds of the historical rule’s loss (0.11 versus 0.31), more than half of it through lower consumption dispersion (0.072 versus 0.188), which nonetheless remains the largest loss term. The Stackelberg plan more than halves the loss that remains under the Nash plan (0.047 versus 0.110), and removes more than 80 percent of the loss relative to the historical rule. Across the rate costs in Table E4, the Stackelberg plan removes between 84 and 88 percent of the historical rule’s loss, while the gap between the two plans falls from 0.17 to 0.004 percent of consumption as moving the short rate becomes more costly. The difference between the two plans comes from consumption dispersion (0.002 versus 0.072), which the Stackelberg plan all but eliminates while accepting a slightly larger output gap term. The

two plans achieve similar macroeconomic outcomes, but the Stackelberg plan also chooses how exposed household portfolios are to aggregate risk.

Third, balance sheet policy is valuable under every planner. Even as a simple rule, balance sheet policy which holds the price of risk at zero removes consumption dispersion entirely (Proposition 7); combined with the natural-rate rule, what remains is mostly the inflation loss, since divine coincidence fails in the quantitative model through the working capital channel and markup shocks. The Stackelberg plan without balance sheet policy does even better than these rules (0.047, compared with 0.061 and 0.098): in the calibrated model, shaping exposures through the path of the short rate is a close substitute for holding the price of risk at zero. With a low balance sheet cost, the Nash and Stackelberg plans achieve the same loss (0.038): when risk prices can be stabilized directly, there is little left for risk-shaping to do.

6.3 Dynamics under Optimal Policy

Figure 2 shows impulse responses under the historical rule and the two commitment plans, when the central bank can use both the short rate and its balance sheet (with the low balance sheet cost of Table 3).

After a bond supply shock, the responses follow the separation principle of Proposition 7; after demand and markup shocks, they show its limits when divine coincidence fails. Under both plans, the short rate follows the familiar commitment path: responding gradually rather than on impact, and stabilizing the output gap and inflation through expected future policy rather than through large changes in current rates. The balance sheet, instead, is used for managing bond risk. After a bond supply shock the household premium rises sharply under the historical rule and stays elevated for years, depressing output; under both plans balance sheet purchases build over the first year, eliminating the premium with small effects on the output gap or the short rate.

The balance sheet is not only used to stabilize the bond market, however. After a demand shock, the plans use balance sheet policy to *raise* the household premium over intermediate horizons: a higher premium raises households' effective borrowing rate and so tightens financial conditions, complementing the gradual increase in the short rate. After a markup shock the balance sheet eases slightly instead. The last row of the figure shows that these movements in the premium are due to the balance sheet response. Balance sheet policy is therefore partly a macroeconomic stabilization tool, and the plans do not aim to hold the premium constant, even when the balance sheet is cheap.

The Nash and Stackelberg plans are similar, consistent with Table 3. They differ mainly in




Figure 2: Impulse Responses under Optimal Policy

Notes: responses to one-standard-deviation demand, markup (cost-push), and bond supply shocks, under the historical rule and the Nash and Stackelberg plans with rate cost $\psi^i i_t^2$ and the low balance sheet cost of Table 3. The last row shows the contribution of the balance sheet response to the household premium, $-a\eta^\sigma qe_t^\sigma$ (the part of the premium in the row above which is due to the balance sheet).

the first months after a shock. Under the Nash plan the instruments do not move on impact; the Stackelberg plan moves both, and after a demand shock it briefly cuts the short rate and lowers the premium before following the gradual tightening of the Nash plan. These early movements reflect the risk-shaping motive: they reduce the exposure of household portfolios

to the shock, rather than the price of risk.

The Role of Policy Costs. Appendix E reports losses and impulse responses for a range of costs of the two instruments. First, without balance sheet policy the Stackelberg plan improves risk-sharing through the path of the short rate (Figure E3). After a demand shock it combines a short-lived cut in the short rate with the gradual tightening of the Nash plan. Because the returns of long-term bonds depend on the whole expected path of the short rate, this makes short-maturity bonds gain and long-maturity bonds lose, so that the household portfolio as a whole is much less exposed to the shock (Table E5).

Second, when moving the short rate is costly, the Nash plan becomes much less reactive; its loadings and risk prices both shrink, and the gap between the Nash and Stackelberg plans closes even without balance sheet policy, at the cost of larger output gap fluctuations. Third, when moving the short rate is costly but balance sheet policy is cheap, the balance sheet takes on a larger role in macroeconomic stabilization (Figure E5). After a demand shock, the balance sheet is used to raise the household premium for several years to tighten financial conditions, and premium volatility is higher than without balance sheet policy.

7 Concluding Remarks

This paper develops a tractable general equilibrium model with market segmentation, nominal frictions, and household heterogeneity. We first derive results for how monetary policy (conventional and unconventional) affects aggregate dynamics. We next show that welfare losses arise from excessive volatility in inflation, output, and consumption dispersion across households. The frictions associated with consumption dispersion arise due to market segmentation and excessive volatility of term premia. When neither instrument is costly and divine coincidence holds, optimal policy satisfies a separation principle: conventional policy stabilizes output and inflation, while balance sheet policy stabilizes term premia. Policy constraints on either instrument imply tradeoffs between these two policy objectives. Under full commitment, policy can also shape the exposure of household portfolios to aggregate risk; in our calibration this motive accounts for about a quarter of the gains from commitment when balance sheet policy is unavailable. For a central bank unable to condition policy on the structure of household portfolios and on the history of shocks, simple balance sheet rules which reduce term premium volatility can achieve most of the same gains.

References

- Acharya, S., Challe, E., and Dogra, K. (2023). Optimal monetary policy according to HANK. *American Economic Review*, 113(7):1741–1782.
- Adrian, T., Crump, R. K., and Moench, E. (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics*, 110(1):110–138.
- Alexandrov, A. and Brunnermeier, M. (2026). Optimal (un)conventional monetary policy. Working paper.
- Andrés, J., López-Salido, J. D., and Nelson, E. (2004). Tobin’s Imperfect Asset Substitution in Optimizing General Equilibrium. *Journal of Money, Credit and Banking*, 36(4):665–690.
- Ang, A. and Piazzesi, M. (2003). A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables. *Journal of Monetary Economics*, 50(4):745–787.
- Angeletos, G.-M., Collard, F., and Dellas, H. (2023). Public Debt as Private Liquidity: Optimal Policy. *Journal of Political Economy*, 131(11):3233–3264.
- Auclert, A. (2019). Monetary Policy and the Redistribution Channel. *American Economic Review*, 109(6):2333–67.
- Backus, D. and Driffill, J. (1986). The consistency of optimal policy in stochastic rational expectations models. CEPR Discussion Papers 124, Centre for Economic Policy Research.
- Bauer, M. D. and Swanson, E. T. (2023). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37(1):87–155.
- Bhandari, A., Evans, D., Golosov, M., and Sargent, T. J. (2021). Inequality, business cycles, and monetary-fiscal policy. *Econometrica*, 89(6):2559–2599.
- Borovička, J. and Hansen, L. P. (2014). Robust preference expansions. 2013 Meeting Papers 1199, Society for Economic Dynamics.
- Bretscher, L., Schmid, L., Sen, I., and Sharma, V. (2022). Institutional Corporate Bond Pricing. Swiss Finance Institute Research Paper 21-07, Swiss Finance Institute.
- Caballero, R. J., Caravello, T. E., and Simsek, A. (2024). Financial Conditions Targeting. NBER Working Papers 33206, National Bureau of Economic Research.
- Caballero, R. J. and Simsek, A. (2024). Central banks, stock markets, and the real economy. *Annual Review of Financial Economics*, 16:179–205.
- Cahill, M. E., D’Amico, S., Li, C., and Sears, J. S. (2013). Duration Risk versus Local Supply Channel in Treasury Yields: Evidence from the Federal Reserve’s Asset Purchase Announcements. FEDS 2013-35, Board of Governors.

- Campbell, J. R. and Nemtyrev, A. (2026). FOMC Responses to Credit Spreads. Technical report.
- Carlstrom, C. T., Fuerst, T. S., and Paustian, M. (2017). Targeting Long Rates in a Model with Segmented Markets. *American Economic Journal: Macroeconomics*, 9(1):205–242.
- Chen, H., Cúrdia, V., and Ferrero, A. (2012). The Macroeconomic Effects of Large-scale Asset Purchase Programmes. *Economic Journal*, 122(564):289–315.
- Cui, W. and Sterk, V. (2021). Quantitative Easing with Heterogeneous Agents. *Journal of Monetary Economics*, 123(C):68–90.
- Cúrdia, V. and Woodford, M. (2011). The Central-Bank Balance Sheet as an Instrument of Monetary Policy. *Journal of Monetary Economics*, 58(1):54–79.
- D’Amico, S. and King, T. B. (2013). Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply. *Journal of Financial Economics*, 108(2):425–448.
- Damjanovic, T., Damjanovic, V., and Nolan, C. (2008). Unconditionally optimal monetary policy. *Journal of Monetary Economics*, 55(3):491–500.
- d’Avernas, A., de Fraisse, A. H., Ning, L., and Vandeweyer, Q. (2024). How Large Is Too Large? A Risk-Benefit Framework for Quantitative Easing. Technical report, Chicago Booth Research Paper.
- Debortoli, D., Galí, J., and Gambetti, L. (2020). On the Empirical (Ir)Relevance of the Zero Lower Bound Constraint. *NBER Macroeconomics Annual*, 34:141–170.
- Debortoli, D. and Galí, J. (2017). Monetary Policy with Heterogeneous Agents: Insights from TANK Models. Economics Working Papers 1686, Department of Economics and Business, Universitat Pompeu Fabra.
- Di Maggio, M., Kermani, A., and Palmer, C. J. (2020). How Quantitative Easing Works: Evidence on the Refinancing Channel. *Review of Economic Studies*, 87(3):1498–1528.
- Ellison, M. and Tischbirek, A. (2014). Unconventional Government Debt Purchases as a Supplement to Conventional Monetary Policy. *Journal of Economic Dynamics and Control*, 43:199–217.
- Federal Open Market Committee (2022). Principles for reducing the size of the Federal Reserve’s balance sheet. Press release, January 26, 2022.
- Fieldhouse, A. J., Mertens, K., and Ravn, M. O. (2018). The Macroeconomic Effects of Government Asset Purchases: Evidence from Postwar U.S. Housing Credit Policy. *Quarterly Journal of Economics*, 133(3):1503–1560.
- Gertler, M. and Karadi, P. (2011). A Model of Unconventional Monetary Policy. *Journal of Monetary Economics*, 58(1):17 – 34. Carnegie-Rochester Conference Series on Public Policy: The Future of Central Banking April 16-17, 2010.

- Gertler, M. and Karadi, P. (2013). QE 1 vs. 2 vs. 3. . . : A Framework for Analyzing Large-Scale Asset Purchases as a Monetary Policy Tool. *International Journal of Central Banking*, 9(1):5–53.
- Gourinchas, P.-O., Ray, W., and Vayanos, D. (2025). A Preferred-Habitat Model of Term Premia, Exchange Rates, and Monetary Policy Spillovers. *American Economic Review*, 115(11):3788–3824.
- Greenwood, R., Hanson, S., Stein, J. C., and Sunderam, A. (2023). A Quantity-Driven Theory of Term Premia and Exchange Rates. *The Quarterly Journal of Economics*, 138(4):2327–2389.
- Greenwood, R., Hanson, S. G., and Vayanos, D. (2016). Forward Guidance in the Yield Curve: Short Rates versus Bond Supply. In Albagli, E., Saravia, D., and Woodford, M., editors, *Monetary Policy through Asset Markets: Lessons from Unconventional Measures and Implications for an Integrated World*. Central Bank of Chile.
- Greenwood, R. and Vayanos, D. (2014). Bond Supply and Excess Bond Returns. *Review of Financial Studies*, 27(3):663–713.
- Gürkaynak, R. S., Sack, B., and Swanson, E. (2005). The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. *The American Economic Review*, 95(1):425–436.
- Gürkaynak, R. S., Sack, B., and Wright, J. H. (2007). The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, 54(8):2291 – 2304.
- Hamilton, J. D. and Wu, J. C. (2012). The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment. *Journal of Money, Credit and Banking*, 44:3–46.
- Hanson, S. G., Lucca, D. O., and Wright, J. H. (2021). Rate-amplifying demand and the excess sensitivity of long-term rates. *The Quarterly Journal of Economics*, 136(3):1719–1781.
- Hanson, S. G. and Stein, J. C. (2015). Monetary policy and long-term real rates. *Journal of Financial Economics*, 115(3):429–448.
- Itskhoki, O. and Mukhin, D. (2023). Optimal Exchange Rate Policy. Working paper.
- Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary Policy According to HANK. *American Economic Review*, 108(3):697–743.
- Karadi, P. and Nakov, A. (2021). Effectiveness and Addictiveness of Quantitative Easing. *Journal of Monetary Economics*, 117:1096–1117.
- Kekre, R. and Lenel, M. (2022). Monetary Policy, Redistribution, and Risk Premia. *Econometrica*, 90(5):2249–2282.

- Kekre, R. and Lenel, M. (2024). Exchange rates, natural rates, and the price of risk. Working Paper 32976, National Bureau of Economic Research.
- Kekre, R., Lenel, M., and Mainardi, F. (2024). Monetary Policy, Segmentation, and the Term Structure. Working Paper 32324, National Bureau of Economic Research.
- Kim, D. H. and Wright, J. H. (2005). An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates. Finance and Economics Discussion Series 2005-33, Board of Governors of the Federal Reserve System (U.S.).
- King, T. B. (2019a). Duration Effects in Macro-Finance Models of the Term Structure. Technical report.
- King, T. B. (2019b). Expectation and Duration at the Effective Lower Bound. *Journal of Financial Economics*, 134(3):736–760.
- Koijen, R. S. J. and Yogo, M. (2019). A Demand System Approach to Asset Pricing. *Journal of Political Economy*, 127(4):1475–1515.
- Koijen, R. S. J. and Yogo, M. (2022). Understanding the Ownership Structure of Corporate Bonds. NBER Working Paper 29679.
- Krishnamurthy, A. and Vissing-Jorgensen, A. (2011). The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy. *Brookings Papers on Economic Activity*, 42(2 (Fall)):215–287.
- Krishnamurthy, A. and Vissing-Jorgensen, A. (2012). The Aggregate Demand for Treasury Debt. *Journal of Political Economy*, 120(2):233–267.
- Li, C. and Wei, M. (2013). Term Structure Modeling with Supply Factors and the Federal Reserve’s Large-Scale Asset Purchase Programs. *International Journal of Central Banking*, 9(1):3–39.
- Marcet, A. and Marimon, R. (2019). Recursive Contracts. *Econometrica*, 87(5):1589–1631.
- McCallum, B. T. (1983). On non-uniqueness in rational expectations models: An attempt at perspective. *Journal of Monetary Economics*, 11(2):139–168.
- Mishkin, F. S. (2007). Monetary policy and the dual mandate. Speech at Bridgewater College, Bridgewater, Virginia, April 10, 2007. Board of Governors of the Federal Reserve System.
- Modigliani, F. and Sutch, R. (1966). Innovations in Interest Rate Policy. *The American Economic Review*, 56(1/2):178–197.
- Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect. *The Quarterly Journal of Economics*, 133(3):1283–1330.

- Peng, C. and Wang, C. (2022). Factor Demand and Factor Returns. Technical report.
- Piazzesi, M. (2005). Bond Yields and the Federal Reserve. *Journal of Political Economy*, 113(2):311–344.
- Ray, W. (2019). Monetary Policy and the Limits to Arbitrage: Insights from a New Keynesian Preferred Habitat Model. Working paper.
- Ray, W., Droste, M., and Gorodnichenko, Y. (2024). Unbundling Quantitative Easing: Taking a Cue from Treasury Auctions. *Journal of Political Economy*, 132(9):3115–3172.
- Silva, D. H. (2026). The Risk Channel of Unconventional Monetary Policy. *Review of Financial Studies*. Forthcoming.
- Sims, E. and Wu, J. C. (2021). Evaluating Central Banks’ Tool Kit: Past, Present, and Future. *Journal of Monetary Economics*, 118:135–160.
- Sims, E., Wu, J. C., and Zhang, J. (2023). The Four-Equation New Keynesian Model. *Review of Economics and Statistics*, 105(4):931–947.
- Vayanos, D. and Vila, J.-L. (2021). A Preferred-Habitat Model of the Term Structure of Interest Rates. *Econometrica*, 89(1):77–112.
- Wallace, N. (1981). A Modigliani-Miller Theorem for Open-Market Operations. *American Economic Review*, 71(3):267–274.

Online Appendix for “Optimal Unconventional Policy in a New Keynesian Preferred Habitat Model”

by Rupal Kamdar and Walker Ray

A Proofs

All proofs use the perturbation of Section 2.2: variables are first-order deviations from the deterministic steady state, the loadings $\sigma(\tau)$ are constant at first order, and welfare is evaluated to second order along first-order dynamics. Household no-Ponzi (transversality) conditions and, except where noted below, the saddle-path conditions of Appendix B.3 are maintained.

Leading order in a . Some one-factor results expand around $a = 0$, separately from the ξ -perturbation; a strictly signed leading-order term keeps its sign for $a \in (0, \bar{a})$ by continuity (Theorem 2). For rules under which a combination of x_t and $\check{c}_t^\sigma$ has a zero root at $a = 0$ (see below), the same argument applies after that combination is solved forward, with the spectrum split at $\rho/2$ instead of $\delta/2$. “Given the composites” means that $(\eta^\sigma, \beta^\sigma, \check{\Sigma}, \Omega_\psi)$ are held fixed.

Equilibrium selection and loss criteria. If the short rate responds to neither x_t nor b_t^σ , the IS curve and the relative Euler equation for $\check{c}_t^\sigma$ load on Λ_t in fixed proportion, and a combination of x_t and $\check{c}_t^\sigma$ has a zero root. As in Section 5.1, we select the minimum-state-variable solution: each jump variable is the forward solution of its equation, so $E_t dy_t = (p y_t + g_t) dt$ with $p \geq 0$ gives $y_t = -\int_0^\infty e^{-ps} E_t g_{t+s} ds$ (the unique bounded solution when $p > 0$). In a stationary equilibrium, $\mathcal{W} = E[\mathcal{L}_t]/(2\rho)$. Rules that offset b_t^σ one-for-one in Λ_t leave b_t^σ with a unit root, and some zero-root rules leave it with a positive root of order a , below $\rho/2$ for small a ; for these we evaluate the discounted loss $\frac{1}{2} E_0 \int_0^\infty e^{-\rho t} \mathcal{L}_t dt$ from $b_0^\sigma = 0$. For the one-factor results of Section 5.1, we write the loss of a rule as $O + D$, the output-gap and dispersion terms of the criterion; a superscript names the rule and an asterisk denotes the optimal rule.

Proof of Proposition 1. The optimal subsidy is

$$1 + \tau^y = \frac{\epsilon}{\epsilon - 1} (1 + \beta\rho), \quad (\text{A1})$$

which offsets both the markup and the steady-state working-capital wedge. Throughout, policy instruments carry no resource costs.

Representative agent. With $a = 0$, (16) gives $\mu_t(\tau) = i_t$ for all τ , whatever the arbitrageur positions, and return diffusions vanish as $\xi \rightarrow 0$: all portfolios earn the same riskless return, so households with equal initial wealth make identical choices.

First best. Let $a = 0$, $\vartheta = 0$, $\epsilon_t \equiv \epsilon$ and $\beta_t(\tau) \equiv \beta(\tau)$. Multiplying (14) by ϑ gives $\text{MC}_t = 1 + \beta\rho$, while $\hat{\mu}_t = \beta i_t$. Hence $W_t = Z_t$ at $i_t = \rho$, and (8) gives $C_t^\varsigma N_t^\varphi = Z_t$ with

$C_t = Z_t N_t$ at every date, equally across households. This is the first best, and it is feasible: with flexible prices the Euler equation holds with $\pi_t = \rho - r_t^n$, where r_t^n is the first-best real rate.

Steady state for $a \geq 0$. With $a > 0$, the deterministic limit of (16) retains the premium $a\sigma(\tau) \int_0^T \sigma(\tau')^\top X(\tau') d\tau'$. Let steady-state policy set $i = \rho$ and $G(\tau) = -\beta(\tau)$, and guess $X(\tau) = 0$ for all τ . Then $\mu(\tau) = \rho$, so the stationary Euler equation gives $\pi = 0$ and $\hat{\mu} = \beta\rho$, and (14) gives $MC = 1 + \beta\rho$, that is, $W = Z$. With $C = Y = ZN$, the steady state is efficient and $WN/(\mathcal{P}Y) = 1$, so bond market clearing (19), scaled by $\mathcal{P}Y$, gives $X(\tau) = G(\tau) + \beta(\tau) - b\eta(\tau) = 0$ for $b = 0$, which verifies the guess. Neither a nor ϑ enters, so the sticky-price, imperfect-arbitrage economy shares the efficient steady state, and equal wealth gives $C(i) = C$ and $N(i) = N$. $\square$

Proof of Theorem 1. Relative consumption function. Relative wealth (28) has no diffusion, so solving it forward under the household's no-Ponzi condition gives $\check{b}_t(i) = \bar{\varsigma} \int_0^\infty e^{-\rho s} E_t \check{c}_{t+s}(i) ds$, where $\bar{\varsigma} \equiv 1 + \varsigma/\varphi$. By (29), $E_t \check{c}_{t+s}(i) = \check{c}_t(i) + \varsigma^{-1}a(\eta_i^\sigma - \eta^\sigma) E_t \int_0^s \Lambda_{t+u} du$. Substituting and exchanging the order of integration,

$$\check{c}_t(i) = \frac{\rho}{\bar{\varsigma}} \check{b}_t(i) - \varsigma^{-1}a(\eta_i^\sigma - \eta^\sigma) \Xi_t, \quad \Xi_t \equiv E_t \int_0^\infty e^{-\rho s} \Lambda_{t+s} ds, \quad (\text{A2})$$

and substituting (A2) back into (28),

$$d\check{b}_t(i) = \bar{\varsigma}\varsigma^{-1}a(\eta_i^\sigma - \eta^\sigma) \Xi_t dt. \quad (\text{A3})$$

(1) If $\check{b}_t(i) = (\eta_i^\sigma - \eta^\sigma) \gamma_t^b$, then by (A2) and (A3), $\check{c}_t(i) = (\eta_i^\sigma - \eta^\sigma) \gamma_t^c$ and $d\check{b}_t(i) = (\eta_i^\sigma - \eta^\sigma) d\gamma_t^b$, where $\gamma_t^c = \frac{\rho}{\bar{\varsigma}} \gamma_t^b - \varsigma^{-1}a\Xi_t$ and $d\gamma_t^b = \bar{\varsigma}\varsigma^{-1}a\Xi_t dt$ are common to all households. Hence spanning at $t = 0$ propagates to all t . Using $E_t d\Xi_t = (\rho\Xi_t - \Lambda_t) dt$, these coefficients satisfy (31). With $b \neq 0$ the argument is identical, using (D3) and (D4).

(2) Since $\int_0^1 (\eta_i^\sigma)^\top (\eta_i^\sigma - \eta^\sigma) di = 0$, $\int_0^1 (\eta_i^\sigma)^\top (\eta_i^\sigma - \eta^\sigma) di = \check{\Sigma}$, and integrating the representations in (1) against $(\eta_i^\sigma)^\top$ gives $b_t^\sigma = \check{\Sigma} \gamma_t^b$ and $\check{c}_t^\sigma = \check{\Sigma} \gamma_t^c$. Multiplying (31) by $\check{\Sigma}$ gives $db_t^\sigma = [\rho b_t^\sigma - \bar{\varsigma} \check{c}_t^\sigma] dt$ and $E_t d\check{c}_t^\sigma = \varsigma^{-1}a\check{\Sigma} \Lambda_t dt$, closed given Λ_t .

(3) $\text{Var}_i \check{c}_t(i) = \int_0^1 [(\eta_i^\sigma - \eta^\sigma) \gamma_t^c]^2 di = \gamma_t^{c\top} \check{\Sigma} \gamma_t^c$, which equals $\check{c}_t^{\sigma\top} \check{\Sigma}^+ \check{c}_t^\sigma$ by (2) and $\check{\Sigma} \check{\Sigma}^+ \check{\Sigma} = \check{\Sigma}$.

(4) Divide bond market clearing (19) by $\mathcal{P}_t Y$ and linearize around the steady state of Proposition 1, where $X(\tau) = 0$, $G(\tau) = -\beta(\tau)$, $b = 0$ and $WN/(\mathcal{P}Y) = 1$:

$$x_t(\tau) = g_t(\tau) - qe_t(\tau) + \beta(\tau)(w_t + n_t) + [\beta_t(\tau) - \beta(\tau)] - \int_0^1 \eta_i(\tau) \check{b}_t(i) di,$$

using $b_t(i) = \check{b}_t(i)$ as $b_t \equiv 0$. Averaging (8) gives $w_t = \varsigma y_t + \varphi n_t$ (with $c_t = y_t$ to first order), and with $n_t = y_t - z_{1,t}$ and $y_t = x_t + \frac{1+\varphi}{\varsigma+\varphi} z_{1,t}$, $w_t + n_t = (1 + \varsigma + \varphi)x_t + \frac{1+\varphi}{\varsigma+\varphi} z_{1,t}$. Integrating against $\sigma(\tau)^\top$ gives (32). $\square$

Proof of Proposition 2. Holding the maturity date fixed, the remaining maturity falls one-

for-one, so Itô's lemma applied to the affine conjecture gives

$$d \log P_t^{(\tau)} = [A'(\tau)^\top s_t + A(\tau)^\top \Gamma s_t + C'(\tau)] dt - A(\tau)^\top \sigma_s dB_t.$$

Hence $\sigma(\tau) = -A(\tau)^\top \sigma_s$ and, up to constants and second-order convexity terms, $\mu_t(\tau) = [A'(\tau) + \Gamma^\top A(\tau)]^\top s_t$. Matching coefficients on s_t in (20), $\mu_t(\tau) - \iota^\top s_t = a\sigma(\tau)Ls_t$, gives $A'(\tau) = \iota - MA(\tau)$ with $M = \Gamma^\top + aL^\top \sigma_s^\top$, and $A(0) = 0$ because a maturing bond pays one dollar. Hence $A(\tau) = \int_0^\tau e^{-Mu} du \iota$, which is (33) when M is invertible. $\square$

Proof of Proposition 3. Substituting the static relations (the risk price (32), the effective rates (26) and the linear policy rules) into (21), (22) and the dynamics of $(b_t^\sigma, \check{c}_t^\sigma)$, obtained by integrating (28) and (29) against $(\eta_i^\sigma)^\top$, gives the linear system (B4). Under the saddle-path condition its unique bounded solution is $\omega_t = \Omega s_t$, and the state rows give $ds_t = -\Gamma s_t dt + \sigma_s dB_t$, with $\sigma_s = (\sigma_z^\top, 0)^\top$ because the wealth statistics have no diffusion when $b = 0$. If $\check{\Sigma}$ is singular, the components of $(b_t^\sigma, \check{c}_t^\sigma)$ outside the range of $\check{\Sigma}$ are identically zero, since $\check{\Sigma}v = 0$ implies $(\eta_i^\sigma - \eta^\sigma)v = 0$ for almost all i , and are dropped. $\square$

Proof of Theorem 2. Loss. By Proposition 3, every argument of (24) is linear in s_t (using Theorem 1(3) for the dispersion term, under spanning, and with $\bar{i}_t$ linear in z_t), so $\mathcal{L}_t = s_t^\top \Psi s_t$ with Ψ positive semidefinite.

Fixed point. Let c collect the composites (25), including any risk-weighted supply or policy loadings. Given c , Proposition 3 gives $(\Omega, \Gamma, \iota, L)$, and Proposition 2 then gives M and the loadings $\sigma(\tau) = -A(\tau)^\top \sigma_s$, and hence the updated composites, which we denote $\mathcal{T}(c; a)$. An affine equilibrium is a fixed point $c = \mathcal{T}(c; a)$. The standard regularity conditions are that, at $a = 0$, (i) the (x, π) block under the policy rule is determinate, with its unstable roots bounded away from the imaginary axis, and (ii) $-\check{\Sigma}L_b^0$ has eigenvalues with positive real parts on the range of $\check{\Sigma}$, where the $N_z \times N_z$ matrix L_b^0 is the loading of Λ_t on b_t^σ in the $a = 0$ equilibrium (it includes policy responses to b_t^σ and to $\check{c}_t^\sigma = \rho b_t^\sigma / (1 + \varsigma/\varphi)$).

At $a = 0$. The composites enter (B3) only through $a\eta^\sigma \Lambda_t$, $a\beta^\sigma \Lambda_t$ and $a\check{\Sigma}\Lambda_t$, and $M = \Gamma^\top + aL^\top \sigma_s^\top$, so $\mathcal{T}(\cdot; 0) \equiv c_0$ and $D_c \mathcal{T}(c_0; 0) = 0$. The wealth block has roots 0 and ρ on the range of $\check{\Sigma}$, and its bounded solution $\check{c}_t^\sigma = \rho b_t^\sigma / (1 + \varsigma/\varphi)$ keeps b_t^σ constant (the permanent-income unit root); since σ_s loads only on z_t , $\sigma(\tau)$ remains bounded.

Near $a = 0$. By (i), the roots at $a = 0$ split into those with nonpositive real parts (including the wealth zeros) and those with real parts of at least some $\delta > 0$. For (c, a) near $(c_0, 0)$ no root crosses $\text{Re} = \delta/2$, so the corresponding invariant subspace, and with it $\mathcal{T}$, is smooth (for $T = \infty$, given maturity weights with finite mean). Since $I - D_c \mathcal{T}(c_0; 0) = I$, the implicit function theorem gives a locally unique fixed point $c(a)$, smooth in a . It is a bounded equilibrium if the perturbed wealth roots are stable. For $a > 0$ the drift of $\check{c}_t^\sigma$ gains $\varsigma^{-1}a\check{\Sigma}L_b^0 b_t^\sigma$, which moves the zero roots to $-\rho^{-1}(1 + \varsigma/\varphi)\varsigma^{-1}a\ell_j + O(a^2)$, where ℓ_j are the eigenvalues of $-\check{\Sigma}L_b^0$ on the range of $\check{\Sigma}$. By (ii) these have negative real parts, so the equilibrium exists and is locally unique for $a \in [0, \bar{a})$. If the rule responds to neither b_t^σ nor $\check{c}_t^\sigma$, then $L_b^0 = -I$ and (ii) holds automatically. $\square$

Proof of Proposition 4. With $\beta(\tau) \equiv 0$ and $\phi_x^\sigma = 0$, (36) reads $\lambda_t = -\epsilon_{q,t}^\sigma - b_t^\sigma$, which does not involve x_t . The $(b^\sigma, \check{c}^\sigma)$ block is therefore autonomous and is forced only by the deterministic QE shocks.

Wealth block, general risk price dynamics. Integrating (A3) against $(\eta_i^\sigma)^\top$ gives $db_t^\sigma = \bar{\varsigma}\varsigma^{-1}a\check{\Sigma}\Xi_t dt$, with $\bar{\varsigma} \equiv 1 + \varsigma/\varphi$ and $E_t d\Xi_t = (\rho\Xi_t - \lambda_t) dt$. Suppose that, following an impulse at date 0 with $b_0^\sigma = 0$, $E_0 \lambda_t = F e^{-\kappa t} + L_b E_0 b_t^\sigma$ for constants F , $\kappa > 0$ and $L_b < 0$. Taking E_0 and eliminating $E_0 \Xi_t$ gives a second-order linear differential equation for $E_0 b_t^\sigma$ with roots $-\nu_b$ and $\rho + \nu_b$, where $\nu_b \geq 0$ solves $\nu_b(\nu_b + \rho) = \bar{\varsigma}\varsigma^{-1}a\check{\Sigma}|L_b|$. Its bounded solution is

$$\begin{aligned} E_0 b_t^\sigma &= -\frac{\nu_b(\nu_b + \rho)}{|L_b|} C (e^{-\kappa t} - e^{-\nu_b t}), \\ E_0 \lambda_t &= C [\kappa(\kappa + \rho)e^{-\kappa t} - \nu_b(\nu_b + \rho)e^{-\nu_b t}], \end{aligned} \quad (\text{A4})$$

where $C \equiv F/[(\kappa - \nu_b)(\kappa + \nu_b + \rho)]$ (for $\nu_b = \kappa$, take limits). If $a\check{\Sigma} > 0$, $E_0 b_t^\sigma$ has the sign of F for all $t > 0$.

Natural rate and conventional shocks. Absent QE shocks the block is homogeneous, so from $b_0^\sigma = 0$ its bounded solution is $b^\sigma = \check{c}^\sigma \equiv 0$. Hence $\lambda \equiv \Xi \equiv 0$, and by (A2) and (A3) with $\check{b}_0(i) = 0$, $\check{b}_t(i) = \check{c}_t(i) = 0$ for all i : borrowing rates equal i_t , and (35) is the representative agent Euler equation. Under $i_t = \phi_x x_t + \epsilon_{i,t}$ its root ϕ_x/ς is positive, and its bounded solution $x_t = v_t/(\phi_x + \varsigma\kappa_z) - \epsilon_{i,t}/(\phi_x + \varsigma\kappa_i)$ gives the four loadings.

QE shock. On impact b_t^σ is predetermined, so $\partial\lambda_t/\partial\epsilon_{q,t}^\sigma = -1$. Solving (35) forward, the response of x_t to the shock is $-\varsigma^{-1}a\eta^\sigma \int_0^\infty e^{-(\phi_x/\varsigma)s} E_t \lambda_{t+s} ds$, and (A4), applied from date t with $F = -\epsilon_{q,t}^\sigma$, $\kappa = \kappa_q$ and $L_b = -1$, gives

$$\frac{\partial x_t}{\partial \epsilon_{q,t}^\sigma} = a\eta^\sigma \frac{\phi_x(\kappa_q + \nu_b + \rho) + \varsigma\kappa_q\nu_b}{(\kappa_q + \nu_b + \rho)(\phi_x + \varsigma\kappa_q)(\phi_x + \varsigma\nu_b)} > 0$$

for every $a > 0$ (at leading order in a , $\nu_b \rightarrow 0$ and this is $a\eta^\sigma/(\phi_x + \varsigma\kappa_q)$). Finally, the shock breaks both features of the equivalence: $\lambda \not\equiv 0$, so borrowing rates differ from i_t ; and if $\check{\Sigma} > 0$, (A4) gives $b_{t+k}^\sigma \neq 0$ for all $k > 0$, so by (A2) consumption differs across households with different exposures. $\square$

Proof of Proposition 5. Write $v_t = v_z z_t$ with $v_z \equiv -\varsigma\kappa_z \frac{1+\varphi}{\varsigma+\varphi} < 0$. In the affine equilibrium with state (z_t, b_t^σ) , write $x_t = \Omega_{xz} z_t + \Omega_{xb} b_t^\sigma$ and $\lambda_t = L_z z_t + L_b b_t^\sigma$. By linearity, it suffices to consider a shock $z_0 > 0$ at date 0 with $b_0^\sigma = 0$.

Loadings at $a = 0$. With $i_t = \phi_x x_t$ and $qe_t^\sigma = 0$ (the last paragraph treats $\phi_x^\sigma < 0$), solving (35) forward gives $\Omega_{xz}^0 = v_z/(\phi_x + \varsigma\kappa_z) < 0$ and $\Omega_{xb}^0 = 0$, and (36) gives $L_b^0 = -1$ and

$$L_z^0 = \beta^\sigma \left[(1 + \varsigma + \varphi)\Omega_{xz}^0 + \frac{1 + \varphi}{\varsigma + \varphi} \right] = \beta^\sigma \frac{1 + \varphi}{\varsigma + \varphi} \cdot \frac{\phi_x - \bar{\phi}_x}{\phi_x + \varsigma\kappa_z} < 0.$$

By continuity, there is $\bar{a} > 0$ such that for $a \in (0, \bar{a})$: $\Omega_{xz} < 0$, $L_z < 0$, $L_b < 0$ and $v_z - a\eta^\sigma L_z < 0$.

Impact. $x_0 = \Omega_{xz} z_0 < 0$, so $i_0 = \phi_x x_0 < 0$ and $\lambda_0 = L_z z_0 < 0$.

Reversal. Along the response, $E_0 \lambda_t = L_z z_0 e^{-\kappa_z t} + L_b E_0 b_t^\sigma$, so (A4) applies with $F = L_z z_0 < 0$ and $\kappa = \kappa_z$, and $\nu_b > 0$ since $\check{\Sigma} > 0$. Then C has the sign of $\nu_b - \kappa_z$, $E_0 b_t^\sigma < 0$ for $t > 0$, and the coefficient on the slower exponential in $E_0 \lambda_t$ is positive. A sum of two exponentials changes sign at most once, and $E_0 \lambda_0 = F < 0$, so $E_0 \lambda_t$ changes sign exactly

once, at $\bar{k} = \ln[\kappa_z(\kappa_z + \rho)/(\nu_b(\nu_b + \rho))]/(\kappa_z - \nu_b) > 0$ (for $\nu_b = \kappa_z$, take limits), and $E_0 \lambda_t > 0$ for all $t > \bar{k}$.

Short rate. Solving (35) forward, $E_0 x_t = \varsigma^{-1} \int_0^\infty e^{-(\phi_x/\varsigma)s} E_0[v_{t+s} - a\eta^\sigma \lambda_{t+s}] ds$, where $E_0[v_t - a\eta^\sigma \lambda_t] = (v_z - a\eta^\sigma L_z)z_0 e^{-\kappa_z t} - a\eta^\sigma L_b E_0 b_t^\sigma < 0$, since $v_z - a\eta^\sigma L_z < 0$, $L_b < 0$ and $E_0 b_t^\sigma < 0$. Hence $E_0 x_t < 0$ and $E_0 i_t = \phi_x E_0 x_t < 0$ at all horizons.

Balance sheet rule. With $qe_t^\sigma = \phi_x^\sigma x_t$, (36) adds $-\phi_x^\sigma x_t$ to λ_t . At $a = 0$ the output gap does not depend on λ_t , so Ω_{xz}^0 is unchanged and the impact loading becomes $L_z^0 - \phi_x^\sigma \Omega_{xz}^0 < L_z^0 < 0$. By continuity, reducing $\bar{a}$ if necessary, the impact decline $|L_z|z_0$ is larger under the balance sheet rule. $\square$

Proof of Corollary 5.1. At $a = 0$ (the strict inequalities below extend to small a by continuity, reducing $\bar{a}$ if necessary), solving (35) forward with $i_t = \phi_x x_t + \epsilon_{i,t}$ gives $x_t = -\epsilon_{i,t}/(\phi_x + \varsigma\kappa_i)$ and $i_t = \varsigma\kappa_i \epsilon_{i,t}/(\phi_x + \varsigma\kappa_i)$, so a shock that lowers i_t ($\epsilon_{i,t} < 0$) raises x_t . On impact b_t^σ is predetermined and z_t does not move, so (36) gives $\lambda_t = \beta^\sigma(1 + \varsigma + \varphi)x_t$: per unit change in the short rate, λ_t moves by $-\beta^\sigma(1 + \varsigma + \varphi)/(\varsigma\kappa_i) < 0$. After a technology shock it moves by $L_z^0/(\phi_x \Omega_{xz}^0) = \beta^\sigma(\bar{\phi}_x - \phi_x)/(\varsigma\kappa_z \phi_x) > 0$ (proof of Proposition 5). The conditional reaction does not depend on ϕ_x , while the unconditional one decreases in ϕ_x and diverges as $\phi_x \rightarrow 0$, so the former is smaller in absolute value for ϕ_x below a threshold. $\square$

Lemma A1 (Consumption dispersion and risk prices). *Risk-weighted consumption can be written as*

$$\check{c}_t^\sigma = \frac{\rho}{1 + \varsigma/\varphi} b_t^\sigma - \varsigma^{-1} a \check{\Sigma} \Xi_t, \quad \Xi_t \equiv E_t \int_0^\infty e^{-\rho s} \Lambda_{t+s} ds, \quad (\text{A5})$$

and the dispersion term of the loss (24) is therefore

$$\text{Var}_i \check{c}_t(i) = \left(\frac{\rho}{1 + \varsigma/\varphi} \right)^2 b_t^{\sigma\top} \check{\Sigma}^+ b_t^\sigma - \frac{2\rho a}{(1 + \varsigma/\varphi)\varsigma} b_t^{\sigma\top} \Xi_t + \frac{a^2}{\varsigma^2} \Xi_t^\top \check{\Sigma} \Xi_t. \quad (\text{A6})$$

In an affine equilibrium with $\Lambda_t = Ls_t$ with dynamics (34) we have $\Xi_t = L(\rho I + \Gamma)^{-1} s_t$. Moreover, with zero initial dispersion wealth ($b_0^\sigma = 0$), we have

$$\int_0^\infty e^{-\rho t} \text{Var}_i \check{c}_t(i) dt = \frac{a^2}{\varsigma^2} \int_0^\infty e^{-\rho t} \Xi_t^\top \check{\Sigma} \Xi_t dt. \quad (\text{A7})$$

A closed-form approximation follows under two further simplifications: (i) assume $E_t[a\Lambda_{t+s}] = e^{-\kappa_p s} a\Lambda_t$ for a scalar mean reversion rate κ_p ; (ii) assume a rank-one portfolio dispersion matrix $\check{\Sigma} = s_0^2(\eta^\sigma)^\top \eta^\sigma$, so that $\tilde{\lambda}_t \equiv a\eta^\sigma \Lambda_t$. Then, with the premium drawn from its stationary distribution, the dispersion loss term is

$$\rho E_0 \int_0^\infty e^{-\rho t} \text{Var}_i \check{c}_t(i) dt \approx \frac{s_0^2 \text{Var}(\tilde{\lambda}_t)}{\varsigma^2 (\rho + \kappa_p)^2}. \quad (\text{A8})$$

Proof of Lemma A1. Integrating (A2) against $(\eta_i^\sigma)^\top$ gives (A5). Substituting it into $\text{Var}_i \check{c}_t(i) = \check{c}_t^{\sigma\top} \check{\Sigma}^+ \check{c}_t^\sigma$ (Theorem 1(3)), and using $\check{\Sigma}^+ \check{\Sigma} b_t^\sigma = b_t^\sigma$ (as $b_t^\sigma = \check{\Sigma} \gamma_t^b$) and $\check{\Sigma} \check{\Sigma}^+ \check{\Sigma} = \check{\Sigma}$,

gives (A6). In an affine equilibrium $E_t \Lambda_{t+s} = L e^{-\Gamma s} s_t$, so $\Xi_t = L(\rho I + \Gamma)^{-1} s_t$ whenever the eigenvalues of Γ have real parts above $-\rho$, as in all equilibria considered here.

For (A7), write $\bar{\varsigma} \equiv 1 + \varsigma/\varphi$. By Theorem 1(2), the first two terms of (A6) equal $\frac{\rho}{\bar{\varsigma}^2} [\rho \gamma_t^{b\top} \check{\Sigma} \gamma_t^b - 2\bar{\varsigma} \varsigma^{-1} a \gamma_t^{b\top} \check{\Sigma} \Xi_t]$. Since $d\gamma_t^b = \bar{\varsigma} \varsigma^{-1} a \Xi_t dt$ has no diffusion,

$$d(e^{-\rho t} \gamma_t^{b\top} \check{\Sigma} \gamma_t^b) = -e^{-\rho t} [\rho \gamma_t^{b\top} \check{\Sigma} \gamma_t^b - 2\bar{\varsigma} \varsigma^{-1} a \gamma_t^{b\top} \check{\Sigma} \Xi_t] dt.$$

Integrating over $[0, \infty)$, with $\gamma_0^{b\top} \check{\Sigma} \gamma_0^b = b_0^{\sigma\top} \check{\Sigma}^+ b_0^\sigma = 0$, with $e^{-\rho t} \gamma_t^{b\top} \check{\Sigma} \gamma_t^b \rightarrow 0$, shows that the first two terms integrate to zero against $e^{-\rho t}$, path by path.

For (A8): under (i), $a\Xi_t = a\Lambda_t/(\rho + \kappa_p)$; under (ii), $a^2 \Lambda_t^\top \check{\Sigma} \Lambda_t = s_0^2 \tilde{\lambda}_t^2$; and for a stationary premium, $\rho \int_0^\infty e^{-\rho t} E_0[\tilde{\lambda}_t^2] dt = \text{Var}(\tilde{\lambda}_t)$. Besides (i), the approximation lies in combining $b_0^\sigma = 0$ with a stationary premium: from a stationary b_0^σ , the right-hand side of (A7) gains $\frac{\rho}{\bar{\varsigma}^2} b_0^{\sigma\top} \check{\Sigma}^+ b_0^\sigma$. $\square$

Proof of Lemma 1. By Theorem 1(4), $\{q e_t(\tau)\}_{\tau=0}^T$ enters the equilibrium conditions, and hence the loadings $\sigma(\tau)$, only through $q e_t^\sigma$, and it enters (24) only through the holding cost. Given $q e_t^\sigma$, the profile therefore solves the strictly convex problem $\min \int_0^T \psi(\tau) q e(\tau)^2 d\tau$ subject to $\int_0^T \sigma(\tau)^\top q e(\tau) d\tau = q e^\sigma$. The first-order condition, with multiplier $2m \in \mathbb{R}^{N_z}$, is $\psi(\tau) q e(\tau) = \sigma(\tau) m$, and the constraint gives $\Omega_\psi m = q e^\sigma$, so $m = \Omega_\psi^{-1} q e^\sigma$ when Ω_ψ is invertible (that is, when the N_z components of $\sigma(\cdot)/\sqrt{\psi(\cdot)}$ are linearly independent in L^2). The minimized cost is $m^\top \Omega_\psi m = q e^{\sigma\top} \Omega_\psi^{-1} q e^\sigma$. $\square$

Proof of Proposition 6. There is no balance sheet policy and, as in the text, $\psi^i = 0$. Let $\varpi \equiv a\eta^\sigma \beta^\sigma (1 + \varsigma + \varphi)/\varsigma > 0$. By (36), $\lambda_t = \beta^\sigma(w_t + n_t) - b_t^\sigma$, where $w_t + n_t = (1 + \varsigma + \varphi)x_t - v_t/(\varsigma\kappa_z)$ is the wage bill (proof of Theorem 1(4)) and firm borrowing is $\beta(\tau)(w_t + n_t)$.

Gap-closing rules. $x \equiv 0$ requires $i_t = v_t - a\eta^\sigma \lambda_t$ with $\lambda_t = -\beta^\sigma v_t/(\varsigma\kappa_z) - b_t^\sigma$, that is, $i_t = \chi_{i,v}^n v_t + a\eta^\sigma b_t^\sigma$. If $\check{\Sigma} = 0$, then $b^\sigma = \check{\Sigma} \gamma^b \equiv 0$ and this is the rule $(\chi_{i,v}^n, 0)$; if $\check{\Sigma} > 0$, it is the only rule of the form (39) with $x \equiv 0$. Under it, (35) reads $E_t dx_t = \varpi x_t dt$, whose unique bounded solution is $x \equiv 0$.

Marginal deviation. Let $\check{\Sigma} > 0$, $\chi_{i,b} = a\eta^\sigma$ and $\chi_{i,v} = \chi_{i,v}^n - \varepsilon$. Then $E_t dx_t = [\varpi x_t - \varsigma^{-1} \varepsilon v_t] dt$, so $x_t = x_v v_t$ and $\lambda_t = \lambda_v v_t - b_t^\sigma$, with $x_v = \varepsilon/(\varsigma(\varpi + \kappa_z))$ and $\lambda_v = \beta^\sigma[(1 + \varsigma + \varphi)x_v - 1/(\varsigma\kappa_z)]$. Given the composites, the $(b^\sigma, \check{c}^\sigma)$ block has $L_b = -1$ and is forced only by $\lambda_v v_t$, so $O = \hat{O} x_v^2$ and $D = \lambda_v^2 \hat{D}$, with $\hat{O}, \hat{D} > 0$ independent of ε . The derivative of $O + D$ at $\varepsilon = 0$ is $2\lambda_v \lambda_v' \hat{D} < 0$, since $\lambda_v < 0 < \lambda_v'$ there. With endogenous composites, O still has zero derivative at $\varepsilon = 0$. At leading order in a , $\hat{D}$ is a times a positive constant that does not depend on the composites: with $\bar{\varsigma} \equiv 1 + \varsigma/\varphi$, γ_t^b is then an Ornstein-Uhlenbeck process with mean reversion $\rho^{-1} \bar{\varsigma} \varsigma^{-1} a \check{\Sigma}$ driven by $\bar{\varsigma} \varsigma^{-1} a \lambda_v v_t/(\rho + \kappa_z)$, and by (A6) the stationary dispersion is $(\rho/\bar{\varsigma})^2 \check{\Sigma} \text{Var}(\gamma_t^b)$ up to terms with zero mean or of order a^2 , in which $\check{\Sigma}$ cancels. The composites therefore enter only through the value $\lambda_v = -\beta^\sigma/(\varsigma\kappa_z)$ at $\varepsilon = 0$, and at leading order β^σ is proportional to $\chi_{i,v}$, since $\sigma(\tau)$ is proportional to the loading of i_t on z_t , so β^σ falls as $\chi_{i,v}$ is lowered. Both channels lower the loss.

Items 2 and 3. Given the composites, the loss along this family is a strictly convex quadratic in $\chi_{i,v}$. At $x_v = 0$ it is the loss of the gap-closing allocation, and at $\lambda_v = 0$ (where $\lambda_t = -b_t^\sigma$ gives $b^\sigma = \check{c}^\sigma \equiv 0$) that of the zero-dispersion allocation; its derivative is nonzero at both points, so both losses exceed the optimum. Any rule with $x \equiv 0$ has the

first allocation, and any rule with zero dispersion has $\check{c}^\sigma \equiv 0$, hence $\lambda \equiv 0$ (as $a\check{\Sigma} > 0$), hence $b^\sigma \equiv 0$ and the second. This gives items 2 and 3 for every $a > 0$. With endogenous composites, both comparisons hold for small a : at the gap-closing rule by the marginal step, and at the zero-dispersion allocation because $D \geq 0$ has zero derivative there while O does not.

Item 1. As b_t^σ is predetermined, the impact response of $\mu_t(\tau) - i_t = a\sigma(\tau)\lambda_t$ is $a\sigma(\tau)\beta^\sigma$ times that of $w_t + n_t$, while that of i_t is $\chi_{i,v}^*$ times that of v_t . For small a , $\chi_{i,v}^* > 0$ and $\sigma(\tau) > 0$: the optimal loss is at most the gap-closing loss, which is of order a , while at $a = 0$, $x_t = (1 - \chi_{i,v})v_t/(\varsigma\kappa_z)$, so $\chi_{i,v}^* \rightarrow 1$ as $a \rightarrow 0$, and $\sigma(\tau)$ is then proportional to $\chi_{i,v}^*$ up to terms of order a . Hence carry trade returns move with i_t if and only if firm borrowing moves with v_t . $\square$

Proof of Proposition 7. Under the stated instruments, (36) gives $\lambda_t = \beta^\sigma(1 + \varsigma + \varphi)x_t - b_t^\sigma$, and (35) gives $E_t dx_t = \varsigma^{-1}a\eta^\sigma \lambda_t dt$. The system has no exogenous forcing, so from $b_0^\sigma = 0$ its minimum-state-variable solution is $x = \lambda = b^\sigma = \check{c}^\sigma \equiv 0$. Hence (1) $x \equiv 0$; (2) $\mu_t(\tau) - i_t = a\sigma(\tau)\lambda_t = 0$ for all τ , so $\mu_t(\tau) = \tilde{\mu}_t = i_t$; and (3) $\Xi \equiv 0$, so by (A2) and (A3), $\check{b}_t(i) = \check{c}_t(i) = 0$ for every household. With $\pi \equiv 0$ and $\psi^i = \psi(\tau) = 0$, every term of (24) is zero. $\square$

Proof of Proposition 8. As $\psi^i \rightarrow \infty$, $i_t = \bar{i}_t = \tilde{\chi}_{i,v}v_t$, and with $\psi(\tau) = 0$ the loss is $O + D$, evaluated with the discounted loss from $b_0^\sigma = 0$. Each rule is evaluated at its own equilibrium composites, at which $\check{\Sigma} > 0$ is assumed. Under the rules below, λ_t loads only on b_t^σ or only on z_t , so the z -component of $A(\tau)$ in Proposition 2 (whose derivation does not use stationarity) decouples: $\sigma(\tau) = -\tilde{\chi}_{i,v}v_z\sigma_z(1 - e^{-m\tau})/m > 0$, with $m \equiv \kappa_z + aL_z\sigma_z$, and $\check{\Sigma}/(\eta^\sigma)^2$ depends on the rule only through m .

Rate equalization ($\mathcal{R}$). The first rule gives $\lambda_t = -b_t^\sigma$, so from $b_0^\sigma = 0$, $\lambda = b^\sigma = \check{c}^\sigma \equiv 0$: $\mu_t(\tau) = \tilde{\mu}_t$ and dispersion is zero. Then $E_t dx_t = -\varsigma^{-1}(1 - \tilde{\chi}_{i,v})v_t dt$, whose minimum-state-variable solution is $x_t^\mathcal{R} = (1 - \tilde{\chi}_{i,v})v_t/(\varsigma\kappa_z) \neq 0$.

The δ -rules. For $\delta \geq 0$, the rules implementing $a\lambda_t = \delta(1 - \tilde{\chi}_{i,v})v_t/\eta^\sigma$ give $x = (1 - \delta)x^\mathcal{R}$ by the same argument. Their composites solve $m + \delta(1 - \tilde{\chi}_{i,v})/(\tilde{\chi}_{i,v}g(m)) = \kappa_z$, where $g(m) \equiv \int_0^T \eta(\tau)(1 - e^{-m\tau})/m d\tau$ is positive and decreasing (the average long-bond weight $\eta(\tau)$ is nonnegative), so the root is unique; it exists when $g(m) \rightarrow \infty$ at the lower end of its domain, as for $T < \infty$ or exponential household weights. This equation does not involve a , so its root, and hence $\sigma(\tau)$ and the composites, are the same for every $a > 0$. So is $\delta(1 - \tilde{\chi}_{i,v})v_t/\eta^\sigma$, which is the premium $a\lambda_t$ that drives the $(b^\sigma, \check{c}^\sigma)$ block: the allocation under these rules, including dispersion, does not depend on a .

Macro stabilization ($\mathcal{S}$). This is $\delta = 1$, which every rule with $x \equiv 0$ implements. At $x = 0$, (36) implements it with $\chi_{q,v}^\sigma = -\beta^\sigma/(\varsigma\kappa_z) - (1 - \tilde{\chi}_{i,v})/(a\eta^\sigma) \neq 0$ and $\chi_{q,b}^\sigma = -1$, under which $E_t dx_t = \varsigma^{-1}a\eta^\sigma\beta^\sigma(1 + \varsigma + \varphi)x_t dt$, so $x \equiv 0$ is the unique bounded solution. Since $\sigma(0) = 0 < \eta^\sigma$, $\mu_t(\tau) - \tilde{\mu}_t = a(\sigma(\tau) - \eta^\sigma)\lambda_t \neq 0$ for short maturities, and since $a\check{\Sigma} > 0$, $\check{c}^\sigma \neq 0$ and $b^\sigma \neq 0$.

The optimum. Along the δ -rules, $O(\delta) = (1 - \delta)^2 O^\mathcal{R}$ and, by (A7), $D(\delta) = \delta^2 d_0 h(\delta)$, where $d_0 > 0$ is a constant and $h(\delta)$ is $\check{\Sigma}/(\eta^\sigma)^2$ under the δ -rule. Let $W \equiv O + D$. A rule with $D = 0$ has $\check{c}^\sigma \equiv 0$, hence $\lambda \equiv 0$ and the allocation of $\mathcal{R}$, while $W'(0^+) = -2O^\mathcal{R} < 0$; so $D^* > 0$ and $O^* \leq O^* + D^* < O^\mathcal{R}$. A rule with $O = 0$ has the allocation of $\mathcal{S}$, while $W'(1) = d_0[2h(1) + h'(1)] \neq 0$ outside a knife-edge case (if all households' long-bond weights

are proportional to a common profile, as in our two-type specification, h is constant and $W'(1) = 2D^S > 0$; so $O^* > 0$ and $D^* \leq O^* + D^* < D^S$.

Versus no balance sheet policy. $\bar{W} \equiv \min_{\delta \in [0,1]} W(\delta)$ satisfies $O^* \leq \bar{W} < O^R$ and does not depend on a . Without balance sheet policy, $x_t - (\eta^\sigma / \check{\Sigma}) \check{c}_t^\sigma$ has the exogenous drift $-\varsigma^{-1}(1 - \check{\chi}_{i,v})v_t$, so $x_t = x_t^R + (\eta^\sigma / \check{\Sigma}) \check{c}_t^\sigma$, and by (A5) and Theorem 1, from $b_0^\sigma = 0$, $(\eta^\sigma / \check{\Sigma}) \check{c}_t^\sigma = \varsigma^{-1} a \eta^\sigma (\rho \int_0^t \Xi_s ds - \Xi_t)$. This is of order a in discounted mean square (one wealth root is of order a , possibly positive, and the other is close to ρ ; the equilibrium exists for small a by the argument of the Appendix A preamble for zero-root rules), so O without balance sheet policy differs from O^R by a term of order a , and exceeds $\bar{W} \geq O^*$ for small a . $\square$

Proof of Proposition 9. (1) Given the composites, at $x \equiv \pi \equiv 0$, the IS curve requires $i_t + a \eta^\sigma \Lambda_t = v_t$, and the Phillips curve requires $\kappa \alpha_{\hat{\mu}} \hat{\mu}_t = -u_t$, that is, $\beta i_t + a \beta^\sigma \Lambda_t = -(\vartheta/\epsilon) u_t$, since $\kappa \alpha_{\hat{\mu}} = \epsilon/\vartheta$ by (B2). Subtracting β times the first equation gives $a \Delta^\sigma \Lambda_t = -\tilde{u}_t$, and then $i_t = v_t - a \eta^\sigma \Lambda_t = v_t + \eta^\sigma \tilde{u}_t / \Delta^\sigma$. At $x = 0$, (32) gives $\Lambda_t = -\beta^\sigma v_t / (\varsigma \kappa_z) - q e_t^\sigma - b_t^\sigma$, which yields the qe rule. Conversely, under (41) the path $x \equiv \pi \equiv 0$ satisfies both equations.

(2) Hold the composites fixed and write $\tilde{u}_t = \tilde{u}_z z_t$. The $-b_t^\sigma$ term in $q e_t^\sigma$ removes b_t^σ from Λ_t , so $\Lambda_t = L_z z_t$ with $L_z = -\tilde{u}_z / (a \Delta^\sigma)$, and $a \Xi_t = -\tilde{u}_z z_t / (\Delta^\sigma (\rho + \kappa_z))$. By (A7), from $b_0^\sigma = 0$,

$$\int_0^\infty e^{-\rho t} \text{Var}_i \check{c}_t(i) dt = \frac{\check{\Sigma} \tilde{u}_z^2}{\varsigma^2 (\rho + \kappa_z)^2 (\Delta^\sigma)^2} \int_0^\infty e^{-\rho t} z_t^2 dt,$$

whose expectation diverges as $\Delta^\sigma \rightarrow 0$ along any sequence with $\tilde{u}_z$ bounded away from zero; the other terms of the loss are nonnegative. The instrument loadings on z_t , $v_z + \eta^\sigma \tilde{u}_z / \Delta^\sigma$ for i_t and $-\beta^\sigma v_z / (\varsigma \kappa_z) - L_z$ for $q e_t^\sigma$, are of order $(\Delta^\sigma)^{-1}$. $\square$

Lemma A2 (Commitment with exogenous volatility). *Under standard regularity conditions, given the bond volatility loadings $\sigma(\tau)$, the full-commitment optimum is characterized by the planner's first-order conditions in which the costates on forward-looking equations act as promise states, given by (44). The promise states $(\theta_x, \theta_\pi, \theta_c)$ are stationary. The shadow price of risk is given by $\theta_{\Lambda,t} \equiv a [\varsigma^{-1} (\eta^\sigma)^\top \theta_{x,t} - \kappa \alpha_{\hat{\mu}} (\beta^\sigma)^\top \theta_{\pi,t} + \varsigma^{-1} \check{\Sigma} \theta_{c,t}]$; and $\theta_{b,t}$, with $d\theta_{b,t} = [\theta_{\Lambda,t} - \rho \theta_{b,t}] dt + S_{\theta_b} dB_t$, is the costate on the wealth statistics, pinned by boundedness. The optimal instruments are linear in the natural states and the promise states. When $\psi^i > 0$ and, if balance sheet policy is used, Ω_ψ is nonsingular, they are linear in the promise states alone (and $\bar{i}_t$); in particular, policy is then inertial, responding to innovations only through the dynamics of the promise costates (and $\bar{i}_t$).*

Proof of Lemma A2. Fix $\sigma(\cdot)$, and hence the composites (if $\check{\Sigma}$ is singular, $\check{c}^\sigma$, b^σ and their costates live on the range of $\check{\Sigma}$). In the class of plans with time-invariant feedback on innovations, every variable y_t has a moving-average representation

$$y_t = \bar{y} + \int_0^\infty k_y(u) dB_{t-u}, \quad \text{so that} \quad dy_t = k_y(0) dB_t + \left(\int_0^\infty k'_y(u) dB_{t-u} \right) dt, \quad (\text{A9})$$

where the square-integrable impulse-response kernel $k_y(u)$ is the response of $E_t[y_{t+u}]$ to dB_t . The means $\bar{y}$ solve a separate static problem, and by the Itô isometry the rest of $E[\mathcal{L}_t]$ is $\int_0^\infty \ell(k(u)) du$, with ℓ the quadratic form (24) (with the cost of Lemma 1) applied to

the kernels. By (A9), each drift restriction $E_t dw_t = f_t^w dt$ in (B3) becomes $k'_w = k_{fw}$ in the horizon u , with free initial values for the forward-looking x , π and $\check{c}^\sigma$, $k_b(0) = 0$, and k_z given. This is a deterministic, undiscounted linear-quadratic problem in u , in which ρ appears only as a coefficient of the constraints. The standard regularity conditions are that it has a unique optimum with square-integrable multipliers for any initial wealth and promise values (b_0, θ_0) , as it does when $\psi^i > 0$, Ω_ψ is nonsingular (if balance sheet policy is used), a stabilizing plan with finite loss exists, and the Hamiltonian of the kernel problem has no roots on the imaginary axis. Standard results then give the necessary and, by convexity, sufficient conditions: with $H \equiv \frac{1}{2}\ell + \sum_w \mu_w^\top k_{fw}$, $\partial H / \partial k_i = \partial H / \partial k_{qe} = 0$, $\mu'_w = -\partial H / \partial k_w$, and $\mu_w(0) = 0$ for $w \in \{x, \pi, \check{c}^\sigma\}$.

Let $\theta_{w,t} \equiv \int_0^\infty \mu_w(u) dB_{t-u}$. By (A9), $d\theta_{w,t} = -(\partial H / \partial w) dt + \mu_w(0) dB_t$, now with $H \equiv \frac{1}{2}\mathcal{L}_t + \sum_w \theta_{w,t}^\top f_t^w$: θ_x , θ_π and θ_c are stationary promise states with no diffusion, while θ_b has diffusion $S_{\theta b} = \mu_b(0)$ and is pinned down by boundedness of b^σ . As the kernel problem is undiscounted, the drifts carry no $+\rho\theta$ term. Collecting the derivatives through Λ_t (which enters the IS curve, the Phillips curve through $\hat{\mu}_t$, and the relative Euler equation) into $\theta_{\Lambda,t}$, and using $\partial \Lambda_t / \partial x_t = (1 + \varsigma + \varphi)(\beta^\sigma)^\top$ and $\partial \Lambda_t / \partial b_t^\sigma = \partial \Lambda_t / \partial qe_t^\sigma = -I$, gives (44), the θ_b equation, and the instrument conditions

$$\psi^i(i_t - \bar{i}_t) + \varsigma^{-1}\theta_{x,t} - \kappa\alpha_{\hat{\mu}}\beta\theta_{\pi,t} = 0, \quad qe_t^\sigma = \Omega_\psi\theta_{\Lambda,t} \quad (\text{A10})$$

(with $\psi(\tau) = 0$, the second reads $\theta_{\Lambda,t} = 0$). Since the solution from any (b_0, θ_0) is unique and linear in (b_0, θ_0) , the continuation of the plan from any date is a time-invariant linear function of (s_t, θ_t) . If $\psi^i > 0$ and, if balance sheet policy is used, Ω_ψ is nonsingular, (A10) gives $i_t - \bar{i}_t$ and qe_t^σ as linear functions of $(\theta_x, \theta_\pi, \theta_c)$ alone (θ_Λ does not involve θ_b), which have no diffusion; policy then responds to innovations only through the drift of the promise states (and $\bar{i}_t$). $\square$

Proof of Theorem 3. Let c^* denote the composites induced by an optimal plan. The standard regularity conditions are that an optimal plan exists and that, near c^* , the conditions of Lemma A2 hold and: (R1) the kernel system of the proof of Lemma A2, with the instrument kernels as inputs, is hyperbolic;¹⁶ (R2) $\det(I + \hat{k}_\Lambda(s)) \neq 0$ for $\text{Re } s \geq \delta'$, with $\hat{k}_\Lambda$ the Laplace transform of k_Λ and $0 < \delta'$ below the slowest decay rate of the portfolio weights and supply profiles ($\delta' = 0$ if k_Λ is integrable); and (R3) $\check{\Sigma}$ has constant rank and the optimum is a regular point of the K constraints below, as when the composite fixed point given p is nondegenerate.

Step 1: pricing consistency. By Appendix C, the loadings generated by a plan satisfy (C2), where k_i and k_Λ are the kernels (A9) of i_t and Λ_t . Given the kernels, this is a linear Volterra equation of the second kind in $\sigma(\cdot)$; by (R2) it has a unique solution $\sigma = \mathcal{S}(k_i, k_\Lambda)$, and $\mathcal{S}$ is continuously differentiable.

¹⁶Under rigid prices the argument is run directly at $\kappa = 0$, rather than as the limit $\vartheta \rightarrow \infty$, with three differences. First, in (R1) the kernel system has one zero root, that of $x - \eta^\sigma \check{\Sigma}^+ \check{c}^\sigma$, whose kernel $-\varsigma^{-1}T(k_i - k_v)$, with $Tk(u) \equiv \int_u^\infty k(s) ds$, is solved forward; and η^σ is required to lie in the row space of $\check{\Sigma}$ near c^* . Second, finite loss then requires $Tk_i \in L^2$, so the norm on p in Step 2 adds $\|Tk_i\|$. Third, in the coercivity argument of Step 4, Q also bounds $\|T\delta k_i\|^2$: it contains the loss on x and on $\check{c}^\sigma$, and $|\eta^\sigma \check{\Sigma}^+ \check{c}^\sigma|^2$ is bounded by a multiple of $\check{c}^{\sigma^\top} \check{\Sigma}^+ \check{c}^\sigma$ because η^σ lies in the row space of $\check{\Sigma}$.

Step 2: finitely many constraints. Treat $c \in \mathbb{R}^K$ as an additional choice variable. Given c and the instrument kernels (k_i, k_{qe}) , the kernel system of the proof of Lemma A2 is a linear system of differential equations in u with constant coefficients that depend on c . The kernels of the states are pinned down by their impact values ($k_b(0) = 0$, k_z given). The impact responses of the forward-looking variables are free, except that the solution must be square integrable, which by (R1) confines them to an affine set of finite dimension; let ξ be coordinates on this set. For example, at $a = 0$ with sticky prices the (x, π) block contributes one free response per shock, the familiar indeterminacy of an exogenous rate path. The planner chooses $p \equiv (k_i, k_{qe}, \xi)$, that is, both the instrument paths and the equilibrium they implement, and the kernels are $k = \mathcal{K}(p; c)$. We norm p by $\|k_i\| + \|k_{qe}\| + |\xi|$ (in L^2). This norm does not depend on c , and by (R1) the stable projector is smooth in c , so $\mathcal{K}$ is affine and bounded in p and smooth in c . Writing $\mathcal{C}(\sigma)$ for the composites (25), the planner solves

$$\min_{p, c} G(p, c) \equiv \frac{1}{2} \mathbb{E}[\mathcal{L}_t](\mathcal{K}(p; c); c) \quad \text{s.t.} \quad \mathcal{C}(\mathcal{S}(\mathcal{K}(p; c))) = c,$$

which has K equality constraints. By (R3) and the Lagrange multiplier rule, at the optimum (p^*, c^*) there is $\zeta \in \mathbb{R}^K$ with

$$D_p G + \zeta^\top D(\mathcal{C} \circ \mathcal{S}) D_p \mathcal{K} = 0, \quad D_c G + \zeta^\top D(\mathcal{C} \circ \mathcal{S}) D_c \mathcal{K} = \zeta^\top. \quad (\text{A11})$$

Step 3: ν and the shadow factors. Let ν solve (C3) with forcing $w \equiv \sum_k \zeta_k w_k$; this adjoint equation is uniquely solvable by (R2). For any kernel direction δk , the adjoint identity (C4) gives

$$\zeta^\top D(\mathcal{C} \circ \mathcal{S})[\delta k] = - \int_0^T [\delta k_i(u) \nu_i(u) + \text{tr}(\delta k_\Lambda(u) \nu_\Lambda(u))] du \equiv -\ell_\nu(\delta k), \quad (\text{A12})$$

with ν_i and ν_Λ as in (46). It replaces the constraints on the continuum of loadings by a linear functional of the kernels. By (A9) and the Itô isometry, $\ell_\nu(k) = \mathbb{E}[i_t v_t^i + \Lambda_t^\top v_t^\Lambda]$, with the shadow factors (45) held fixed. By (A12), the first condition in (A11) is the first-order condition in p of the augmented problem $\min_p F(p; c^*)$, where

$$F(p; c) \equiv G(p, c) - \ell_\nu(\mathcal{K}(p; c)) = \frac{1}{2} \mathbb{E}[\mathcal{L}_t] - \mathbb{E}[i_t v_t^i + \Lambda_t^\top v_t^\Lambda].$$

Since F is a convex quadratic in p , p^* minimizes it. The multiplier function ν is thus the adjoint of the linearized pricing map.

Step 4: $\zeta = \partial J / \partial c$. With ζ , and hence ν and the shadow factors, held fixed, let $J(c) \equiv \min_p F(p; c)$, measured in the units of $\frac{1}{2} \mathbb{E}[\mathcal{L}_t]$ in which the costates of Lemma A2 are defined. The quadratic part $Q(h) \equiv \frac{1}{2} D_{pp} F[h, h]$ is coercive on the choice set: $Q(h) \geq m \|h\|^2$ for some $m > 0$.¹⁷ By the Lax-Milgram theorem and the implicit function theorem, the minimizer

¹⁷For $h = (\delta k_i, \delta k_{qe}, \delta \xi)$: (i) Q is a sum of positive semidefinite terms that include $\psi^i \|\delta k_i\|^2$ and δk_{qe} weighted by Ω_ψ^{-1} ($\psi^i > 0$, Ω_ψ nonsingular); (ii) at fixed instruments, a nonzero $\delta \xi$ changes the impact response, and hence the kernel, of some forward-looking variable, so $Q(0, 0, \delta \xi) \geq m_\xi |\delta \xi|^2$ with $m_\xi > 0$ because ξ is finite-dimensional. If $Q(h_n) \rightarrow 0$ with $\|h_n\| = 1$, then (i) (and, under rigid prices, footnote 16) forces the instrument parts of h_n to zero, so $|\delta \xi_n| \rightarrow 1$ while $Q(0, 0, \delta \xi_n) \leq 2Q(h_n) + 2Q(h_n - (0, 0, \delta \xi_n)) \rightarrow 0$, contradicting (ii).

$p(c)$ is unique and continuously differentiable near c^* . By the envelope theorem, (A12) and the second condition in (A11),

$$\frac{\partial J}{\partial c} = D_c G - \ell_\nu(D_c \mathcal{K} \cdot) = D_c G + \zeta^\top D(\mathcal{C} \circ \mathcal{S}) D_c \mathcal{K} = \zeta^\top.$$

Step 5: the optimality system. The augmented problem is the problem of Lemma A2 with H replaced by $H - i_t v_t^i - \Lambda_t^\top v_t^\Lambda$. Every derivative through Λ_t therefore collects $\theta_{\Lambda,t} - v_t^\Lambda$ instead of $\theta_{\Lambda,t}$, and the short-rate condition in (A10) gains $-v_t^i$. The added term is linear and pointwise in u , so the transversality conditions, and hence the diffusion-free promise states, are unchanged. The system is linear with exogenous stationary forcing (v^i, v^Λ) ; its bounded solution is linear in (s_t, θ_t) plus conditional expectations of the forcing, which by (A9) and (45) are convolutions of past dB with kernels linear in ν . This gives (47), with f_i and F_q linear in ζ given k_Λ and $\sigma(\cdot)$, and with constant ϵ_i, ϵ_q under the unconditional criterion.

Step 6: finite span and finite dimension. Since (C3) is linear in its forcing, $\nu = \sum_k \zeta_k \nu^{(k)}$, where $\nu^{(k)}$ solves (C3) with forcing w_k ; setting $\zeta = 0$ gives the plan of Lemma A2. If the plan has a finite-state realization, $k_\Lambda(u) = L^+ e^{-\Gamma^+ u} \sigma_s^+$ and $\sigma(\cdot)$ are summarized by the extended pricing matrix M^+ (L^+ is read off M^+ when $a > 0$ and σ_z is invertible), (C5) gives the $\nu^{(k)}$ in closed form, and every remaining object is a linear solve given (c, ζ, M^+) . The optimum is therefore a fixed point in $(c, \zeta) \in \mathbb{R}^{2K}$ and M^+ . The problem in (p, c) is not convex, so these conditions are necessary only, and their solutions must be ranked by value. Appendix C gives sufficient conditions for a finite-state realization. $\square$

B Model Details

B.1 Household Optimality

Let $\zeta_t(i)$ denote the costate on nominal wealth, with $d\zeta_t(i) = \zeta_t(i)[\mu_{\zeta,t}(i) dt + \sigma_{\zeta,t}(i) dB_t]$. The maximum principle gives the static conditions $e^{-\rho t + \varrho t} C_t(i)^{-\varsigma} = \zeta_t(i) \mathcal{P}_t$ and $e^{-\rho t + \varrho t} N_t(i)^\varphi = \zeta_t(i) \mathcal{W}_t$ and, since wealth enters the budget constraint linearly with return $d\tilde{R}_t(i)$, the costate condition $\mu_{\zeta,t}(i) = -[\tilde{\mu}_t(i) + \sigma_{\zeta,t}(i) \tilde{\sigma}_t(i)^\top]$. Applying Itô's lemma to $\zeta_t(i) = e^{-\rho t + \varrho t} e^{-\varsigma c_t(i)} \mathcal{P}_t^{-1}$ (the price level carries no diffusion under Rotemberg pricing) and substituting yields the exact Euler equation

$$\begin{aligned} \mu_{\zeta,t}(i) &= \frac{1}{\varsigma} [\tilde{\mu}_t(i) - \pi_t - \rho - \kappa_\varrho \varrho_t] \\ &\quad + \frac{\varsigma}{2} |\sigma_{c,t}(i)|^2 - \sigma_{c,t}(i) \tilde{\sigma}_t(i)^\top + \frac{1}{\varsigma} \left[\frac{1}{2} |\sigma_\varrho|^2 - \varsigma \sigma_\varrho \sigma_{c,t}(i)^\top + \sigma_\varrho \tilde{\sigma}_t(i)^\top \right], \end{aligned} \quad (\text{B1})$$

where the second line is equal to $\varsigma^{-1} \xi_t(i)$ in (9); and where $\mu_{c,t}(i)$ and $\sigma_{c,t}(i)$ are the drift and diffusion of $c_t(i)$. Every term in the second line is a product of two first-order diffusion loadings, hence second order under the perturbation of Section 2.2; to first order, $\mathbb{E}_t dc_t(i) = \varsigma^{-1} [\tilde{\mu}_t(i) - \pi_t - \kappa_\varrho \varrho_t] dt$ in deviations from steady state. The discount shock enters only through its expected change $-\kappa_\varrho \varrho_t$.

B.2 Firm Optimality and the Phillips Curve

Let $\gamma_t(j)$ be the costate on $\mathcal{P}_t(j)$ and $Q_t^{\mathcal{H}} \equiv \mathcal{P}_t^{-1} \int_0^1 C_t(i)^{-\varsigma} di$. Because the stochastic integral of the financing diffusion has zero expectation, only the drift of $d\Pi_t(j)$ is priced, so the firm's flow objective depends on financing costs only through $\hat{\mu}_t$. Let F denote the drift of $d\Pi_t(j)$ in (10), net of the financing diffusion (so $E_t d\Pi_t(j) = F dt$). The Hamiltonian is $e^{-\rho t} Q_t^{\mathcal{H}} F + \gamma_t(j) \mathcal{P}_t(j) \pi_t(j)$, with first-order condition $\gamma_t(j) \mathcal{P}_t(j) = e^{-\rho t} Q_t^{\mathcal{H}} \vartheta \pi_t(j) \mathcal{P}_t Y_t$ and costate equation $E_t d\gamma_t(j) = -[e^{-\rho t} Q_t^{\mathcal{H}} F_{\mathcal{P}(j)} + \gamma_t(j) \pi_t(j)] dt$. Differentiating and substituting along with symmetry gives (14).

Linearization. Around the zero-inflation steady state, $E_t d(Q^{\mathcal{H}} Y \pi) / (Q^{\mathcal{H}} Y) = E_t d\pi + O(\xi^2)$, and log real marginal cost is $mc_t = (\varsigma + \varphi)x_t + \frac{\hat{\mu}_t + \rho \hat{\beta}_t}{1 + \beta \rho}$, using $w_t - z_{1,t} = (\varsigma + \varphi)x_t$ and splitting the financing rate into the portfolio component $\hat{\mu}_t$ (steady-state weights) and the financing shock component $\rho \hat{\beta}_t$. Substituting into (14) gives (22) with

$$\kappa = \frac{\epsilon(1 + \beta\rho)(\varsigma + \varphi)}{\vartheta}, \quad u_t = -\frac{1 + \tau^y}{\vartheta \epsilon} (\epsilon_t - \epsilon) = -\frac{1 + \beta\rho}{\vartheta (\epsilon - 1)} (\epsilon_t - \epsilon). \quad (\text{B2})$$

Note that a fall in the demand elasticity is interpreted as a rise in the desired markup. Since $E \epsilon_t = \epsilon$ and ϵ_t enters the welfare derivation only through the symmetric-equilibrium aggregator, markup shocks leave the steady state, the subsidy (A1), and the form of the loss (24) unchanged.

B.3 The First-Order Equilibrium System

Eliminating the static relations (risk prices (Theorem 1(4)), effective rates (26), financing shocks, and the policy rules $i_t = \phi_i^\top s_t$, $qe_t^\sigma = \Phi_q s_t$) gives dynamics

$$\begin{aligned} E_t dx_t &= \varsigma^{-1} [i_t + a\eta^\sigma \Lambda_t - \pi_t - v_t] dt, \\ E_t d\pi_t &= \left[\rho\pi_t - \kappa \left(x_t + \alpha_{\hat{\mu}} \hat{\mu}_t + \alpha_{\hat{\beta}} \hat{\beta}_t \right) - u_t \right] dt, \\ E_t d\check{c}_t^\sigma &= \varsigma^{-1} a \check{\Sigma} \Lambda_t dt, \\ db_t^\sigma &= [\rho b_t^\sigma - (1 + \varsigma/\varphi) \check{c}_t^\sigma] dt, \end{aligned} \quad (\text{B3})$$

together with the factor dynamics (1), where Λ_t and $\hat{\mu}_t$ are the linear functions of (ω_t, s_t) given by (32) and (26). Stacking the jump variables $\omega_t = (x_t, \pi_t, \check{c}_t^{\sigma^\top})^\top$ ($2 + N_z$ variables) on the states $s_t = (z_t^\top, b_t^{\sigma^\top})^\top$ ($2N_z$ variables, plus any policy states) yields

$$\begin{pmatrix} E_t d\omega_t / dt \\ E_t ds_t / dt \end{pmatrix} = H \begin{pmatrix} \omega_t \\ s_t \end{pmatrix}, \quad \sigma_s = \begin{pmatrix} \sigma_z \\ 0 \end{pmatrix}, \quad (\text{B4})$$

where H collects the coefficients of (B3) after the static substitutions (a known function of deep parameters, the composites (25), the shock loading maps, and the policy coefficients). Assume standard saddle-path conditions: exactly $2 + N_z$ eigenvalues of H with positive real part (the wealth statistics are predetermined when $b = 0$; for $b \neq 0$, see Appendix D). Then the unique bounded solution is $\omega_t = \Omega s_t$, and Γ in (34) is the implied equilibrium dynamics matrix.

Fixed-point algorithm. (1) Guess loadings $\{\sigma(\tau)\}_{\tau=0}^\infty$; (2) form the composites (25) and risk-weighted policy terms; (3) solve (B4) for $(\Omega, \Gamma, \iota, L)$; (4) update $A(\tau)$ from Proposition 2 and set $\sigma(\tau) = -A(\tau)^\top \sigma_s$; (5) iterate to convergence. Theorem 2 guarantees a locally unique fixed point near $a = 0$; away from that neighborhood, we select the solution continuously connected in a to the $a = 0$ equilibrium; this same homotopy simultaneously serves as the solution algorithm and the selection device.

B.4 Microfoundation of Fixed Portfolio Weights

Fixed heterogeneous weights are individually optimal under even small fixed rebalancing costs, because portfolio choice is a *second-order* problem (and this remains second-order even with the perturbation approach utilized in this paper, since households are not directly affected by the arbitrageur frictions). With $b = 0$, relative wealth is locally deterministic, so $d(e^{-\rho t} \check{b}_t(i)) = -e^{-\rho t}(1 + \varsigma/\varphi) \check{c}_t(i) dt$ holds under any portfolio weight choice, even with time-variation in weights. Additionally, $\int_0^\infty e^{-\rho t}(1 + \varsigma/\varphi) \check{c}_t(i) dt = \check{b}_0(i)$, so utility is independent of portfolio weights at first order, and utility differences across weight policies are $O(\xi^2)$. Hence for any fixed cost $\varkappa > 0$ there exists $\bar{\xi}(\varkappa) > 0$ such that for all $\xi < \bar{\xi}$, never rebalancing is optimal for every household. However, while each individual household's loss due to fixed rebalancing costs is second order, the dispersion term (24) is the population sum of losses across all individuals; policy that stabilizes term premia removes these costs without individual households paying the fixed cost.

C Full Commitment: Computation

Fix a candidate loading function $\sigma(\cdot)$; every composite is computed from it directly.

Pricing consistency. Follow a bond bought at date t with maturity τ , so that its remaining maturity at $t + u$ is $\tau - u$. Along this holding path, $d \log P_{t+u}^{(\tau-u)} = \mu_{t+u}(\tau - u) du + \sigma(\tau - u) dB_{t+u}$ at first order (the Itô correction is second order). Integrating from $u = 0$ to τ , using $\log P^{(0)} = 0$, taking E_t , and using $\mu_t(\tau) - i_t = a\sigma(\tau)\Lambda_t$ gives

$$\log P_t^{(\tau)} = -\rho\tau - E_t \int_0^\tau [i_{t+u} + a\sigma(\tau - u)\Lambda_{t+u}] du. \quad (C1)$$

Write i_t and Λ_t as in (A9), with impulse-response kernels k_i and k_Λ (the responses of $E_t[i_{t+u}]$ and $E_t[\Lambda_{t+u}]$ to dB_t). The right-hand side of (C1) is then a moving average of past dB with kernel $K_\tau(v) = -\int_0^\tau [k_i(u+v) + a\sigma(\tau - u)k_\Lambda(u+v)] du$. Its loading on dB_t , $K_\tau(0)$, differs from the loading $\sigma(\tau)$ along the holding path only by the drift term $-\partial_\tau \log P_t^{(\tau)} dt$, so the two coincide. Consistency of the plan with the candidate loadings therefore requires

$$\sigma(\tau) = -\int_0^\tau [k_i(u) + a\sigma(\tau - u)k_\Lambda(u)] du \quad \forall \tau, \quad (C2)$$

a continuum of linear equality constraints on the plan (equivalent to the fixed point of Proposition 2).

Multipliers. Because the loss depends on $\sigma(\cdot)$ only through the K composites c , once (C2) is solved for $\sigma(\cdot)$ given the plan, the planner faces only the K constraints that the composites be consistent with $\sigma(\cdot)$. The multiplier function $\nu(\tau)$ is the adjoint of the linearized pricing map (C2): interchanging the order of integration produces the shadow factors (46), which enter the objective as $-E[i_t v_t^i + \Lambda_t^\top v_t^\Lambda]$, and $\nu = \sum_k \zeta_k \nu^{(k)}$, where $\nu^{(k)}$ solves the linear Volterra equation of the second kind

$$\nu^{(k)}(\tau) + a \int_\tau^T k_\Lambda(s - \tau) \nu^{(k)}(s) ds = w_k(\tau) \equiv \partial c_k / \partial \sigma(\tau)^\top, \quad (\text{C3})$$

which for $T < \infty$ is solved in a single backward pass from the long end ($\nu^{(k)}(T) = w_k(T)$); for $T = \infty$, as in our calibration, we use the closed form below.

Adjoint identity. Let $\nu = \sum_k \zeta_k \nu^{(k)}$, which solves (C3) with forcing $w \equiv \sum_k \zeta_k w_k$. Consider any direction $\delta k = (\delta k_i, \delta k_\Lambda)$ of the kernels. By (C2), the induced change in the loadings solves

$$\delta \sigma(\tau) = - \int_0^\tau [\delta k_i(u) + a \sigma(\tau - u) \delta k_\Lambda(u) + a \delta \sigma(\tau - u) k_\Lambda(u)] du,$$

and the induced change in c_k is $\delta c_k = \int_0^T \delta \sigma(\tau) w_k(\tau) d\tau$. Multiply (C3), with forcing w , on the left by $\delta \sigma(\tau)$ and integrate over τ . Substituting the expression for $\delta \sigma$ into the term $\int_0^T \delta \sigma(\tau) \nu(\tau) d\tau$, the change of variables $\tau' = \tau - u$ turns the part in $\delta \sigma(\tau - u) k_\Lambda(u)$ into exactly minus the integral term of (C3). Exchanging the order of integration over $0 \leq u \leq \tau \leq T$ in what remains, and using (46),

$$\sum_k \zeta_k \delta c_k = - \int_0^T [\delta k_i(u) \nu_i(u) + \text{tr}(\delta k_\Lambda(u) \nu_\Lambda(u))] du. \quad (\text{C4})$$

By (A9) and the Itô isometry, the right-hand side is the change in $-E[i_t v_t^i + \Lambda_t^\top v_t^\Lambda]$ with the shadow factors (45) held fixed: a change in the plan changes the shadow-factor term by exactly the value of the induced change in the composites.

Finite-state plans. When the plan admits a finite-state realization (Theorem 3), $k_\Lambda(u) = L^+ e^{-\Gamma^+ u} \sigma_s^+$, and the substitution $\phi(\tau) \equiv \int_\tau^T e^{-\Gamma^+(s-\tau)} \sigma_s^+ \nu^{(k)}(s) ds$ converts (C3) into the adjoint of the pricing ODE, $\phi' = M^{+\top} \phi - \sigma_s^+ w_k$ with $\phi(T) = 0$, with closed form

$$\nu^{(k)}(\tau) = w_k(\tau) - a L^+ \int_\tau^T e^{M^{+\top}(\tau-s)} \sigma_s^+ w_k(s) ds, \quad (\text{C5})$$

where $+$ superscripts denote extended-state objects. A finite-state realization exists, for instance, when $T = \infty$, the household, firm and supply maturity weights are exponential-polynomials in a finite common set of rates $\{\delta_p\}$, and the balance sheet cost does not depend on the loadings. The forcings w_k then have these rates. For any plan, $\int_\tau^\infty k_\Lambda(s - \tau) q e^{-\delta s} ds = \hat{k}_\Lambda(\delta) q e^{-\delta \tau}$ for any vector q , where $\hat{k}_\Lambda$ is the Laplace transform of k_Λ . So for $w_k = q e^{-\delta \tau}$, the unique solution of (C3) is $\nu^{(k)}(\tau) = (I + a \hat{k}_\Lambda(\delta))^{-1} q e^{-\delta \tau}$, where the inverse exists by condition (R2) in the proof of Theorem 3 (for finite-state k_Λ this is (C5)); exponential-

polynomials follow by differentiating in δ . The $\nu^{(k)}$, and hence the kernels (46), therefore have the rates $\{\delta_p\}$, and the shadow factors are linear combinations of memory states with $dm_{p,t} = -\delta_p m_{p,t} dt + dB_t$ (with Jordan chains for polynomial factors), whose rates are known in advance. By Step 5 of that proof, the plan is then linear feedback on (s_t, θ_t, m_t) . The weights and the balance sheet cost in our calibration (Section 6) satisfy these conditions.

Algorithm (fixed point in $(c, \zeta) \in \mathbb{R}^{2K}$ and the extended pricing matrix M^+): given (c, ζ, M^+) , (1) solve the K basis equations (C5) and form $\nu = \sum_k \zeta_k \nu^{(k)}$ and the shadow kernels (46); (2) solve the amended linear optimality system of Theorem 3 for the plan; (3) update $\sigma(\cdot)$, hence c ; (4) update $\zeta = \partial J / \partial c$ from envelope formulas (stationary bilinear moments of the solved system). Every step is a linear solve or an explicit formula; setting $\zeta = 0$ reproduces the plan of Lemma A2. When balance sheet costs are specified directly on the risk-weighted instrument (an exogenous quadratic cost), the induced cost metric is invariant to the loadings and drops from the composite list; otherwise c also includes the entries of $\Omega_\psi \equiv \int_0^T \sigma(\tau)^\top \sigma(\tau) / \psi(\tau) d\tau$.

Conditional criterion. If the planner instead minimizes $E_t \int_t^\infty \frac{1}{2} e^{-\rho(s-t)} \mathcal{L}_s ds$ from a given date-0 state, the same construction applies with three changes: each costate drift in Lemma A2 gains the current-value term $+\rho\theta$ and the promise states start at $\theta_0 = 0$; the shadow factors are $v_t^i = \int_0^t e^{\rho u} \nu_i(u)^\top dB_{t-u}$ and $v_t^\Lambda = \int_0^t e^{\rho u} \nu_\Lambda(u)^\top dB_{t-u}$, the factor $e^{\rho u}$ undoing the planner's discounting of a horizon- u effect on the undiscounted consistency constraint (multiplier scale absorbed into ν); and the marginal values of the composites ζ are $\int_0^\infty e^{-\rho t}$ -weighted transitional moments. The Volterra equation (C3) and its closed form (C5) are unchanged. The extension of Theorem 3 to this criterion is heuristic: the conditionally optimal plan is time-varying, so its loadings are time-varying and the time-invariant consistency constraint (C2) does not apply exactly.

D Nonzero Steady-State Wealth

This appendix relaxes the assumption of zero net steady-state wealth: households hold a common real bond wealth $b \geq 0$ in the deterministic steady state, and zero net supply at each maturity now requires $G(\tau) = b\eta(\tau) - \beta(\tau)$, so that arbitrageur positions remain zero and Proposition 1 is unchanged. We first characterize the first-order dynamics in full. We then show that the second-order loss retains its form and derive the welfare consequences of revaluations under perfect arbitrage.

D.1 Revaluation Effects and First-Order Dynamics

Linearizing the budget constraint (4) with nonzero steady-state wealth b and arbitrary transfer drift f_t and diffusion σ_f gives

$$db_t(i) = [\rho b_t(i) + b(\tilde{\mu}_t(i) - \pi_t) + w_t + n_t(i) - c_t(i) + f_t] dt + [b\eta_i^\sigma + \sigma_f] dB_t. \quad (D1)$$

The diffusion term $b\eta_i^\sigma dB_t$ is the *revaluation* of household i 's steady-state portfolio following a shock. The term is first order under the perturbation of Section 2.2: the fund's return

diffusion is $\xi \hat{\eta}_i^\sigma + O(\xi^2)$ and real wealth is $b + O(\xi)$, so their product has diffusion $\xi b \hat{\eta}_i^\sigma + O(\xi^2)$ (the same order as $b \tilde{\mu}_t(i)$). We adopt value-targeting issuance throughout (government supply $G_t(\tau)$ is specified directly as a real market value), so no revaluation term arises from government holdings.¹⁸

Averaging (D1) over i and using $w_t + n_t - c_t = (\varsigma + \varphi)x_t$, aggregate wealth evolves according to

$$db_t = [\rho b_t + b(\tilde{\mu}_t - \pi_t) + (\varsigma + \varphi)x_t + f_t] dt + [b\eta^\sigma + \sigma_f] dB_t, \quad (D2)$$

which is now a genuine state variable (due to its impact on aggregate bond demand). The transfer process (f_t, σ_f) is arbitrary; note that the wealth-stabilizing rule $\sigma_f = -b\eta^\sigma$, $f_t = -b(\tilde{\mu}_t - \pi_t) - (\varsigma + \varphi)x_t$ delivers $b_t \equiv 0$. Differencing (D1) from its cross-sectional mean and using relative labor supply,

$$d\check{b}_t(i) = [\rho \check{b}_t(i) + b\check{\lambda}_t(i) - (1 + \varsigma/\varphi)\check{c}_t(i)] dt + b(\eta_i^\sigma - \eta^\sigma) dB_t, \quad (D3)$$

with $\check{\lambda}_t(i) = a(\eta_i^\sigma - \eta^\sigma)\Lambda_t$ the relative portfolio premium as in the main text: with $b > 0$, relative wealth earns first-order relative premium income and carries first-order relative revaluations. The relative Euler equation (29) is unchanged. Solving (D3) forward gives the relative consumption function

$$\check{c}_t(i) = \frac{\rho\varphi}{\varphi + \varsigma} \check{b}_t(i) + \left(\frac{\rho\varphi}{\varphi + \varsigma} b - \frac{1}{\varsigma} \right) a(\eta_i^\sigma - \eta^\sigma) E_t \int_t^\infty e^{-\rho(s-t)} \Lambda_s ds. \quad (D4)$$

Thus, relative consumption equals the permanent-income value of relative wealth (the first term) plus two premium effects: an income effect on the wealth level (when $b \neq 0$) and the intertemporal substitution effect $(-1/\varsigma)$.

Proposition 10 (Sufficient statistics, $b \geq 0$). *Suppose initial relative wealth is spanned by exposures, $\check{b}_0(i) = (\eta_i^\sigma - \eta^\sigma)\gamma_0^b$. Then for all t , $\check{b}_t(i) = (\eta_i^\sigma - \eta^\sigma)\gamma_t^b$ and $\check{c}_t(i) = (\eta_i^\sigma - \eta^\sigma)\gamma_t^c$, where*

$$d\gamma_t^b = [\rho\gamma_t^b + ab\Lambda_t - (1 + \varsigma/\varphi)\gamma_t^c] dt + b dB_t, \quad E_t d\gamma_t^c = a\varsigma^{-1}\Lambda_t dt; \quad (D5)$$

the risk-weighted statistics satisfy $b_t^\sigma = \check{\Sigma}\gamma_t^b$, $\check{c}_t^\sigma = \check{\Sigma}\gamma_t^c$, with

$$db_t^\sigma = [\rho b_t^\sigma + ab\check{\Sigma}\Lambda_t - (1 + \varsigma/\varphi)\check{c}_t^\sigma] dt + b\check{\Sigma} dB_t, \quad E_t d\check{c}_t^\sigma = \varsigma^{-1}a\check{\Sigma}\Lambda_t dt; \quad (D6)$$

consumption dispersion is $\text{Var}_i \check{c}_t(i) = \gamma_t^{c\top} \check{\Sigma} \gamma_t^c = \check{c}_t^{\sigma\top} \check{\Sigma}^+ \check{c}_t^\sigma$; and risk prices are

$$\Lambda_t = (\beta^\sigma)^\top \left[(1 + \varsigma + \varphi)x_t + \frac{1 + \varphi}{\varsigma + \varphi} z_{1,t} \right] + g_t^\sigma - qe_t^\sigma + \hat{\beta}_t^\sigma - (\eta^\sigma)^\top b_t - b_t^\sigma. \quad (D7)$$

¹⁸Under face-value (quantity) targeting, the real value of outstanding supply gains the endogenous first-order term $G(\tau)p_t(\tau)$, with $p_t(\tau)$ the log-price deviation, which adds a further feedback block to (D7) of the form $\sigma_s^\top [\int_0^T G(\tau)A(\tau)A(\tau)^\top d\tau] s_t$. Either approach is valid; we focus on value targeting for simplicity as it does not generate new composites in the volatility fixed point.

Hence the household side enters equilibrium only through $(b_t^\sigma, \check{c}_t^\sigma)$ and the scalar b_t .

The proof is identical to that of Theorem 1.

State-space form. The state is $s_t = (z_t^\top, b_t^{\sigma^\top}, b_t)^\top$ ($2N_z + 1$ variables, plus any policy states), the jumps are $\omega_t = (x_t, \pi_t, \check{c}_t^{\sigma^\top})^\top$ as before, and the stacked system (B4) holds with the b^σ -row of (D6), the additional b_t -row (D2), and

$$\sigma_s = \begin{pmatrix} \sigma_z \\ b\check{\Sigma} \\ b\eta^\sigma + \sigma_f \end{pmatrix}. \quad (\text{D8})$$

Proposition 2 applies, but there are three subtleties. First, the wealth statistics are *backward-looking but not predetermined*: wealth statistics depend on initial conditions and drift restrictions, but their innovations are prescribed revaluations. In the saddle-path condition they are counted with the predetermined variables, so the requirement remains $2 + N_z$ eigenvalues of H with positive real part regardless of b . In the spanned coordinates of Proposition 10, γ_t^b now has diffusion exactly bI (exogenous). Second, the wealth rows of σ_s depend on the composites $\check{\Sigma}$ and η^σ , which are themselves built from the endogenous loadings $\sigma(\tau) = -A(\tau)^\top \sigma_s$; this dependence simply becomes part of the fixed-point algorithm of Appendix B.3; the same a -homotopy applies. Third, the revaluation feedback adds a solvability margin specific to $b \neq 0$: given $A(\cdot)$, the b_t -row of σ_s solves a linear fixed point and the b^σ -row a quadratic one; near the boundary of this margin small changes in loadings trigger large revaluation responses.

D.2 Welfare

The per-period loss (24) holds for any $b \geq 0$: expanding flow utility to second order, $u(C_t(i)) - u(C) = C^{1-\varsigma}[c_t(i) + \frac{1-\varsigma}{2}c_t(i)^2]$ and $v(N_t(i)) - v(N) = N^{1+\varphi}[n_t(i) + \frac{1+\varphi}{2}n_t(i)^2]$, with $N^{1+\varphi} = C^{1-\varsigma}$ by steady-state efficiency; averaging over i with $\int c_t(i)^2 di = c_t^2 + \text{Var}_i \check{c}_t(i)$ (likewise for n); using the cross-sectional Jensen corrections $\log(C_t/C) = c_t + \frac{1}{2} \text{Var}_i \check{c}_t(i)$ and $\log(N_t/N) = n_t + \frac{1}{2} \text{Var}_i \check{n}_t(i)$ in goods and labor market clearing, so that $c_t = y_t - \frac{\vartheta}{2}\pi_t^2 - \Psi_t^S - \Psi_t^i - \frac{1}{2} \text{Var}_i \check{c}_t(i)$ and $n_t = y_t - z_{1,t} - \frac{1}{2} \text{Var}_i \check{n}_t(i)$; collecting the dispersion terms with relative labor supply $\check{n}_t(i) = -\frac{\varsigma}{\varphi}\check{c}_t(i)$ into $-\frac{\varsigma(\varsigma+\varphi)}{2\varphi} \text{Var}_i \check{c}_t(i)$; and finally writing the quadratic in $(y_t, z_{1,t})$ in gap form around the first-best output y_t^n , at which its derivative vanishes identically in $z_{1,t}$, gives $-\frac{1}{2}\mathcal{L}_t$ plus a policy-invariant first-best term. Only preferences, market clearing, production, and labor supply enter; b matters for welfare only through the dynamics of the loss arguments, and the sufficient-statistic representation $\text{Var}_i \check{c}_t(i) = \check{c}_t^{\sigma^\top} \check{\Sigma}^+ \check{c}_t^\sigma$ continues to hold by Proposition 10.

However, the results on attainability of the first-best (Proposition 7) do not carry through to $b \neq 0$.

Proposition 11 (Revaluation-driven dispersion under zero premia). *Let $b \geq 0$ and consider any equilibrium in which risk prices are identically zero, $\Lambda_t \equiv 0$. Then from spanned initial*

relative wealth γ_0^b ,

$$d\gamma_t^b = b dB_t, \quad \gamma_t^c = \frac{\rho\varphi}{\varphi + \varsigma} \gamma_t^b, \quad (\text{D9})$$

and the discounted dispersion component of the social loss is

$$\mathbb{E}_0 \int_0^\infty e^{-\rho t} \frac{\varsigma(\varsigma + \varphi)}{\varphi} \text{Var}_i \check{c}_t(i) dt = \frac{\varsigma\varphi\rho}{\varsigma + \varphi} \gamma_0^{b\top} \check{\Sigma} \gamma_0^b + \frac{\varsigma\varphi}{\varsigma + \varphi} b^2 \text{tr} \check{\Sigma}. \quad (\text{D10})$$

In particular, if $b \neq 0$ and $\check{\Sigma} \neq 0$, the first-best is not attainable by any policy, even from equalized initial wealth.

Proof. With $\Lambda \equiv 0$, the relative Euler equation gives $\mathbb{E}_t d\gamma_t^c = 0$, and the relative consumption function (D4) gives $\gamma_t^c = \frac{\rho\varphi}{\varphi + \varsigma} \gamma_t^b$. Substituting into (D5) gives (D9): relative wealth is a random walk driven by heterogeneous duration exposure (yield volatility does not vanish under perfect arbitrage). Hence $\mathbb{E}_0 \text{Var}_i \check{c}_t(i) = (\frac{\rho\varphi}{\varphi + \varsigma})^2 [\gamma_0^{b\top} \check{\Sigma} \gamma_0^b + b^2 t \text{tr} \check{\Sigma}]$, and integrating against $e^{-\rho t}$ gives (D10).

For the last claim, the first best requires $\check{c}_t(i) \equiv 0$ for all i , i.e. $\check{\Sigma} \gamma^c \equiv 0$. By the relative Euler equation this requires $a \check{\Sigma} \Lambda \equiv 0$, so all premium terms drop out and (D4) gives $\check{\Sigma} \gamma_t^c = \frac{\rho\varphi}{\varphi + \varsigma} \check{\Sigma} \gamma_t^b$. But by (D5), $\check{\Sigma} \gamma_t^b$ has diffusion $b \check{\Sigma} \neq 0$, a contradiction. (The condition $\check{\Sigma} \neq 0$ is generic: $\check{\Sigma} = 0$ requires all households to have the same risk-weighted exposure, i.e. $\sigma(\cdot)$ orthogonal to every $\eta_i(\cdot) - \eta(\cdot)$. With heterogeneous portfolios this essentially requires bond loadings to vanish, i.e. a short rate that does not respond to shocks, under which macroeconomic stabilization fails.) $\square$

E Additional Tables and Figures

Table E1: Risk-Sharing Terms by SCF Wave

SCF wave		$SD_i(D_i)$	wedge	$SD_i(\check{c})$	loss		
					simple	κ_p -consistent	$D_{eq} = 4$
2022	ACM	9.0	303	13.2	1.74	3.30	6.13
	Kim-Wright	9.0	195	6.5	0.43	0.98	1.51
2019	ACM	8.8	294	12.8	1.64	3.11	5.35
	Kim-Wright	8.8	189	6.4	0.40	0.92	1.32

Notes: as in Table 1, Panel B, with ACM and Kim-Wright premia over 1990-2019. $SD_i(D_i)$ is in years; the wedge $SD_i(D_i)SD(\ell)$ is in basis points; $SD_i(\check{c})$ is in percent.




Figure E1: Unconditional Pass-Through by Maturity, Short and Long Horizons

Notes: estimates of $\beta(\tau, h)$ from (37), unconditional specification, for maturities $\tau = 2$ to 20 years at horizons $h = 0$ and $h = 60$ trading days (about three months). OLS on all trading days, 1982-2026. Bands are ± 2 Newey-West standard errors (90-day bandwidth).




Figure E2: Unconditional Pass-Through to the 20-Year Forward at Long Horizons, Rolling Windows

Notes: estimates of $\beta(\tau, h)$ for $\tau = 20, h = 60$ from (37), unconditional specification, over rolling 5-year windows dated by the end of each window. Bands are ± 2 Newey-West standard errors (90-day bandwidth).

Table E2: Targeted and Untargeted Moments

	Data	Model
<i>Panel A: Targeted moments</i>		
SD(x) (%)	1.85	1.85
SD(π), core PCE y/y (%)	0.38	0.38
corr(x, π)	0.53	0.53
SD(household term premium) (bp)	84	84
corr($\Delta_{12}i, \Delta_{12}$ premium)	-0.53	-0.53
<i>Panel B: Untargeted moments</i>		
SD(i) (%)	–	1.31
SD(2-year term premium) (bp)	–	56
SD(10-year term premium) (bp)	–	144
<i>Panel C: Loss under the historical rule (% of consumption)</i>		
output gap	–	0.034
inflation	–	0.071
consumption dispersion	–	0.188
total	–	0.293

Notes: model moments under the historical rule $i_t = 1.5\pi_t + 0.5x_t$. The household term premium is the premium on the household bond portfolio $\eta(\tau)$. Panel C reports the terms of the social loss in percent of steady-state consumption, as in Table 1; the total excludes the rate cost (0.016), which Table 3 includes.

Table E3: Taylor-Type Rules, with and without Balance Sheet Policy

	loss (% of consumption)				volatility			response of i to		
	total	x	π	dispersion	SD(x)	SD(π)	premium	v	u	$a_{\max}/a$
<i>Panel A: Short-rate rules</i>										
Historical Taylor rule	0.309	0.034	0.071	0.188	1.85	0.38	84	0.87	1.52	1.67
Taylor, $\phi_\pi = 3$	0.361	0.028	0.062	0.248	1.67	0.35	100	0.91	2.54	1.72
Taylor, $\phi_x = 0$	0.459	0.188	0.175	0.089	4.33	0.59	53	0.60	1.72	1.16
Gap-stabilizing Taylor	0.357	0.000	0.046	0.283	0.00	0.30	109	1.05	1.17	≥ 3.0
Inflation-stabilizing Taylor	0.469	0.028	0.055	0.353	1.68	0.33	126	0.95	3.82	1.69
Natural-rate rule, $\chi_{i,v} = 1$	0.406	0.021	0.051	0.310	1.45	0.32	116	1.05	1.52	1.42
Optimal simple rule	0.266	0.033	0.079	0.141	1.81	0.40	70	0.83	0.22	–
<i>Panel B: Balance sheet holds $\Lambda_t = 0$, with the short-rate rule</i>										
Natural-rate rule, $\chi_{i,v} = 1$	0.061	0.000	0.045	0.000	0.11	0.30	0	1.05	1.52	∞
$\chi_{i,v} = 1.1$, $\chi_{i,u} = 1.5$	0.062	0.002	0.042	0.000	0.43	0.29	0	1.07	2.43	∞
Historical Taylor rule	0.098	0.021	0.066	0.000	1.43	0.36	0	0.88	1.52	∞
Gap-stabilizing Taylor	0.061	0.000	0.046	0.000	0.00	0.30	0	1.05	1.17	∞

Notes: losses in percent of steady-state consumption; the total includes the rate cost $\psi^i i_t^2$ of Section 6.2, which is not shown separately. SD(x) and SD(π) in percent; “premium” is the standard deviation of the household term premium, in basis points. The responses of i to v and u are the coefficients of the short rate, in equilibrium, on the natural rate and the markup factor. $a_{\max}/a$ is the largest value of a , relative to its calibrated value, for which the equilibrium exists ($\geq$ marks the limit of the search). Taylor-type rules are $i_t = \phi_\pi \pi_t + \phi_x x_t$ with $(\phi_\pi, \phi_x) = (1.5, 0.5)$ (historical), $(3, 0.5)$, $(1.5, 0)$, $(1.5, 500)$ (gap-stabilizing) and $(5, 0.5)$ (inflation-stabilizing); the natural-rate rules add $\chi_{i,v} v_t + \chi_{i,u} u_t$ to the historical rule. Panel B: balance sheet policy holds $\Lambda_t = 0$ at no cost, so that the equilibrium exists for all a .

Table E4: Social Loss by Rate and Balance Sheet Cost

		loss (% of consumption)					
ψ^i	QE cost	Rule	Simple	Nash	Stackelberg	Nash – Stack.	cut (pp)
<i>Panel A: rate cost $\psi^i i_t^2$</i>							
0.1	∞	0.294	0.253	0.205	0.034	0.171	−6.56
	1	0.294	–	0.155	0.034	0.121	−6.48
	1/10	0.294	–	0.071	0.034	0.037	−6.05
	1/240	0.294	–	0.032	0.031	0.001	−2.40
1.87	∞	0.309	0.266	0.110	0.047	0.064	−1.02
	1	0.309	–	0.092	0.046	0.045	−0.98
	1/10	0.309	–	0.057	0.044	0.013	−0.77
	1/240	0.309	–	0.038	0.038	<0.001	−0.12
5	∞	0.336	0.289	0.083	0.055	0.028	−0.45
	1	0.336	–	0.074	0.054	0.020	−0.42
	1/10	0.336	–	0.057	0.052	0.005	−0.30
	1/240	0.336	–	0.046	0.046	<0.001	−0.04
18.7	∞	0.454	0.371	0.074	0.070	0.004	−0.10
	1	0.454	–	0.073	0.070	0.003	−0.09
	1/10	0.454	–	0.069	0.068	0.001	−0.05
	1/240	0.454	–	0.063	0.063	<0.001	0.00
<i>Panel B: rate cost $\psi^i(i_t - v_t)^2$</i>							
530	1	1.803	–	n.a.	0.229	–	−0.05
	1/10	1.803	–	0.093	0.077	0.016	−0.01
	1/240	1.803	–	0.033	0.032	<0.001	0.00

Notes: the loss $\frac{1}{2} E[\mathcal{L}_t]$, including the costs of both instruments, in percent of steady-state consumption. “Rule” is the historical Taylor rule and “Simple” the optimal simple rule, both without balance sheet policy. The QE cost is relative to the reference cost defined in the notes to Table 3; ∞ means no balance sheet policy. Panel B: the rate cost is $\psi^i(i_t - v_t)^2$, so that the short rate is effectively tied to the natural rate v_t ; the historical rule, which does not track v_t , pays a large rate cost. “Cut” is the impact response of the short rate to a demand shock under the Stackelberg plan, in percentage points (Panel B: of $i_t - v_t$). n.a.: the Nash fixed point has no solution.

Table E5: Loss Decomposition for Selected Rate and Balance Sheet Costs

	loss (% of consumption)						relative to the rule			
	total	x	π	dispersion	rate	QE	SD(i)	premium	$ \eta^\sigma $	SD($\hat{\eta}'\Lambda$)
$\psi^i = 1.87$, no QE										
Historical rule	0.309	0.034	0.071	0.188	0.016	–	1.31	84	1.00	1.00
Optimal simple rule	0.266	0.033	0.079	0.141	0.014	–	1.22	70	0.91	0.89
Nash plan	0.110	0.008	0.023	0.072	0.008	–	0.93	46	0.72	0.72
Stackelberg plan	0.047	0.015	0.023	0.002	0.007	–	0.84	6	0.25	0.28
$\psi^i = 1.87$, QE cost: 1/240										
Historical rule	0.309	0.034	0.071	0.188	0.016	0.000	1.31	84	1.00	1.00
Nash plan	0.038	0.007	0.023	<0.001	0.006	0.002	0.81	17	0.74	0.25
Stackelberg plan	0.038	0.008	0.023	<0.001	0.006	0.001	0.78	15	0.66	0.25
$\psi^i = 18.7$, no QE										
Historical rule	0.454	0.034	0.071	0.188	0.161	–	1.31	84	1.00	1.00
Optimal simple rule	0.371	0.065	0.105	0.102	0.099	–	1.03	57	0.81	0.81
Nash plan	0.074	0.021	0.027	0.007	0.019	–	0.45	11	0.35	0.37
Stackelberg plan	0.070	0.025	0.028	0.002	0.015	–	0.40	6	0.25	0.27
$\psi^i = 18.7$, QE cost: 1/240										
Historical rule	0.454	0.034	0.071	0.188	0.161	0.000	1.31	84	1.00	1.00
Nash plan	0.063	0.017	0.026	0.003	0.016	0.001	0.41	14	0.42	0.39
Stackelberg plan	0.063	0.017	0.026	0.003	0.016	0.001	0.41	14	0.41	0.38
$\psi^i = 0.1$, no QE										
Historical rule	0.294	0.034	0.071	0.188	0.001	–	1.31	84	1.00	1.00
Optimal simple rule	0.253	0.029	0.076	0.147	0.001	–	1.26	72	0.92	0.90
Nash plan	0.205	0.006	0.022	0.176	0.001	–	1.37	81	0.97	0.91
Stackelberg plan	0.034	0.009	0.022	0.001	0.002	–	2.22	3	0.16	0.19

Notes: losses in percent of steady-state consumption; SD(i) in percent; “premium” is the standard deviation of the household term premium, in basis points. The last two columns are relative to the historical rule: $|\eta^\sigma|$ is the size of the household portfolio’s loading on aggregate risk, and SD($\hat{\eta}'\Lambda$) is the standard deviation of the price of risk in the direction $\hat{\eta}$ of the historical rule’s household loading.




Figure E3: Impulse Responses under Optimal Policy: Short Rate Only

Notes: as in Figure 2, with the rate cost of Section 6.2 and no balance sheet policy.




Figure E4: Impulse Responses under Optimal Policy: Nearly Free Short Rate

Notes: as in Figure 2, with the lowest rate cost of Table E4 and no balance sheet policy.




Figure E5: Impulse Responses under Optimal Policy: Costly Short Rate, Cheap Balance Sheet

Notes: as in Figure 2, with the highest rate cost of Panel A of Table E4 and the low balance sheet cost.




Figure E6: Impulse Responses under Optimal Policy: Short Rate Tied to the Natural Rate

Notes: as in Figure 2, with the rate cost $\psi^i(i_t - v_t)^2$ of Panel B of Table E4 and the low balance sheet cost. Under the plans the short rate tracks the natural rate v_t .