Monetary and financial policy with privately optimal risk taking *

Alfred Duncan† University of Kent

João Pedro De Camargo Mainente‡ University of Kent

Charles Nolan§ University of Glasgow

December 2025

We study how monetary and macroprudential policy should interact in a New Keynesian model with privately optimal, but socially risky, leverage. Entrepreneurs finance idiosyncratically risky projects under imperfect state verification generating endogenous leverage, factor wedges and an equity premium, even when aggregate risk is traded. Higher household risk aversion produces a "paradox of safety": stronger demand for safe assets concentrates risk on entrepreneurial balance sheets, raising the welfare cost of cycles. Macroprudential tools act on leverage and balance-sheet exposures but cannot fully internalise associated externalities. Thus, a financial-stability role for monetary policy remains–especially in the face of temporary shocks–yet a regime that leans too hard on this role risks losing control of inflation.

JEL Classification: D52, E32, E52. **Keywords**: Macroeconomics, Incomplete Markets, Monetary Policy.

---

*We would like to thank participants at the Baltic Economic Conference 2023, the Association of Southern European Economic Theorists Conference 2023, and the Scottish Economic Society Conference 2024, for helpful comments and discussions. All errors are our own.

†Corresponding author. Email: a.j.m.duncan@kent.ac.uk

‡Email: J.P.De-camargo-mainente@kent.ac.uk

§Email: charles.nolan@glasgow.ac.uk

Introduction

If macroprudential policy can implement a constrained-efficient allocation of risk and leverage, and if markets for aggregate risk are open, does monetary policy have a role beyond stabilising inflation and the output gap? Should financial stability considerations be left almost entirely to macroprudential policy? Put more prosaically: in a world where prudential regulators wield powerful macroprudential tools, should monetary policy "take financial conditions into account"?

The answer is not obvious. One strand of recent work suggests a neat division of labour: macroprudential tools internalise the externalities associated with systemic risk, while monetary policy focuses on nominal distortions. Another strand argues that monetary policy has a more fundamental role in shaping risk-taking, leverage and financial stability, so that "sticking to inflation" is neither realistic nor desirable. We take a simple New Keynesian benchmark and ask: under what conditions, and through which channels, does monetary policy continue to matter for financial conditions when (i) aggregate risk can be traded and (ii) a prudential authority is actively using its instruments?

We build a New Keynesian model with monopolistic competition and either flexible or sticky prices, augmented with a microfounded financial friction based on an imperfect state verification problem. Entrepreneurs finance binary risky projects with a mix of inside wealth and outside funds from households. Project outcomes contain both aggregate and idiosyncratic components. Idiosyncratic outcomes are privately observed and can be verified only through a costly and imperfect audit technology. At the macro level, households and entrepreneurs can trade state-contingent claims on aggregate shocks, so business-cycle risk can, in principle, be widely shared.

This structure delivers three ingredients that are central to our analysis. First, entrepreneurs choose an interior leverage ratio: because audits are imperfect and costly, pushing arbitrarily much risk onto lenders raises expected audit costs and becomes privately unattractive. Second, entrepreneurs optimally discount the marginal products of labour and capital relative to factor prices, generating a wedge in factor markets that grows with both leverage and idiosyncratic risk. Third, the presence of undiversifiable project risk implies an endogenous equity premium that compensates households for bearing aggregate risk in the shadow of imperfect idiosyncratic insurance.

On top of this we introduce a simple but microfounded representation of macroprudential policy. A prudential authority can constrain the way aggregate shocks feed into entrepreneurial leverage and net worth—for example, by imposing risk-based capital requirements, exposure limits or stress-test-style constraints—but cannot fully re-engineer risk sharing because entrepreneurs can always hide wealth in a safe asset. At a high level, we treat macroprudential policy as about where aggregate risk sits when the downturn arrives, rather than about any single tool or institutional arrangement. Instead of modelling a detailed regulatory architecture, we ask what a planner with realistic information constraints can achieve by acting on balance-sheet exposures ex ante. The key limitation is that entrepreneurs can hide resources "off balance sheet", so the policymaker cannot impose different intertemporal prices on them than on households without being circumvented. That rules out a whole class of interest-rate-based interventions. What remains are constraints on observable leverage and risk-weighted exposures that make the system less fragile when bad shocks hit. In the model, we capture those constraints as a macroprudential wedge in the aggregate risk-sharing relation; the algebra is in the appendices, but the economics is simple: macroprudential policy is more like a building code than a bailout, shaping resilience before the storm rather than moving wealth around afterwards.

We use this structure to make three main points. First, the model generates what we call a *paradox of safety*: when households become more risk averse and try to insure themselves by holding safe or countercyclical assets, aggregate risk is pushed onto entrepreneurial balance sheets, amplifying fluctuations in leverage, hours and consumption. Even with open markets for aggregate risk, attempts to "be safer" at the individual level can raise the volatility of equilibrium consumption. One convenient way to see this is through a Lucas-style calculation: in a flexible-price benchmark without prudential intervention, the familiar consumption-equivalent cost of business cycles inherits Lucas's form but the mapping from shocks to consumption volatility is distorted by leverage and the factor wedge. As risk aversion rises, this mapping becomes more volatile rather than less. Our full welfare analysis later in the paper shows that this is not just a reduced-form curiosity but reflects a genuine deterioration in risk sharing between households and entrepreneurs.

Second, we characterise the joint roles of monetary and macroprudential policy when aggregate risk markets are open and prudential tools are used actively. Financial stress appears as a cost-push-like disturbance in the Phillips curve, via the factor wedge, and as a demand wedge in the IS curve, via the equity premium and the distribution of wealth. Even when macroprudential policy is aimed at a constrained optimum for leverage and risk, the optimal monetary reaction function continues to respond to financial conditions. The relative roles of the two sets of instruments depend critically on the persistence and nature of shocks: monetary accommodation is well suited to transitory disturbances, whereas prudential tools are better suited to persistent changes in risk and technology. Relying too heavily on prudential policy to manage short-lived shocks can generate undesirable medium-run fluctuations in leverage and consumption inequality; relying too heavily on monetary policy to manage persistent risk can entrench financial fragilities.

Finally, we show that assigning monetary policy a narrowly defined financial-stability objective can be dangerous, even in the presence of active macroprudential tools. In our framework, a central bank that insists on keeping financial conditions "too safe" at all times can generate permanently high inflation in response to a temporary recessionary shock. Prudential policy can dampen but cannot remove this inflationary bias, because the underlying information frictions limit how cleanly nominal and financial objectives can be separated. In this sense, the model cautions against the view that macroprudential policy can fully "take financial stability off the table" for monetary policy, even when aggregate risk markets are open and prudential tools are used in a sophisticated way.

Related literature

Our paper connects to several strands of work on monetary policy, macroprudential policy and risk sharing. Martin et al. (2021) and Laeven et al. (2022) provide recent and comprehensive overviews. The post-2008 consensus that "monetary policy cannot do it all" led to broad support for macroprudential tools.[1] A central question in this literature is how these tools should be coordinated with monetary policy, and

[1]See, for example, Allen and Rogoff (2011) on the need for "as many tools as possible" when dealing with asset-price booms.

when—if ever—a clean separation of roles is appropriate.

One strand of theory argues that, if macroprudential policy can efficiently target the relevant externalities, then monetary policy should largely ignore financial conditions and stick to traditional objectives. Representative contributions include Korinek and Simsek (2016) and Caballero and Simsek (2019), where optimally designed prudential instruments internalise systemic-risk externalities and the central bank is assigned to stabilising inflation and the output gap. In these frameworks, any residual financial-stability role for monetary policy reflects constraints or imperfections in the prudential toolkit.

A second strand emphasises that monetary policy has a more structural financial-stability role. Stein (2012, 2013) argue that short-term interest rates influence risk-taking and leverage in ways that cannot easily be replicated by macroprudential measures. Quantitative models in the tradition of Gertler and Kiyotaki (2010) and Akinci et al. (2021) study regimes in which monetary policy explicitly leans against financial stress. Our contribution is closest in spirit to this work, but differs in that we allow for open aggregate risk markets and an explicit macroprudential wedge, and provide a welfare interpretation of how risk aversion, leverage and prudential policy interact.

A third, closely related strand studies how monetary policy can improve risk sharing when contracts are imperfect. Bhandari et al. (2021) and Sheedy (2014) show that countercyclical monetary policy can replicate missing state-contingent markets when nominal contracts cannot be indexed to aggregate conditions. In those models, private agents and the policymaker agree on the desired state contingency; the role of policy is to recreate it through prices. In our framework, by contrast, contracts can be written contingent on aggregate states and markets for macro risk are open. The key friction is at the micro level—imperfect verification of idiosyncratic outcomes—which generates wedges in factor markets and in aggregate risk sharing. Monetary policy improves welfare not by creating missing aggregate markets, but by changing the marginal value of entrepreneurial equity and hence the privately optimal leverage choice.

Our modelling of the financial friction builds on the imperfect state verification literature. Duncan and Nolan (2019) extend Townsend (1979), addressing critiques by Border and Sobel (1987) and Mookherjee and Png (1989). This class of models is useful here for three reasons. First, it provides robust microfoundations for debt-like contracts even when agents can adjust their exposure to macroeconomic risk. Second, it implies that financial conditions respond to both productivity shocks and risk/uncertainty shocks, even when aggregate risk markets are open. Third, it yields a reduced-form representation with an interior leverage choice and an explicit factor wedge, which can be embedded in a New Keynesian environment at relatively low cost.

Our treatment of macroprudential externalities is related to Di Tella (2017), Duncan and Nolan (2021), Farhi and Werning (2016) and Schmitt-Grohe and Uribe (2012). As in Di Tella (2017), private agents do not internalise the higher social costs of financial stress when leverage is high in downturns. In our setup, however, this externality arises not only from "pure" risk shocks but also from technology and monetary policy shocks, because leverage feeds directly into marginal costs and aggregate demand. As in Farhi and Werning (2016) and Schmitt-Grohe and Uribe (2012), changes in the distribution of wealth create aggregate-demand externalities; in our setting, these interact with the equity-premium and macroprudential wedges to shape the IS curve.

Finally, our "paradox of safety" is closely related to the "safety trap" described by Caballero and Farhi (2017). In canonical New Keynesian or real-business-cycle models, higher household risk aversion typically dampens volatility via stronger consumption smoothing and wealth effects on labour supply. In our model, a stronger taste for safety leads households to load up on safe or countercyclical assets. Entrepreneurs take the other side of these trades, and aggregate risk becomes concentrated on leveraged balance sheets. The cost of business cycles rises with risk aversion once this mechanism is taken into account.

Preview of results

In the next Section the model is introduced, then Section 2 derives the paradox of safety. Higher household risk aversion pushes aggregate risk onto entrepreneurial balance sheets, amplifying fluctuations in labour demand, output and consumption. Section 3 characterises optimal monetary and macroprudential policy under log utility for both households and entrepreneurs. Financial stress—summarised by the equity premium and associated factor wedge—acts like a cost-push disturbance in the Phillips curve and a demand disturbance in the IS curve. Even with constrained-efficient prudential policies and active aggregate risk markets, optimal monetary policy still responds to financial conditions. Section 4 shows that accommodative monetary policy is best suited to transitory shocks and macroprudential policy better suited to persistent disturbances; relying on prudential tools to manage short-lived shocks induces long-lived swings in leverage and consumption inequality. Section 5 considers a monetary authority targeting financial stability directly, in the spirit of Akinci et al. (2021). Keeping financial stress low risks permanently higher inflation after a temporary recessionary shock, as the required monetary accommodation is never fully unwound. Prudential policy can dampen but not eliminate this inflationary bias, as the information frictions prevent a clean separation of nominal and financial objectives. Section 6 turns to uncertainty shocks. The same ingredients that generate the paradox of safety—strong demand for safe assets, leverage-induced concentration of risk, and an increasing welfare cost of cycles—mean risk shocks influence the optimal policy mix. Appendices A - D provide details of the core derivations.[2]

1 The model

The model consists of a representative worker–household and a continuum of entrepreneurs. The household supplies labour and capital and owns the retail sector. Entrepreneurs operate a risky technology with both aggregate and idiosyncratic productivity components. They finance production partly from their own net worth and partly by borrowing from the household. Because entrepreneurs observe their firm-specific shocks privately, outside investors cannot obtain full insurance against idiosyncratic risk at zero cost. Risk sharing is therefore incomplete at the firm level, even though agents can trade securities contingent on aggregate states of the world.

Monopolistically competitive retailers buy a homogeneous wholesale good from entrepreneurs and sell differentiated final goods to both households and entrepreneurs. Prices in the retail sector are subject to Calvo (1983) stickiness. Monetary policy sets the nominal interest rate, while a macroprudential authority chooses a state-

[2]Comprehensive appendices are also available on-line.

contingent policy that shapes the allocation of aggregate risk between households and entrepreneurs.

We first describe the agents, technologies, markets, and shocks in levels. We then explain how risk sharing and leverage are determined, and how macroprudential policy appears as a wedge in the aggregate risk-sharing condition. Finally, we present the log-linear equilibrium conditions that form the three-equation core of the model: an IS curve, a Phillips curve, and a leverage curve.

1.1 Agents, technologies and shocks

1.1.1 Households

The representative worker–household enters period t with real financial wealth q_t. It consumes c_t, supplies labour n_t, and trades a complete set of securities contingent on aggregate states $s \in S$. Preferences are given by

$$v(q_t) = \max_{\{c_t, n_t, q_{t+1}, z_t(\cdot)\}} \mathbb{E}_t \left[\frac{c_t^{1-\sigma}}{1-\sigma} - \frac{n_t^{1+\varphi}}{1+\varphi} + \beta v(q_{t+1}) \right],$$

where $\sigma > 0$ is the household's coefficient of relative risk aversion, $\varphi > 0$ is the inverse Frisch elasticity of labour supply and $0 < \beta < 1$ is the discount factor.

The household's budget constraint in real terms is

$$q_{t+1} = (1 + r_{t+1})q_t + w_t n_t + \Pi_t - c_t - \int_{s \in S} p_t(s) z_t(s)\, ds + z_t(s_{t+1}),$$

where r_{t+1} is the gross real return to their period t wealth, w_t is the real wage, and Π_t denotes aggregate profits from the retail sector (which is owned by the household). The Arrow–Debreu securities $z_t(s)$ pay one unit of the consumption good in aggregate state $s_{t+1} = s$. In our analysis s can include productivity shocks, uncertainty shocks, markup and cost push shocks, government purchases shocks, and monetary policy shocks. In practice, we consider this aggregate risk trade as a proxy for a wide range of financial decisions that shift agents' exposure to business cycle shocks, and shift risks between groups.[3] Wealth q_{t+1} is determined by

[3]The decision between mortgage fixed rate terms is an example. A longer fixed rate term will reduce the household's exposure to aggregate shocks that result in high interest rates, which would

decisions made in period t, but is contingent on time $t+1$ outcomes of exogeneous state variables, and is therefore measurable in the $t+1$ state-space.

1.1.2 Entrepreneurs and production

There is a continuum of measure one of entrepreneurs, indexed by their idiosyncratic state $\theta_t \in \Theta$. Entrepreneurs have log utility over their own consumption and discount factor β^e:

$$v^e(q_t^e) = \max_{z_t^e, c_t^e, q_{t+1}^e} \mathbb{E}_{\Theta,t} \left\{ \log c_t^e + \beta^e v^e(q_{t+1}^e) \right\}$$

subject to

$$q_{t+1}^e = R_t(\theta_t, s_t) q_t^e - c_t^e - \int_{s \in S} p_t(s) z_t^e(s_t) ds + z_t^e(s_{t+1})$$

Superscript e denotes the entrepreneur, and $v^e(q^e)$ is the value function. q_t^e denotes entrepreneurial wealth at the start of period t, and expectations $\mathbb{E}_{\Theta,t}$ are taken over the distribution of idiosyncratic shocks θ_t and aggregate shocks. $R(\theta, s)$ is the return to entrepreneurial wealth, q_t^e, and is the outcome of a privately optimal external finance contract, determined at the beginning of the period, and conditional on idiosyncratic states realised within the period. Trade in aggregate risk markets is captured by the quantities $z^e(s)$, denoting the amount purchased of an asset with payoff 1 conditional upon the future state of the world being realised as state s. The current period price of this security is denoted $p(s)$. As indicated earlier, trade in securities indexed by the aggregate state are not hampered by any problem of asymmetric information; unlike idiosyncratic states, aggregate states are costlessly observed and verified by all agents. These markets are active.

Each entrepreneur operates a constant-returns-to-scale wholesale production technology

$$y_t^e(\theta_t) = A_t \, \theta_t \, (k_t^e)^{\alpha} (n_t^e)^{1-\alpha},$$

be harmful to households with short mortgage fixed rate terms. In this way, a longer mortgage fixed rate provides insurance against aggregate shocks that increase interest rates. This insurance doesn't remove risk from the aggregate economy, but it does shift the risk from the mortgage borrowing household to other agents who are happy to accept interest rate risk at an agreeable price.

where A_t is the aggregate technology level, θ_t is an idiosyncratic productivity shock, k_t^e is the capital rented from the household, n_t^e is hired labour, and $0 < \alpha < 1$ is the capital share.

Entrepreneurs finance their input purchases using their own wealth q_t^e and external funds raised from the household. Realised entrepreneurial wealth at the start of the next period, q_{t+1}^e, is the net payoff from the production and financial contract, plus returns on aggregate risk securities $z_t^e(\cdot)$.

They choose a financial contract and an exposure to aggregate risks before observing the realisation of the idiosyncratic shock θ_t. The contract specifies state-contingent repayments to outside investors, subject to incentive-compatibility and limited-enforcement constraints that arise from private information about θ_t and costly monitoring of misreports.

Specifically, our financial friction is derived from a costly and imperfect state verification problem in the spirit of Duncan and Nolan (2019). Entrepreneurs finance risky projects using their own wealth and external funds from the representative household, but individual cash flows are only imperfectly observable: lenders must audit reported defaults at a cost, and the audit technology makes type-I errors (truthful 'low' reports are incorrectly tagged as untrue and penalised). The optimal static contract is therefore debt-like, with a promised repayment in the good state and a lower recovery in the bad state, together with costly audits of defaults. Because audits are imperfect and anonymity prevents history-dependent contracts, entrepreneurs choose an interior leverage ratio: borrowing more both scales up expected profits and raises the expected cost of monitoring. In aggregate, the friction can be summarised by an endogenous leverage ratio l_t, a factor-market wedge τ_t that increases with leverage and idiosyncratic risk, and an equity risk premium ρ_t that compensates households for bearing entrepreneurial risk.

From individual contracts to aggregate wedges

At the micro level, each entrepreneur solves a static contracting and factor-demand problem. Given wages, the rental rate of capital and the price of equity, entrepreneur i chooses factor inputs, a scale of external finance and a state-contingent repayment schedule to maximise expected utility, subject to: (i) a budget constraint; (ii) an incentive-compatibility constraint that prevents low-type projects from mimicking high types; and (iii) a participation (break-even) constraint for households. Appendix A shows that the optimal contract has three key features: it is debt-like (a high promised repayment in good states and a lower recovery in bad states with audit), it implies an interior choice of leverage (because pushing more risk onto lenders raises expected audit costs), and it induces wedges between expected marginal products and factor prices.

These wedges can be written as

$$\tau_{Nti} = \frac{\bar{Y}_{Nti} - W_t}{\bar{Y}_{Nti}}, \qquad \tau_{Kti} = \frac{\bar{Y}_{Kti} - r_{Kti}}{\bar{Y}_{Kti}},$$

where $\bar{Y}_{jti}$ is the expected marginal product of factor j for entrepreneur i. The first-order conditions imply that both wedges increase with the scale of the risky project and with idiosyncratic risk, and that $\tau_{Nti} < \tau_{Kti}$ because audit costs load on capital but not on labour. The same contract pins down an optimal individual leverage ratio L_{ti} and an equity premium ρ_t as functions of these wedges and of the spread between good and bad project outcomes (see Appendix A).

Constant returns to scale in production and monitoring, together with homothetic preferences, imply that the optimal contract is scale-free: leverage, wedges and the equity premium do not depend on individual wealth. All entrepreneurs choose the same leverage ratio and face the same factor wedges, so we can drop the i index and work with aggregate leverage and aggregate wedges. Around the steady state, this aggregation delivers simple reduced-form relationships of the form

$$\tau_t \approx \tau_l \, l_t + \tau_\xi \, \xi_t, \qquad \rho_t \approx \psi \, (l_t + \xi_t),$$

where l_t is (log) leverage, ξ_t measures idiosyncratic risk (see below), and $\tau_l, \tau_\xi, \psi > 0$ are functions of the contract primitives. These two objects — a factor wedge increasing in leverage and risk, and an equity premium increasing in leverage and risk — are the only channels through which the underlying contracting problem enters the IS and Phillips curves and, ultimately, the welfare analysis. The algebra behind them is in the appendix; the main text uses only these aggregated, reduced-form relationships.

1.1.3 Retailers, price setting and monetary policy

A unit mass of monopolistically competitive retailers transforms the homogeneous wholesale good into differentiated final goods using a one-for-one technology. Retailers are owned by the household and face Calvo pricing frictions. A retailer that can reset its price at date t chooses a common price P_t^* to maximise the present value of expected profits, given nominal marginal cost and the probability of not being able to adjust in future periods. The familiar New Keynesian Phillips curve arises after log-linearising the optimal price-setting condition.

Monetary policy sets the short-term nominal interest rate i_t. In later sections we will consider both simple interest-rate rules and fully optimal monetary policy. The real interest rate is $r_t = i_t - E_t\pi_{t+1}$, where π_t is inflation.

1.1.4 Shocks

The economy is subject to aggregate technology shocks and aggregate uncertainty shocks. We denote the log-deviation of technology from its steady state by a_t and the uncertainty shock by ξ_t. For simplicity, ξ_t follows a stationary AR(1) process,

$$\xi_t = \rho_\xi \xi_{t-1} + \epsilon_{\xi,t}, \qquad |\rho_\xi| < 1,$$

and measures persistent fluctuations in the dispersion of idiosyncratic productivity draws θ_t across entrepreneurs. Higher ξ_t raises the variance of firm-specific profitability and therefore the cost of bearing idiosyncratic risk.

1.2 Risk sharing, leverage and macroprudential policy

1.2.1 Leverage and the equity risk premium

Entrepreneurs are risk averse and cannot fully offload firm-specific risk to the household because of private information and costly monitoring. In equilibrium, entrepreneurs pay an equity risk premium ρ_t to the household in order to attract outside funds. This premium compensates the household for bearing some of the aggregate component of entrepreneurial risk and for imperfect diversification of idiosyncratic risk.

We measure leverage by the ratio of expected entrepreneurial output to entrepreneurial net worth. After log-linearisation around the deterministic steady state, leverage can be written as

$$l_t = x_t - c_t^e + \rho_t,$$

where x_t is the log-deviation of aggregate output from steady state, c_t^e is the log of average entrepreneurial consumption, and ρ_t is the log equity risk premium. Given entrepreneurial net worth, higher leverage corresponds to raising more external finance from the household, and it uniquely determines the distribution of consumption between entrepreneurs and the household in equilibrium.

1.2.2 Aggregate risk markets and risk sharing

As noted, households and entrepreneurs can trade Arrow–Debreu securities contingent on aggregate states $s \in S$. These claims allow agents to share aggregate business cycle risk, while idiosyncratic risk remains imperfectly insured because it is privately observed at the firm level and only partly contractible through the imperfect-audit technology described in Appendix A.

Recall that c_t is household consumption and c_t^e is entrepreneurial consumption. In a frictionless complete-markets benchmark with household CRRA preferences and no private information, competitive risk sharing over aggregate states implies the familiar condition

$$\sigma\, \Delta c_t = \Delta c_t^e,$$

so that the growth rates of consumption across types are perfectly aligned up to the risk aversion parameter.

In our environment, incomplete insurance against privately observed idiosyncratic shocks implies that entrepreneurs bear extra risk at the firm level. To induce households to hold claims on entrepreneurial cash flows, they must therefore pay an endogenous equity risk premium ρ_t, which is increasing in leverage and in the dispersion of idiosyncratic productivity (Appendix A). At the aggregate level this shows up as a wedge in the risk-sharing condition:

$$\sigma\,\Delta c_t = \Delta c_t^e - \rho_t.$$

When $\rho_t > 0$, entrepreneurial consumption moves more than household consumption across aggregate states: the combination of leverage and imperfect verification concentrates aggregate risk on entrepreneurial balance sheets. In the limiting case with no leverage effects and no idiosyncratic risk, ρ_t vanishes and we return to the textbook complete-markets condition.

1.2.3 Macroprudential policy and the wedge δ_t

Macroprudential policy in our model changes where aggregate risk sits in the economy by restricting entrepreneurs' exposure to macroeconomic shocks. We are not interested in the detailed implementation of, say, Basel-style regulations. Instead, we take a mechanism-design approach: a macroprudential authority chooses an ex ante state-contingent constraint on feasible financial contracts, subject to the same information and enforcement frictions as private agents. The resulting allocation is best read as the constrained-efficient benchmark for macroprudential policy; our question is then how monetary policy should behave around that benchmark, not as against precise regulatory tool is used.

Formally, the authority can tilt the risk profile of entrepreneurial balance sheets—by limiting leverage or risky exposures in bad states—but cannot freely reallocate wealth between households and entrepreneurs. Entrepreneurs, as we explain below, have access to "hidden storage": they can secretly save at the risk-free rate and conceal part of their income and wealth from outside financiers and the regulator. Any attempt to engineer permanent, deterministic transfers between groups would therefore be undone in equilibrium. In this sense macroprudential policy in our framework is closer to a building code than to a bailout: it cannot rebuild balance sheets after a shock, but it can make them less fragile before the shock arrives.

The setup is isomorphic to a model in which banks make risky, state-contingent loans to firms. High promised returns in recessions require banks to take on substantial undiversifiable risk, so each bank chooses leverage that is privately optimal but too high from a social perspective: system-wide balance sheets are too exposed going into a downturn, and lending capacity is too weak once the downturn hits. A regulator acting optimally would then rein in risky lending—for example through higher capital requirements or tighter loan-to-value limits—applied symmetrically across institutions. For tractability, we work directly with risk-sharing constraints between households and entrepreneurs, but the macroprudential problem is the same: design state-contingent bounds on risky exposures, not pick a particular instrument, and then study how that constrained-efficient arrangement interacts with monetary policy.

The impact of macroprudential policy on risk sharing is summarised by a wedge δ_t that enters the aggregate risk-sharing condition:

$$\sigma \Delta c_t = \Delta c_t^e - \rho_t - \big(1 + \sigma\omega(1-\psi)\big)\delta_t, \tag{1.6}$$

where ω is the steady-state ratio of household to entrepreneurial consumption, and ψ is the elasticity of the equity risk premium with respect to leverage and risk (defined more precisely in Appendix A).

Relative to the benchmark condition $\sigma\Delta c_t = \Delta c_t^e$, equation (1.6) has two additional terms. The first, ρ_t, reflects the endogenous equity premium that arises because entrepreneurial returns combine aggregate and non-diversifiable idiosyncratic risk. The second, proportional to δ_t, captures how macroprudential policy tilts the exposure of entrepreneurial net worth and leverage to aggregate shocks. Rearranging (1.6) gives

$$\Delta c_t^e - \sigma \Delta c_t = \rho_t + \big(1 + \sigma\omega(1-\psi)\big)\delta_t,$$

so that positive δ_t corresponds to policies that load more aggregate risk onto entrepreneurs (increasing the gap between entrepreneurial and household consumption growth), while negative δ_t corresponds to policies that buffer entrepreneurial balance sheets and compress that gap.

By construction, δ_t parameterises the effect of a wide class of feasible prudential instruments on the joint dynamics of (c_t, c_t^e, ρ_t) without committing to a single regulatory technology. Hidden storage then implies that only the unanticipated component of δ_t can matter for intertemporal risk sharing; we formalise this in Lemma 1 below.

1.2.4 Hidden storage and the limits of prudential policy

Entrepreneurs can hide resources from outside investors and from the prudential authority at the risk-free rate. We formalise this through the following constraint.

Constraint 1 (Hidden storage). Entrepreneurs can secretly store wealth across periods at the market risk-free real interest rate. Within periods, they can hide part of their realised income and consumption from external financiers. Across periods, they can hide accumulated wealth from the macroprudential authority.

Hidden storage implies that the policymaker cannot directly impose group-specific interest rates or confiscate hidden wealth. Prudential policy must instead operate through restrictions on the risk characteristics of observable contracts (for example, constraints on promised repayments or on risk-weighted assets). As a result, macroprudential policy can alter the exposure of entrepreneurial net worth to unanticipated shocks but cannot, in expectation, redistribute wealth permanently between households and entrepreneurs.

Given Constraint 1, both groups discount future consumption at the same risk-free rate in expectation, and intertemporal risk sharing over aggregate states implies

$$\sigma E_t[\Delta c_{t+1}] = E_t[\Delta c^e_{t+1}] - E_t[\rho_{t+1}]. \tag{1.7}$$

Combining equations (1.6) and (1.7) yields the key restriction on macroprudential policy:

Lemma 1 (Unpredictability of the macroprudential wedge). The macroprudential wedge is a martingale difference:

$$E_t[\delta_{t+1}] = 0.$$

Lemma 1 says that macroprudential policy can change how entrepreneurial net worth responds to *unanticipated* shocks, but cannot reliably engineer expected transfers between households and entrepreneurs. Equivalently, macroprudential policy can make crises less damaging ex post by requiring balance sheets to be built "above the flood-line" ex ante, but once the water has risen and net worth has already been hit, the damage cannot be undone with prudential tools alone.

1.3 Log-linear equilibrium conditions

We now collect the log-linear equilibrium conditions that define the core New Keynesian part of the model augmented with leverage and uncertainty.

Let x_t denote the (log) deviation of aggregate output from its steady state, π_t inflation, i_t the nominal interest rate, ξ_t the uncertainty shock, and l_t the (log) leverage ratio. All variables are expressed as deviations from their deterministic steady state unless stated otherwise. The principal equations are:

IS curve. The IS curve is derived from the household's intertemporal Euler equation, combined with the definition of the output gap and the law of motion for the household's consumption share. It takes the form

$$x_t = E_t[x_{t+1}] - \frac{1}{\sigma}(i_t - E_t[\pi_{t+1}]) + \omega(1-\psi)E_t[\Delta l_{t+1}] + \omega\psi(1-\rho_\xi)\xi_t. \quad (1.1)$$

Relative to the standard New Keynesian IS curve, two additional terms appear. First, expected changes in leverage $E_t[\Delta l_{t+1}]$ affect aggregate demand because leverage determines the distribution of consumption between households and entrepreneurs, whose intertemporal elasticities differ. Second, the uncertainty shock ξ_t affects aggregate demand directly, since higher uncertainty shifts consumption towards entrepreneurs and changes effective risk premia.

Phillips curve. Log-linearising Calvo price setting yields

$$\pi_t = \beta E_t[\pi_{t+1}] + \lambda p_t^p, \quad (1.2)$$

where $\lambda := \frac{(1-\theta)(1-\beta\theta)}{\theta}\frac{1-\alpha}{1-\alpha+\alpha\epsilon}$ is the standard composite parameter, θ is the Calvo parameter, α is the capital share, ϵ is the elasticity of substitution across varieties, and p_t^p denotes (log) real marginal costs (producer prices, hence the superscript).

Real marginal costs are given by

$$p_t^p = (\tilde{\sigma} + \chi)x_t - \chi a_t + \mu_l l_t - \mu_\xi \xi_t, \tag{1.4}$$

where

$$\chi := \frac{1+\phi}{1-\alpha}, \qquad \tilde{\sigma} := \sigma - 1,$$

and $\mu_l > 0$, $\mu_\xi > 0$ are composite elasticities that capture both wealth-effect and labour-wedge channels. Higher leverage increases marginal costs through two mechanisms: it raises the household's consumption share and thereby the real wage demanded (a wealth effect), and it increases the sensitivity of entrepreneurial profits to firm-specific risk, amplifying the labour wedge. Higher uncertainty ξ_t increases the dispersion of entrepreneurial outcomes and thus the shadow cost of bearing idiosyncratic risk.

Leverage curve. The leverage curve summarises how leverage responds to technology, uncertainty, aggregate demand and macroprudential policy. Using the risk-sharing condition (1.6), the definition of leverage, and the dynamics of the equity risk premium, we obtain

$$l_t = \phi l_{t-1} + (1-\phi)\left(\omega\sigma\Delta\xi_t - \xi_{t-1} - \frac{\tilde{\sigma}}{\psi}\Delta x_t\right) - \delta_t, \tag{1.3}$$

where $0 < \phi < 1$ governs the persistence of leverage in response to shocks, and $(1-\phi)$ is the elasticity of the equity risk premium with respect to the ratio of marginal utilities. The parameter $\psi > 0$ is the elasticity of the equity risk premium with respect to leverage and uncertainty (derived in Appendix A), and $\tilde{\sigma} = \sigma - 1$ emphasises that the leverage dynamics are driven by the difference between household and entrepreneurial risk aversion.

Equation (1.3) highlights three forces. First, leverage is persistent because higher leverage today increases the vulnerability of entrepreneurial balance sheets going forward. Second, leverage rises with increases in uncertainty and with falls in output (through Δx_t), reflecting the concentration of risk on entrepreneurial balance sheets when households insure themselves against aggregate shocks. Third, macroprudential policy enters directly via δ_t: a countercyclical prudential stance that tightens in booms and relaxes in busts corresponds to a δ_t process that dampens the response of l_t to ξ_t and x_t.

Equations (1.1)–(1.3), together with the laws of motion for a_t and ξ_t and a monetary policy rule for i_t, form a closed system describing the joint dynamics of output, inflation and leverage. In subsequent sections we use this system to study the paradox of safety and to characterise optimal monetary and macroprudential policies.

1.4 Welfare

To evaluate policy we use a utilitarian social welfare criterion that aggregates the utilities of households and entrepreneurs using Negishi (1960) weights. Let U_t^h and U_t^e denote the period utilities of the representative household and of a representative entrepreneur, respectively. Social welfare is

$$W_0 = E_0 \sum_{t=0}^{\infty} \beta^t \left[U_t^h + \omega\, U_t^e\right],$$

where ω is the steady-state ratio of household to entrepreneurial consumption. We choose the Negishi weight ω so that, conditional on the monitoring technology and other frictions, the deterministic competitive steady state also solves the planner's problem. Intuitively, the policymaker has no desire to reallocate wealth permanently between households and entrepreneurs in the steady state; policy is motivated purely by efficiency around this benchmark.

We log-linearise equilibrium conditions around a zero-inflation deterministic steady state and take a second-order approximation to W_0 using standard methods. Up to a second order, overall welfare can be written as a constant minus the discounted sum of a quadratic per-period loss Λ_t. The derivation is provided in Appendix C; here we report the resulting loss function:

$$\begin{aligned}2\Lambda_t = {} & (1+\omega)\,\frac{\epsilon}{\lambda}\,\pi_t^2 + (1+\omega)\,\chi\, x_t(x_t - 2a_t) + \tilde{\sigma} x_t^2 \\ & + \omega\Big[(1+\sigma\omega)(1-\psi)\, l_t + \tilde{\sigma} x_t\Big]\Big[(1-\psi)\, l_t - \psi \xi_t\Big] + \omega\, l_t(\kappa_{ll} l_t + \kappa_{l\xi}\xi_t) + \text{t.i.p.}\end{aligned} \tag{1.8}$$

Here ϵ is the elasticity of substitution across varieties, λ is the usual Calvo composite, $\chi := \frac{1+\phi}{1-\alpha}$, and $\tilde{\sigma} := \sigma - 1$. The parameters $\kappa_{ll} > 0$ and $\kappa_{l\xi} > 0$ are composite elasticities that capture how the cross-sectional dispersion of entrepreneurial consumption responds to leverage and to the uncertainty shock; they are functions of the underlying monitoring technology and the distribution of idiosyncratic shocks (see Appendix C). The term "t.i.p." denotes components that are independent of policy and can therefore be ignored for optimal policy design.

The first line of (1.8) is closely related to the benchmark New Keynesian loss function, scaled by the population factor $(1+\omega)$. The term $(1+\omega)\frac{\epsilon}{\lambda}\pi_t^2$ captures the welfare cost of price dispersion. The term $(1+\omega)\chi x_t(x_t - 2a_t)$ measures the cost of inefficient fluctuations in hours and output; up to terms independent of policy it can be written as $(1+\omega)\chi(x_t - a_t)^2$, with a_t playing the role of the efficient level of output. The additional term $\tilde{\sigma}x_t^2$ reflects the contribution of consumption risk borne by the more risk-averse household when the aggregate output gap moves.

The second line captures how fluctuations in leverage l_t and uncertainty ξ_t interact with the distribution of consumption between households and entrepreneurs. When leverage rises, entrepreneurs absorb more macroeconomic risk; with incomplete insurance at the firm level, this redistributes consumption risk towards entrepreneurs in bad states and away from the household. Because the social planner values both groups, these reallocations generate extra welfare costs that depend on the level and volatility of leverage, as well as on aggregate uncertainty.

The final term in (1.8) reflects strictly financial-stability costs: higher leverage and higher uncertainty increase the dispersion of entrepreneurial outcomes and thus the expected monitoring costs and deadweight losses associated with the imperfect state verification problem. These costs are convex in both leverage and uncertainty (through κ_{ll} and $\kappa_{l\xi}$), making episodes of high leverage in recessions particularly costly.

It is useful to emphasise why (1.8) contains no first-order (linear) terms in $x_t - a_t$, π_t, l_t or ξ_t. Given the Negishi weights, the deterministic competitive steady state satisfies the planner's first-order conditions, so linear terms around that point cancel in the welfare expansion. For example, the output term can be written as

$$(1+\omega)\chi x_t(x_t - 2a_t) = (1+\omega)\chi\big[(x_t - a_t)^2 - a_t^2\big],$$

so that only the squared output gap $(x_t - a_t)^2$ is welfare relevant, while the linear and constant components are absorbed into the "t.i.p." term. Lemma 1 implies $E_t[\delta_{t+1}] = 0$, so there is no first-order gain from manipulating the predictable component of the macroprudential wedge either. As in standard New Keynesian analysis, this justifies evaluating welfare using the linearised equilibrium conditions together with the quadratic loss in (1.8).

1.5 The special case of log household utility

When the representative household has log utility, $\sigma = 1$, the welfare criterion simplifies considerably. In this case the planner's loss function collapses to the familiar New Keynesian form, up to terms that are independent of policy (see Appendix C, equation (C.5)):

$$\Lambda_{\sigma=1,t} = \frac{1}{2}\left[\frac{\epsilon}{\lambda}\,\pi_t^2 + \chi\, x_t(x_t - 2a_t)\right] + \text{t.i.p.} \tag{1.9}$$

Equivalently, up to policy-irrelevant constants, (1.9) can be written as

$$\Lambda_{\sigma=1,t} = \frac{1}{2}\left[\frac{\epsilon}{\lambda}\,\pi_t^2 + \chi\,(x_t - a_t)^2\right] + \text{t.i.p.},$$

so that only inflation and the welfare-relevant output gap matter for monetary policy.

The intuition is that with log utility the labour supply decision is locally independent of the distribution of consumption between households and entrepreneurs. Individual entrepreneurs still face a labour wedge and monitoring costs because they cannot perfectly insure idiosyncratic risk, but in the aggregate an increase in hours raises both entrepreneurial income and their net worth in a way that leaves the expected monitoring costs unchanged. Under competitive risk sharing over aggregate states, changes in leverage and in the equity risk premium induced by uncertainty shocks do not feed back into the marginal rate of substitution between consumption and leisure for the representative household.

From the perspective of the monetary policymaker, it is therefore efficient for labour supply not to respond to fluctuations in leverage and uncertainty over and above what is implied by the efficient output gap. Macroprudential policy still matters for the allocation of risk and for the welfare of entrepreneurs, but these effects are of second order for the monetary transmission mechanism and do not show up directly in the quadratic loss (1.9). This log-utility benchmark is thus useful for connecting our framework to the canonical New Keynesian model and for isolating the pure macroprudential trade-offs before reintroducing the full feedback from risk aversion and leverage in later sections.

2 Paradox of safety and aggregate risk sharing

This section studies a flexible-price version of the model without macroprudential intervention and shows how higher household risk aversion can make the economy *more* volatile once leverage and imperfect risk sharing are taken into account. When households insure themselves by holding safe or countercyclical assets, entrepreneurs absorb more aggregate risk on their balance sheets, amplifying fluctuations in labour demand, output and consumption. We use a Lucas-style consumption-equivalent calculation as a convenient way to summarise this amplification, while deferring the full welfare analysis to later sections.

Caballero and Farhi (2017) introduce the idea of a safety trap: in a severe liquidity trap, attempts by households to eliminate risk end up exacerbating the shortage of safe assets and amplifying real volatility. Our model exhibits a closely related mechanism, but in a fully micro-founded New Keynesian environment with endogenous leverage and macro risk sharing. Risk-averse households try to protect themselves from business-cycle risk by holding safe or countercyclical financial claims. Entrepreneurs take the other side of these trades and so absorb the aggregate risk. As a result, macro risk becomes concentrated on entrepreneurial balance sheets, generating large procyclical movements in entrepreneurial net worth and a financial accelerator acting through leverage and the labour wedge.

The key twist is that higher household risk aversion can make the economy *more* volatile. When households become more risk averse, their demand for safe assets rises. In equilibrium, this pushes more aggregate risk onto entrepreneurs, increasing the volatility of leverage. A recessionary technology shock then has a doubly adverse effect. First, it lowers output directly. Second, because entrepreneurs are bearing more risk, it delivers a larger proportional hit to their net worth, producing a bigger jump in leverage. Higher leverage raises marginal production costs via (1.4) and depresses aggregate demand via (1.1). For typical monetary policy specifications, or under flexible prices, this feedback loop further amplifies the fall in output.

In benchmark New Keynesian and real business cycle models, by contrast, higher risk aversion dampens fluctuations: households smooth consumption more aggressively, and the wealth effect on labour supply stabilises hours and output. In our environment those familiar stabilising channels are still present, but they are offset—and for plausible calibrations more than offset—by the effect of risk aversion on the concentration of risk in the firm sector and the induced response of leverage and the labour wedge.

Proposition 1 (The paradox of safety) *For a range of parameter values consistent with our baseline calibration, an increase in household risk aversion σ raises the unconditional volatility of output and consumption in competitive equilibrium. In particular,*

$$\frac{\partial}{\partial\sigma}\mathrm{Var}(x_t) > 0 \quad \textit{and} \quad \frac{\partial}{\partial\sigma}\mathrm{Var}(c_t) > 0$$

in a neighbourhood of the baseline σ.

A flexible-price illustration

The mechanism is easiest to see in the flexible-price version of the model with technology shocks only and no macroprudential wedge (so $\xi_t = \delta_t = 0$). Appendix E shows that aggregate supply can then be written as

$$(\chi + \tilde{\sigma})x_t = \chi a_t - \mu_l l_t, \tag{2.1}$$

where $\tilde{\sigma} := \sigma - 1$, χ is the composite curvature parameter from labour supply and production, and

$$\mu_l := \sigma\omega(1-\psi) + \tau_l > 0$$

is the composite elasticity of marginal costs with respect to leverage, capturing both wealth effects and the labour-wedge channel. The flexible-price leverage curve reduces to

$$l_t = \phi l_{t-1} - (1-\phi)\frac{\tilde{\sigma}}{\psi}(x_t - x_{t-1}), \tag{2.2}$$

with $\phi \in (0,1)$ the persistence of leverage.

A negative productivity shock a_t reduces output x_t through (2.1). Under competitive macro risk sharing, this also depresses entrepreneurs' net worth and so raises leverage l_t via (2.2). Higher leverage pushes up marginal costs and reduces labour demand, feeding back into output. Combining (2.1) and (2.2), and using the technology process

$$a_t = \rho_a a_{t-1} + \epsilon_t^a, \qquad E_t[\epsilon_{t+1}^a] = 0, \quad \mathrm{Var}(\epsilon_t^a) = \sigma_a^2,$$

one can solve for a reduced-form law of motion for output of the form (see the on-line Appendix)

$$x_t = \nu_x(\sigma)\, x_{t-1} + \nu_a(\sigma)\big(a_t - \phi a_{t-1}\big), \tag{2.3}$$

where the coefficients are

$$\nu_a(\sigma) = \frac{\chi\psi}{\psi(\chi + \tilde{\sigma}) + \mu_l\tilde{\sigma}(\phi - 1)}, \tag{2.4}$$

$$\nu_x(\sigma) = \frac{\chi\phi\psi + \mu_l\phi\tilde{\sigma} - \mu_l\tilde{\sigma}}{\psi(\chi + \tilde{\sigma}) + \mu_l\tilde{\sigma}(\phi - 1)}. \tag{2.5}$$

Household consumption moves proportionally with output under flexible prices:

$$c_t = \kappa_x x_t, \qquad \kappa_x > 0,$$

so the volatility of consumption is proportional to the volatility of output.

Equations (2.4)–(2.5) highlight how risk aversion enters. In a canonical flexible-price model without leverage effects one has

$$\nu_a^{\mathrm{NK}}(\sigma) = \frac{\chi}{\chi + \tilde{\sigma}}, \qquad \nu_x^{\mathrm{NK}}(\sigma) = 0,$$

so the on-impact response of output to technology shocks is *decreasing* in risk aversion. In our model, leverage and the labour wedge appear through μ_l and ϕ. For typical parameterisations with $\tau_l > 1$ and ϕ moderately high, both μ_l and ϕ are increasing in σ, and the denominator in (2.4) falls as σ rises. As a result, $\nu_a(\sigma)$ and $\nu_x(\sigma)$ both increase with risk aversion: technology shocks induce larger and more persistent movements in output the more risk averse households become. Appendix E.4 and Figure 6 provide a numerical example, showing that output volatility is increasing in σ over a wide range of values.

A Lucas-style welfare interpretation

For welfare, it is natural to relate the paradox of safety to Lucas's (1987) famous calculation of the consumption-equivalent cost of business cycles. In his experiment, a CRRA representative agent with relative risk aversion σ faces exogenous log consumption,

$$\log c_t = \log \bar{c} + \hat{c}_t, \qquad \hat{c}_t \sim \mathcal{N}(0, s^2),$$

and the question is: by what fraction λ would we have to raise the level of *deterministic* consumption in order to make the agent indifferent between the stochastic path and a flat path with the same mean? For small s^2, Lucas shows that

$$\lambda \;\approx\; \frac{\sigma}{2}\, s^2.$$

In our environment, consumption is endogenous and inherits its volatility from technology through the leverage channel. The on-line Appendix shows that, under flexible prices and technology shocks only, we can write log consumption in reduced form as an ARMA process driven by ϵ_t^a, and the stationary variance of log consumption takes the form

$$\mathrm{Var}(\hat{c}_t) = \Xi(\sigma)\, \sigma_a^2,$$

where $\Xi(\sigma)$ is an explicit function of $\nu_a(\sigma)$, $\nu_x(\sigma)$ and hence of the leverage and labour-wedge parameters. A second-order Lucas expansion then delivers a consumption-equivalent welfare cost

$$\lambda(\sigma) \;\approx\; \frac{\sigma}{2}\, \Xi(\sigma)\, \sigma_a^2.$$

Two observations follow immediately. First, when leverage is shut down ($\mu_l = 0$ and ϕ independent of σ), $\nu_a(\sigma)$ collapses to the familiar flexible-price New Keynesian coefficient $\chi/(\chi + \tilde{\sigma})$, $\nu_x(\sigma) = 0$, and $\Xi(\sigma)$ reduces to Lucas's s^2 (up to the constant factor κ_x^2). We literally recover his formula. Second, with leverage active, μ_l and ϕ rise with risk aversion and $\Xi(\sigma)$ becomes an *increasing* function of σ. Higher risk aversion now does two things at once: it raises the marginal disutility of risk (the prefactor $\sigma/2$) and it increases the variance of equilibrium consumption via $\Xi(\sigma)$.

This is the sense in which our paradox of safety is a Lucas-style welfare result. The form of the welfare cost is unchanged–it is still $\lambda \approx (\sigma/2)\times$variance of log consumption–but the mapping from technology shocks to consumption variance is itself distorted by leverage and risk sharing. As households try to make themselves safer by loading risk onto entrepreneurs, they raise the equilibrium volatility of their own consumption and increase the Lucas cost of business cycles.

3 Optimal monetary and macroprudential policy: log-utility benchmark

We now characterise the joint roles of monetary and macroprudential policy in a benchmark case where both households and entrepreneurs have log utility. This shuts down the paradox of safety while preserving the leverage and risk-sharing channels. In this environment, financial stress—summarised by the equity premium and the associated factor-market wedge—acts like a cost-push disturbance in the Phillips curve and a demand disturbance in the IS curve, and we show that optimal monetary policy still responds to financial conditions even when macroprudential tools are used in a constrained-efficient way.

Household risk aversion is central to the financial amplification mechanism in our model, and restricting households to log utility means that fluctuations in leverage and the equity risk premium are solely the result of uncertainty shocks. We find this log utility benchmark to be a useful starting point for our analysis, before returning to the general model with greater household risk aversion in later sections.

In this section we present three optimal policy results. First, we characterise optimal macroprudential policy in a flexible price benchmark economy. Second, we characterise optimal monetary and macroprudential policy under sticky prices. Optimal monetary policy stimulates the economy following uncertainty shocks, which complements the optimal macroprudential response. Third, we characterise optimal macroprudential policy under an interest rate rule. In this case, we focus on technology shocks, where the interest rate rule does not optimally manage aggregate demand, creating a role for macroprudential policy that differs from the earlier regimes.

Assumption 1 *All of the results in this section only rely on the assumption that household utility is logarithmic, $\sigma = 1$.*

Assumption 1 is strong. It improves tractability at the cost of removing an important feedback mechanism from the model. Under log utility, the feedback from output to leverage (and financial stability) is broken. Net worth moves one-for-one with aggregate output. Broadly speaking, this assumption means that the monetary policy authority can treat leverage as exogenous. This makes the model very tractable and helps us identify costs and benefits of policy interventions. The financial sector does not amplify technology shocks under log utility, and financial distress reflects the net impact of uncertainty shocks and macroprudential interventions. Macroprudential policy may optimally choose to generate a link between financial stability and technology shocks in order to dampen the output response to technology shocks. In Section 4 we analyse optimal policy responses to technology shocks with greater household risk aversion and financial amplification.

The full derivations for this section are available in the main Appendix Section F.

3.1 The flexible price benchmark

We start with a flexible price benchmark before re-introducing nominal rigidities and monetary policy. Section E in the Appendix derives the flexible price aggregate demand and supply equilibrium relationship,

$$\chi(x_t - a_t) = -\mu_l l_t - \mu_\xi \xi_t. \tag{3.1}$$

Real output is increasing with technology, but decreases with leverage and uncertainty shocks. Both leverage and uncertainty increase risk borne by entrepreneurs, reducing labour demand. In addition, an increase in leverage reflects an increase in household wealth, generating a negative wealth effect on labour supply. Holding all else equal, an increase in uncertainty increases the entrepreneurs' share of consumption, generating a positive wealth effect on labour supply which dampens the effect of uncertainty shocks on output.

By Lemma 1, the prudential policymaker is constrained to policies that satisfy $\mathbb{E}_t[\delta_{t+1}] = 0$.[4] It is convenient to incorporate this constraint into the leverage curve (1.3) in order to arrive at the constraint specified by (3.2), expressed in terms of the feasible paths of leverage that the prudential policymaker can implement:

$$\mathbb{E}_t[\Delta l_{t+1}] = -(1-\phi)(l_t + \xi_t) + \omega(1-\phi)\mathbb{E}_t[\Delta \xi_{t+1}]. \tag{3.2}$$

The flexible price macroprudential policymaker's problem is described by Programme 1.

Programme 1

$$\min_{x,l} \mathbb{E} \sum_{t=0}^{\infty} \beta^t \frac{1}{2} \left[\begin{array}{l} (1+\omega)\chi\left(x_t^2 - 2x_t a_t\right) + \omega\left(\kappa_{ll} + (1+\sigma\omega)(1-\psi)^2\right) l_t^2 \\ +2\omega\left(\kappa_{l\xi} - (1+\sigma\omega)(1-\psi)\psi\right) l_t \xi_t \end{array} \right]$$

Subject to (3.1), (3.2).

Solving Programme 1 yields the following optimal macroprudential policy wedge:

$$\delta_t = \left(\frac{\omega\hat{\kappa}_{l\xi} + (1+\omega)\chi^{-1}\mu_l\mu_\xi}{\omega\hat{\kappa}_{ll} + (1+\omega)\chi^{-1}\mu_l^2} \left(\frac{\phi' - \phi}{\phi' - \rho_\xi} \right) - \frac{1 - \omega(\phi' - 1)}{\phi' - \rho_\xi} (1-\phi) \right) \epsilon_{\xi,t}, \tag{3.3}$$

where ϕ'^{-1}, the explosive eigenvalue associated with the shadow cost of the leverage constraint.

First, recall that the uncertainty process is

$$\xi_{t+1} = \rho_\xi \xi_t + \epsilon_{\xi,t+1}, \qquad \epsilon_{\xi,t+1} = \xi_{t+1} - \mathbb{E}_t[\xi_{t+1}],$$

so that $\epsilon_{\xi,t+1}$ is the innovation (one–step–ahead forecast error) in uncertainty. By construction $\epsilon_{\xi,t+1}$ is $\mathcal{F}_{t+1}$–measurable and orthogonal to $\mathcal{F}_t$, and our measurability

[4]Recall that this is due to macroprudential policies affecting the allocation of risk and the responses of wealth to shocks, and cannot redistribute wealth between agents in expectation.

restriction on macroprudential policy implies that the implementable wedge must be a martingale difference, loading only on this unanticipated component. In other words, prudential policy can change how leverage responds to new, unexpected uncertainty shocks, but cannot engineer predictable transfers across agents or across dates; any systematic component would be arbitraged away in equilibrium.

Next, note that the ratio

$$\frac{\omega\hat{\kappa}_{l\xi} + (1+\omega)\chi^{-1}\mu_l\mu_\xi}{\omega\hat{\kappa}_{ll} + (1+\omega)\chi^{-1}\mu_l^2}$$

is the current period marginal rate of transformation between the social costs of uncertainty and the social costs of leverage. Without loss of generality, $\hat{\kappa}_{l\xi}$ is the individual cost of increased entrepreneurial risk bearing resulting from greater covariance between leverage and uncertainty, and is weighted by the entrepreneurial Negishi weight ω. The product $\chi^{-1}\mu_l\mu_\xi$ captures the cost of reduced hours resulting from the labour demand and supply effects of leverage and uncertainty, which are particularly high when the labour margin is more elastic (when χ is small). The resulting costs are borne by all and are therefore Negishi weighted $(1+\omega)$. In sum, the numerator captures the extent to which a change in leverage can offset the marginal social costs of uncertainty, and the denominator captures the social costs of the resulting volatility of leverage. These relative costs are weighted by the relative persistences of uncertainty and leverage. If the persistence of leverage ϕ is high relative to the persistence of uncertainty ρ_ξ, then the policymaker will moderate their prudential response to uncertainty shocks.

In the competitive equilibrium, uncertainty shocks increase current period leverage but they reduce leverage over longer time horizons. When uncertainty is high, the return to inside wealth is also high, and entrepreneurs' inside wealth grows quickly. As leverage is persistent, macroprudential policy has a enduring effect on the path of leverage, and can exacerbate the medium term decrease in leverage in response to a contractionary uncertainty shock. This persistence may not be desirable. Finally, observe that the second term,

$$-\frac{1-\omega(\phi'-1)}{\phi'-\rho_\xi}(1-\phi),$$

reflects the persistent effect of current period uncertainty on future leverage, and dampens the optimal macroprudential response to uncertainty shocks.

Optimal macroprudential policy does not respond to technology shocks in this economy. Under log utility, technology shocks do not generate fluctuations in leverage. The competitive allocation appropriately adjusts hours worked in response to changes in technology. We show in Section 3.3 that deviations from optimal aggregate demand management can generate a motivation for macroprudential policy even in the absence of feedback from output to leverage, and we show in Section 4 that when the representative household is more risk averse, financial amplification of technology shocks generates fluctuations in leverage and in turn motivates macroprudential policy.

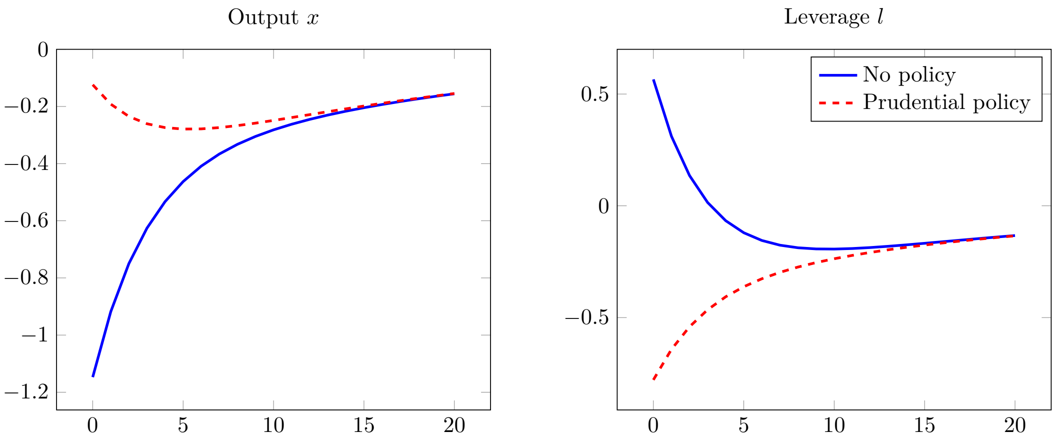

Figure 1: Responses to a recessionary uncertainty shock.

Figure 1 presents responses to a recessionary uncertainty shock, with and without macroprudential policy. In the absence of policy, entrepreneurial net wealth decreases sharply in response to the uncertainty shock, with leverage increasing as a result. The combination of high leverage and high uncertainty decreases labour demand, and output decreases in response. Under the optimal prudential policy, entrepreneurial net wealth is protected, and leverage decreases with output. Falling leverage helps to dampen the response of labour demand to the uncertainty shock, and as a result, the output response is dampened.

3.2 Optimal monetary and prudential policy with nominal rigidities

In this section we reintroduce nominal rigidities and solve for jointly optimal monetary and prudential policy under commitment. We separate the problem into two parts. Under log utility, the effect of the monetary policymaker's action on leverage is mediated through the optimal policy of the prudential policymaker. So, we first solve for the paths of output and inflation as functions of leverage, uncertainty and technology shocks—we interpret this as monetary policy—then we solve for the optimal path of leverage—we interpret this as the prudential policy.

The combined policymaker solves the following programme:

Programme 2 *Joint optimal monetary and prudential policy under log utility.*

$$\min_{\pi,x,l} \mathbb{E} \sum_{t=0}^{\infty} \beta^t \frac{1}{2} \left((1+\omega) \left(\frac{\epsilon}{\lambda} \pi_t^2 + \chi \left(x_t^2 - 2x_t a_t \right) \right) + \omega \hat{\kappa}_{ll} l_t^2 + 2\omega \hat{\kappa}_{l\xi} l_t \xi_t \right)$$

Subject to (3.2), and

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \lambda\chi x_t - \lambda\chi a_t + \lambda\mu_l l_t + \lambda\mu_\xi \xi_t.$$

The divine coincidence holds for technology shocks under log utility, so we focus our analysis on uncertainty shocks.

Leverage and uncertainty enter the Phillips curve in a similar way to traditional New Keynesian cost-push shocks. Given the absence of feedback from monetary policy to leverage under log utility, optimal monetary policy faces similar trade-offs to monetary policy under cost-push shocks. Optimal inflation resembles a price-level targeting rule. Inflation increases in response to leverage and uncertainty, but eventually turns negative in order to restore the original price level (which is normalised to zero).

$$p_t = \phi_1 p_{t-1} + \frac{\beta^{-1}\lambda}{\phi_2 - \phi} \left(\mu_l l_t + \mu_\xi \left(1-\gamma\right) \xi_t \right), \tag{3.4}$$

where ϕ_1, ϕ_2 are the stable and explosive eigenvalues associated with optimal aggregate demand management familiar to New Keynesian models, and γ reflects the policymakers internalisation of expected effect of current uncertainty on future leverage.[5]

By allowing prices to increase on impact following recessionary uncertainty shocks, the monetary policy authority bears a welfare cost from inflation but generates an increase in welfare by smoothing the path of output, consumption, and hours worked. This countercyclical monetary policy has no impact on leverage and risk bearing, with firms' net wealth increasing one for one with output to ensure that leverage remains invariant to monetary stimulus.

Optimal prudential policy is countercyclical, reducing leverage whenever the expected path of prices or the risk-bearing cost of uncertainty rises:

$$\delta_t = \frac{(1+\omega)\epsilon\mu_l}{\omega\hat{\kappa}_{ll}}(\phi' - \phi)\sum_{j=0}^{\infty}(\beta\phi)^{j+1}(\mathbb{E}_t[p_{t+j}] - \mathbb{E}_{t-1}[p_{t+j}])$$
$$+ \left(\frac{\hat{\kappa}_{l\xi}}{\hat{\kappa}_{ll}}\left(\frac{\phi' - \phi}{\phi' - \rho_\xi}\right) - \frac{1 - \omega(\phi' - 1)}{\phi' - \rho_\xi}(1-\phi)\right)\epsilon_{\xi t}$$

Under commitment, optimal monetary policy trades off the inflation it must tolerate against the output gaps it can smooth in response to uncertainty shocks. The prudential authority evaluates its own actions against this benchmark: cutting leverage by one unit lowers marginal costs by μ_l on impact, with the effect on future prices decaying at rate ϕ and discounted at β, while the welfare cost of any extra inflation scales with the elasticity of substitution ϵ and the aggregate Negishi weight $1 + \omega$.

Equation 3.5 expresses the optimal prudential response purely as a function of the uncertainty shock. When $\epsilon \to \infty$, even small price distortions are very costly, so optimal monetary policy keeps inflation near zero and the prudential rule collapses to the flexible-price benchmark in (3.3). When $\epsilon \to 0$, countercyclical monetary policy can almost fully neutralise the marginal-cost channel of uncertainty shocks,

[5]

$$\phi_1 = \frac{(1+\beta+\lambda\chi\epsilon) - \sqrt{(1+\beta+\lambda\chi\epsilon)^2 - 4\beta}}{2\beta}, \qquad \phi_2 = \frac{1}{\beta\phi_1},$$
$$\gamma = \frac{\phi - \rho_\xi + \frac{\mu_l}{\mu_\xi}(1+\omega(1-\rho_\xi))(1-\phi)}{\phi_2 - \rho_\xi}, \qquad \lim_{\epsilon\to\infty}\gamma = 0.$$

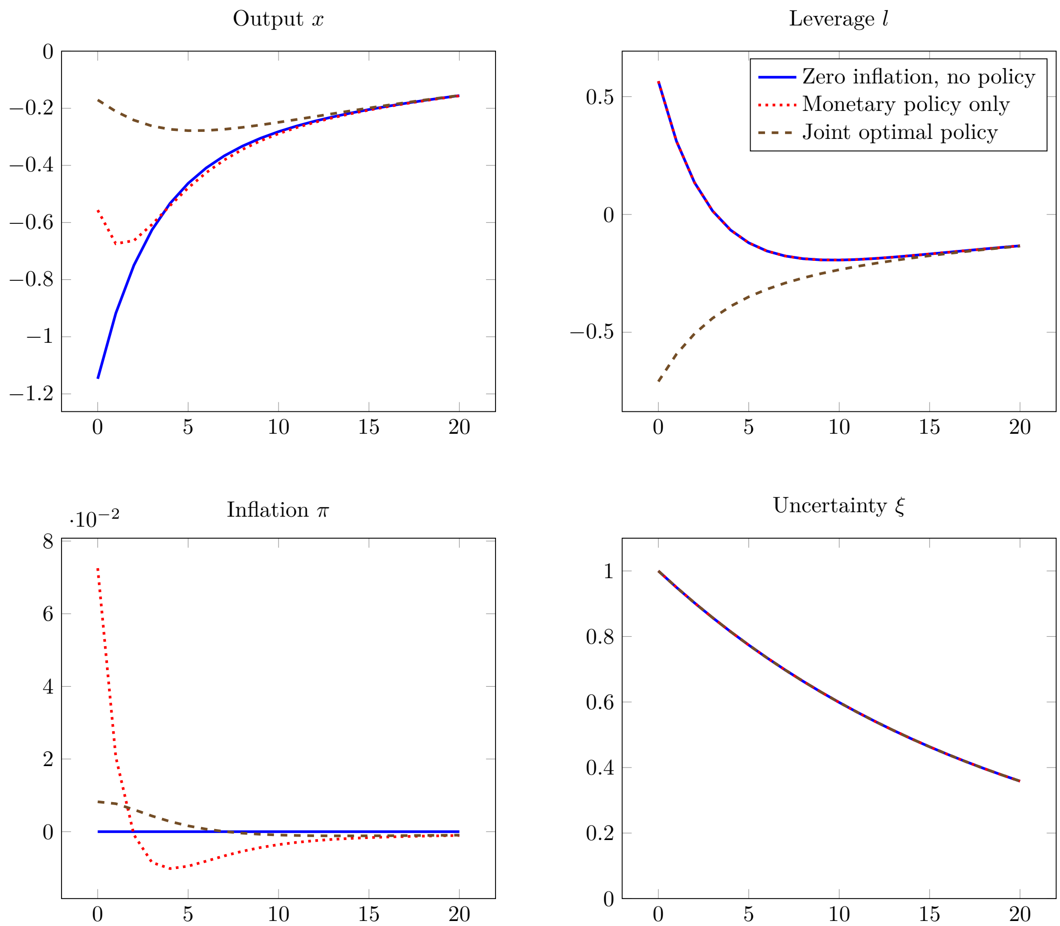

Figure 2: Monetary and prudential responses to a recessionary uncertainty shock.

and prudential policy optimally focuses on the remaining risk-bearing and distributional distortions from uncertainty and leverage.

$$\delta_t = \left(\frac{\chi\omega\hat{\kappa}_{l\xi} + (1+\omega)\mu_l\mu_\xi\varsigma(1-\gamma)}{\chi\omega\hat{\kappa}_{ll} + (1+\omega)\mu_l^2\varsigma} \left(\frac{\phi' - \phi}{\phi' - \rho_\xi} \right) - \frac{1 - \omega(\phi' - 1)}{\phi' - \rho_\xi}(1 - \phi) \right) \epsilon_{\xi t} \tag{3.5}$$

where

$$\varsigma = \frac{\lambda\chi\epsilon}{\beta} \frac{\phi'}{(\phi_2 - \phi)(\phi' - \phi_1)}, \qquad \lim_{\epsilon \to 0} \varsigma = 0, \quad \lim_{\epsilon \to \infty} \varsigma = 1.$$

Figure 2 compares three allocations after a recessionary uncertainty shock: flexible prices with no policy, optimal monetary policy alone, and joint optimal monetary–prudential policy. With log utility, optimal monetary policy lets inflation rise on impact, softening the output fall; because there is no feedback from monetary policy to leverage in this case, this is a pure inflation–output trade-off. When prudential policy is also used optimally, the required prudential response is slightly smaller than in the flexible-price benchmark (see Figure 1), and the output decline is smaller than under prudential policy alone.

Without prudential tools, the optimal monetary response features a sharp rise in inflation at the onset of the shock, followed by a period of below-target inflation to bring the price level back to path. With prudential policy in place, the initial inflation spike is muted and the subsequent undershooting is much smaller. The central bank still restores the original price level, but does so with substantially smaller deviations of inflation from zero in both directions.

3.3 Optimal prudential policy with an interest rate rule

In both the flexible-price benchmark and the optimal-monetary-policy case with log utility, there is no reason for the prudential authority to react to technology shocks. Well-designed monetary policy can handle the demand response on its own, and technology shocks do not threaten financial stability in that environment.

Macroprudential policy becomes relevant for technology shocks only when the aggregate-demand response is off the Ramsey path—for example, under a fixed exchange rate or monetary union, when the central bank follows a simple Taylor rule, or when it optimises under discretion.[6] In what follows we work with an interest-rate rule, but present our results in terms of the implied output and inflation elasticities to technology shocks so that the arguments apply more broadly.

The macroprudential trade-offs for uncertainty shocks are analogous to those already discussed. To avoid repetition, in this section we shut down uncertainty shocks and focus on technology shocks alone. We assume the policy interest rate follows the simple rule

$$i_t = \phi_\pi \pi_t, \qquad \text{where } \phi_\pi > 1.$$

[6] Chen, Kirsanova, and Leith (2017) show that US monetary policy is well described by an optimising but discretionary central bank.

We then solve the system

$$x_t = \mathbb{E}[x_{t+1}] - (\phi_\pi \pi_t - \mathbb{E}_t[\pi_{t+1}]) - \sigma\omega(1-\psi)(1-\phi)l_t \tag{IS}$$

$$\pi_t = \beta\mathbb{E}_t[\pi_{t+1}] + \lambda\chi(x_t - a_t) + \lambda\mu_l l_t \tag{PC}$$

to arrive at a general solution with the form

$$x_t = \eta_{xa} a_t + \eta_{xl} l_t, \tag{3.6}$$

$$\pi_t = \eta_{\pi a} a_t + \eta_{\pi l} l_t, \tag{3.7}$$

where

$$\begin{aligned} \eta_{\pi a} &= -\frac{\lambda\chi}{(1-\beta\rho_a) + \frac{\phi_\pi - \rho_a}{1-\rho_a}\lambda\chi}, \qquad & \eta_{xa} &= -\left(\frac{\phi_\pi - \rho_a}{1-\rho_a}\right)\eta_{\pi a}, \\ \eta_{\pi l} &= \frac{\lambda\mu_l - \sigma\omega(1-\psi)\lambda\chi}{(1-\beta\phi) + \frac{\phi_\pi - \phi}{1-\phi}\lambda\chi}, \qquad & \eta_{xl} &= -\left(\frac{\phi_\pi - \phi}{1-\phi}\right)\eta_{\pi l} - \sigma\omega(1-\psi). \end{aligned} \tag{3.8}$$

Summarising the above solution, both technology and leverage generate fluctuations in marginal costs, with elasticities χ and μ_l respectively. Leverage however also reduces aggregate demand, with elasticity $\sigma\omega(1-\psi)(1-\phi)$. The resulting decrease in output, represented by the term $-\sigma\omega(1-\psi)$ in the expression for η_{xl}, dampens or potentially reverses the response of marginal costs and inflation to fluctuations in leverage ($-\sigma\omega(1-\psi)\lambda\chi$). The responses of inflation to technology and leverage fluctuations are decreasing in ϕ_π, but are increasing in the persistence of technology and leverage fluctuations respectively ρ_a and ϕ.

We impose the solution (3.6, 3.7) as a constraint on the macroprudential policymaker. The macroprudential policymaker then solves Programme 3.

Programme 3

$$\min_{\pi,x,l} \mathbb{E}\sum_{t=0}^{\infty}\beta^t\frac{1}{2}\left((1+\omega)\left(\frac{\epsilon}{\lambda}\pi_t^2 + \chi\left(x_t^2 - 2x_t a_t\right)\right) + \omega\hat{\kappa}_{ll} l_t^2\right)$$

subject to (3.7), (3.6), and

$$\mathbb{E}_t[\Delta l_{t+1}] = -(1-\phi) l_t.$$

The optimal macroprudential wedge ω can be expressed as follows:

$$\delta_{t+1} = \left(\frac{\phi' - \phi}{\phi' - \rho_a}\right) \frac{(1+\omega)\left(\frac{\epsilon}{\lambda}\eta_{\pi l}\eta_{\pi a} + \chi\eta_{xl}(\eta_{xa} - 1)\right)}{\omega\hat{\kappa}_{ll} + (1+\omega)\left(\frac{\epsilon}{\lambda}\eta_{\pi l}^2 + \chi\eta_{xl}^2\right)} \epsilon_{at+1}.$$

The sign of the prudential policy response to technology shocks is given by the sign of the following expression,

$$\frac{\epsilon}{\lambda} \underset{-/+}{\eta_{\pi l}} \ \underset{-}{\eta_{\pi a}} + \chi \ \underset{-}{\eta_{xl}} \ (\underset{-}{\eta_{xa}} - 1).$$

Moving from right to left, under the interest rate rule a positive technology shock raises output ($\eta_{xa} > 0$) but not enough to close the efficient output gap ($\eta_{xa} - 1 < 0$). Higher leverage reduces output ($\eta_{xl} < 0$), and this is always undesirable from a demand-management point of view (there is no offsetting term in the welfare function that rewards output–leverage comovement). Hence the product $\eta_{xl}(\eta_{xa} - 1)$ is positive: by raising ω and making leverage move more countercyclically with output, the prudential authority can lower the welfare cost of an under-reacting output gap—at the price of intentionally introducing financial amplification where none existed under the raw interest-rate rule.

Once inflation enters, the trade-off changes. A positive technology shock is disinflationary $\eta_{\pi a} < 0$, reflecting the fact that output does not rise enough to prevent marginal costs from falling. The effect of leverage on inflation, $\eta_{\pi l}$, can be positive or negative. If the demand channel dominates, higher leverage lowers inflation $\eta_{\pi l} < 0$, so prudentially induced financial amplification helps offset the disinflationary impact of technology shocks. If instead leverage mainly raises marginal costs so that $\eta_{\pi l} < 0$, amplification worsens inflation volatility; in that case, especially when the welfare cost of inflation is high, the prudential authority should lean against technology shocks, engineering a more countercyclical path for leverage and reducing the elasticity of net worth with respect to output below one.

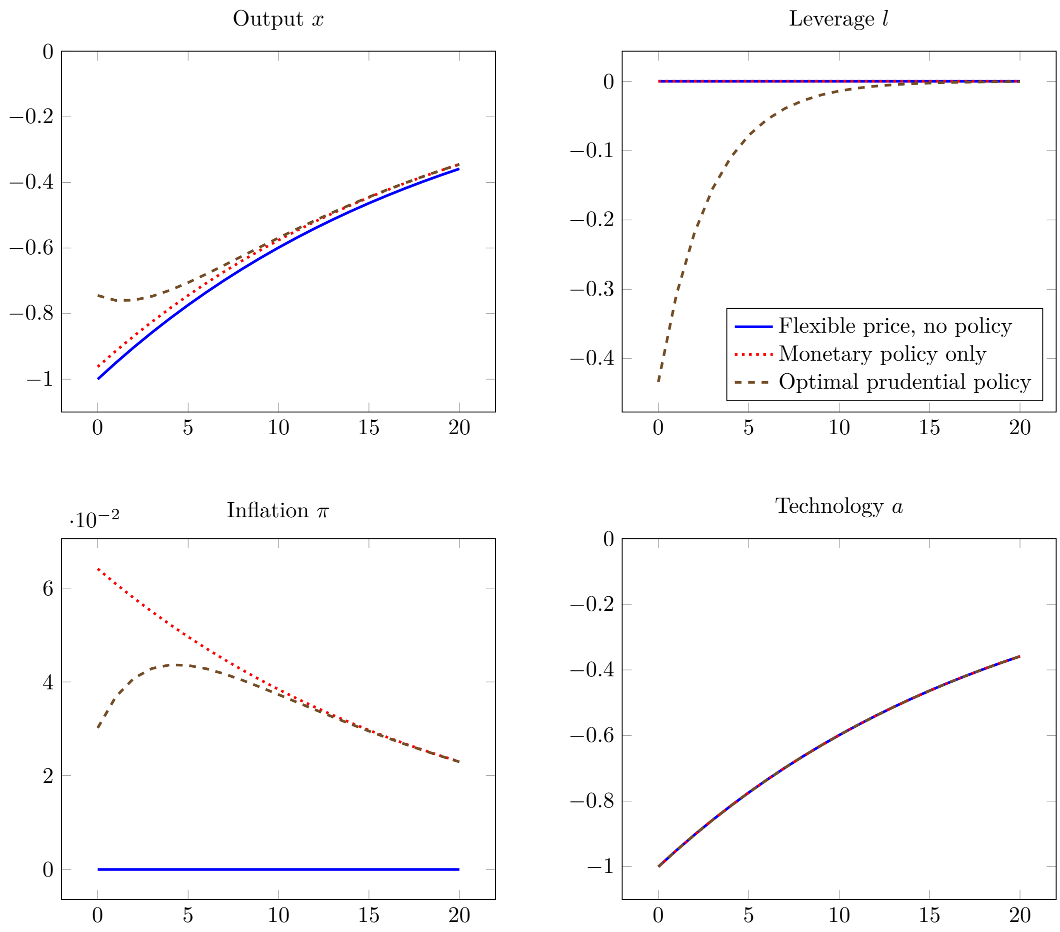

Figure 3: Prudential response to a recessionary technology shock ($\phi_\pi = 1.7$).

Figure 3 compares responses to a recessionary technology shock with and without prudential policy under a simple interest-rate rule. With log utility, the divine coincidence holds under optimal monetary policy: inflation is kept at zero and constrained-efficient prudential policy does not react to technology shocks. To create a role for macroprudential policy, we instead impose a Taylor-type rule, $i_t = \phi_\pi \pi_t$ with $\phi_\pi = 1.7$, which generates a small positive output gap after a negative technology shock. Given this suboptimal rule, the optimal prudential response is countercyclical: leverage falls in recessions and rises in expansions, further damping the output response relative to the flexible-price benchmark. Lower leverage in downturns reduces marginal production costs—both by cutting expected monitoring costs and by shifting part of the wealth loss onto households, whose higher labour supply offsets some of the demand shortfall. This moderates the inflation response and reduces the welfare cost of deviating from fully optimal monetary policy.

Under this interest rate rule, the optimal prudential response is countercyclical in our calibration: leverage falls in recessions and rises in expansions, further damping the output response relative to the flexible-price (optimal-policy) path. Lower leverage in response to the recessionary technology shock reduces marginal production costs—both by cutting expected monitoring costs and by shifting part of the wealth loss onto households, whose higher labour supply offsets some of the demand shortfall. This moderates the inflation response and reduces the welfare cost of deviating from fully optimal monetary policy. More generally, whether prudential policy dampens or amplifies technology shocks depends on how strongly leverage feeds into inflation and how costly inflation variability is, but the figure illustrates the dampening case that is most relevant for our calibration.

4 Risk aversion, the paradox of safety, and policy trade-offs

This section relaxes log utility for the worker–household and reintroduces the paradox of safety. Higher household risk aversion strengthens the demand for safe or countercyclical assets, concentrates aggregate risk on entrepreneurial balance sheets, and makes the leverage and factor wedges more sensitive to shocks. We use this richer risk environment to study how monetary and macroprudential policy should trade off stabilising inflation, output and leverage when the central bank maintains zero anticipated inflation.

The basic amplification logic is straightforward. When a negative technology shock hits, entrepreneurs' net worth falls by more than proportionately because they are holding the risky side of aggregate risk. Leverage rises, the financial friction tightens, and marginal costs increase. This shows up in the Phillips curve as an extra term linking marginal costs to leverage, and in the leverage equation as a feedback from current output to future leverage. Putting the two together, a given fall in output leads to a more than proportionate increase in marginal costs: the financial system amplifies technology shocks.

To focus on this amplification channel and on the relative roles of monetary and prudential policy, we impose the following simplifications.

Assumption 2 *For the analysis in this section:*

(i) Uncertainty shocks are switched off. Only technology shocks drive the economy.

(ii) Monetary policy is restricted to deliver zero expected future inflation. That is, policy is parameterised so that

$$E_t[\pi_{t+1}] = 0,$$

and any departure of inflation from zero in response to current technology shocks must be unwound in expectation.

Assumption 2(i) keeps the shock structure close to the benchmark New Keynesian case studied earlier. Assumption 2(ii) is a deliberate restriction on the monetary authority. It converts the analysis into a question about whether small departures from strict inflation stabilisation can improve outcomes by dampening financial amplification, rather than a full characterisation of unconstrained optimal policy.

Under these assumptions, the flexible-price Phillips curve and the leverage equation can be written as

$$(\tilde{\sigma} + \chi)x_t = -\mu_l \, l_t + \chi a_t + \gamma \varepsilon_{a,t}, \tag{4.1}$$

$$l_t = \varphi l_{t-1} - (1 - \varphi)\frac{\tilde{\sigma}}{\psi}\Delta x_t - \delta \varepsilon_{a,t}, \tag{4.2}$$

where x_t is output, l_t is leverage, a_t is technology, $\varepsilon_{a,t}$ is the innovation to technology, $\tilde{\sigma} = \sigma - 1$ captures household risk aversion, and μ_l, φ, ψ are the leverage and wedge parameters derived from the financial friction. The coefficients in (4.1) and (4.2) depend on risk aversion σ: higher σ makes leverage more responsive to shocks and strengthens the effect of leverage on marginal costs.

Equations (4.1)–(4.2) show that technology shocks affect marginal costs both directly, through a_t, and indirectly, through their impact on leverage. Monetary policy can respond by allowing inflation to move (via γ), while macroprudential policy can respond by altering how leverage reacts to the shock (via δ). Solving this simple two-equation system yields closed-form expressions for the impulse responses of x_t and l_t to technology shocks and to each policy instrument (see Appendix C).

Two results follow.

First, both monetary and macroprudential policy should be used. For generic parameterisations, no single instrument can simultaneously stabilise inflation, output and leverage. Monetary policy is relatively effective at moving output on impact but has more transient effects on leverage, because any inflationary response must be reversed to preserve zero expected inflation. Macroprudential policy has a larger and more persistent effect on leverage, but smaller effects on output. Using both instruments allows the policymaker to better balance the marginal welfare costs of fluctuations in inflation, output and leverage.

Second, the relative advantage of each instrument depends on the persistence of technology shocks. When shocks are highly persistent, the leverage response to macroprudential policy has a similar time profile to the leverage response to the shock itself, while the leverage response to monetary policy is more short-lived. Prudential tools are then better suited to containing financial amplification. When technology shocks are short-lived, the leverage effects of monetary policy more closely match the leverage response induced by the shocks, while prudential policy tends to induce overshooting in later periods. In that case, monetary policy is better suited to dampening leverage volatility. These differences in impact and persistence are illustrated in Figure 4, where we choose γ and δ so that both tools deliver the same period-1 response of leverage.

Both instruments share a common cost: using them countercyclically pushes hours above the level where the marginal contribution of work to output compensates for the sum of disutility of labour and monitoring costs. At the same time, both reduce the feedback from output to leverage, lowering the volatility of monitoring costs and increasing welfare. Monetary stabilisation carries the extra cost of inflation, which prudential stabilisation does not. Nevertheless, monetary policy remains part of the optimal mix because prudential policy alone leaves excess volatility in leverage and output that can be reduced at a second-order welfare cost of temporary inflation. The higher is the elasticity of substitution between goods the

Figure 4: Responses to a unit recessionary technology shock.

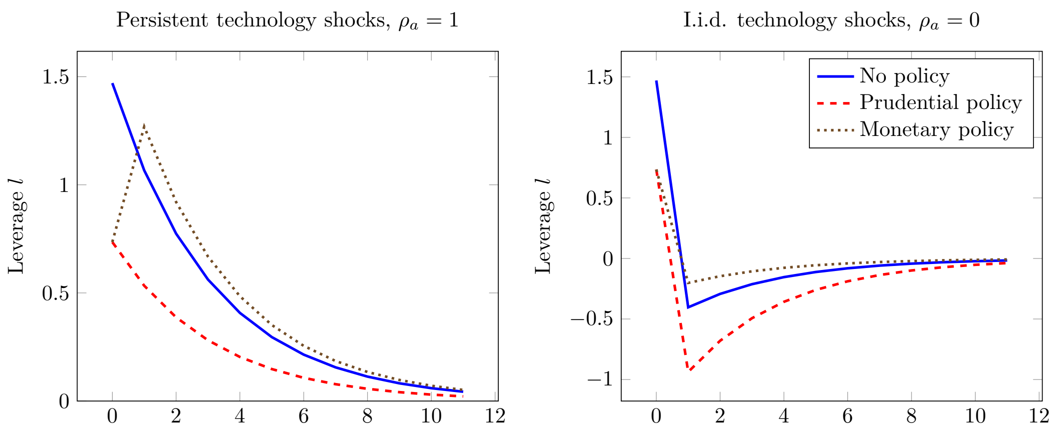

Policy parameters γ, δ are chosen such that the period 1 response of leverage to the shock is constant across both policy tools.

more costly inflation becomes and the more attractive prudential tools are relative to monetary policy. When the labour margin is inelastic (high χ), larger movements in inflation are required to offset the effect of technology shocks on output and leverage, again tilting the balance towards prudential policy.

In short, policymakers should use both prudential policy and monetary policy to dampen fluctuations in leverage resulting from technology shocks. Policymakers should put greater reliance on prudential policy when technology shocks are persistent, the labour-output margin is inelastic, and the costs of inflation are high.

5 Financial-stability-focused monetary policy

We now study a regime in which the monetary authority targets financial stability directly, in the spirit of policies that lean aggressively against tightening financial conditions. Concretely, we ask what interest-rate path would be required to stabilise our measure of financial stress, the equity risk premium ρ_t, and what macroeconomic side-effects such a policy entails.

In the benchmark New Keynesian model, there is a unique "natural" real interest rate r_t^* that closes the output gap and stabilises inflation. In our framework, the equity premium reflects a risk-bearing wedge that depends on leverage and idiosyncratic risk. Following Akinci et al. (2021), we define an analogous financial-stability interest rate r_t^{**} as the real rate that stabilises this wedge. Appendix D shows that, in our linear approximation, stabilising ρ_t is equivalent to choosing monetary policy so that leverage moves one-for-one against the uncertainty shock: $l_t = -\xi_t$ (Remark 1 in the Appendix). Monetary policy must therefore expand aggregate demand when uncertainty rises, so as to raise entrepreneurial net worth and reduce leverage enough to offset the increase in risk.

The ability to implement such a rule is itself a consequence of financial amplification. When the representative household has log utility ($\sigma = 1$), there is no financial-stability interest rate that can keep ρ_t constant (Remark 2 in the Appendix). With $\sigma > 1$, the interaction between the leverage curve and the equity premium allows the central bank to engineer $l_t = -\xi_t$, but only at the cost of large movements in output and inflation.

Appendix D characterises these dynamics in detail. Proposition 2 shows that, in response to technology shocks, a financial-stability rule keeps the real interest rate and output at their steady-state levels. Recessionary technology shocks then manifest purely as higher inflation: the central bank prevents the deterioration of entrepreneurial balance sheets by holding the output gap at zero, so marginal costs and hence inflation rise. In contrast, following an increase in uncertainty, the same rule raises output and lowers the real rate in order to cut leverage. Because uncertainty shocks also act as cost-push shocks through the factor wedge, the nominal rate may move in either direction: in some calibrations, the financial-stability rule raises the policy rate even as it expands real activity.

The more fundamental problem is the time-series behaviour of such a policy. Proposition 3 shows that the financial-stability interest rate has a random-walk property: any temporary departure from r_t^{**}, or any use of prudential policy alongside it, generates permanent output and inflation gaps. Intuitively, once the economy is pushed into a high-inflation, positive-output-gap equilibrium that is consistent with financial stability, returning inflation to target would require the central bank to tolerate a period of financial stress. A policy that refuses to do so keeps financial conditions "too safe" at the cost of permanently higher inflation.

The interaction with macroprudential policy reinforces this point. Prudential tools can improve the on-impact response of leverage to uncertainty shocks, but under a financial-stability interest-rate rule they cannot realign the dynamic paths of leverage and the output gap. At some point, leverage will have returned to steady state while output and inflation have not; maintaining financial stability then means accepting a persistent deviation of inflation from target. In our framework, therefore, a narrowly defined financial-stability objective for monetary policy is dangerous: prudential policy can dampen but not remove the inflationary bias that arises from trying to keep financial stress low in all states.

6 Uncertainty shocks, safe assets, and the policy mix

We now return to uncertainty shocks. In our model these enter as shocks to the technology process, and macroprudential policy responds by adjusting leverage and balance-sheet exposures. The same ingredients that drive the paradox of safety in Section 2 — strong demand for safe or countercyclical assets, leverage-induced concentration of risk, and an increasing welfare cost of cycles — give uncertainty shocks first-order welfare consequences. The question in this section is how the policy mix between monetary and macroprudential tools should respond.

A growing empirical and theoretical literature finds that microeconomic uncertainty shocks can generate sizable macroeconomic fluctuations (for example, Arellano, Bai, and Kehoe, 2019; Bloom et al., 2018; Christiano, Motto, and Rostagno, 2014; Di Tella (2017)). In our framework, uncertainty shocks reduce aggregate demand and create a labour wedge: even with flexible prices, hours fall below the level implied by the marginal revenue product of labour. From the perspective of policy, they generate a trade-off analogous to a New Keynesian cost-push shock: a monetary authority that optimises under timeless commitment is willing to tolerate some inflation volatility in order to dampen volatility in hours and output.

Static and dynamic leverage constraints

At the start of period t, entrepreneurial leverage satisfies the static identity

$$l_t = x_t - q_t^e - i_{t-1} + \pi_t, \tag{6.1}$$

where x_t is (log) output, q_t^e is entrepreneurial net wealth at the start of the period, i_{t-1} is investment chosen in $t-1$, and π_t is inflation. Taken in isolation, (6.1) suggests that monetary stimulus in response to an uncertainty shock is counterproductive: for given q_t^e and i_{t-1}, higher output and inflation raise leverage, concentrating more risk on entrepreneurial balance sheets and increasing the wedges in labour and capital markets. In a one-period version of the model the competitive allocation is therefore constrained efficient and there is no role for policy beyond respecting this trade-off.

In the dynamic model, however, entrepreneurial net wealth q_t^e is endogenous to anticipated policy. If the central bank is expected to ease in recessions driven by uncertainty shocks, entrepreneurs choose to carry more equity into those states. The relevant constraint for policy is therefore the dynamic law of motion for leverage. Adapting the leverage curve from earlier, Equation (1.3), leverage evolves according to

$$l_t = \phi l_{t-1} + (1-\phi)\big(\omega\sigma\Delta\xi_t - \xi_{t-1}\big) - (1-\phi)\frac{\tilde{\sigma}}{\psi}\Delta x_t,$$

where ξ_t is idiosyncratic risk, $\Delta\xi_t = \xi_t - \xi_{t-1}$, $\tilde{\sigma} = \sigma - 1$, and $\phi \in (0,1)$ captures the persistence of leverage. Given this dynamic constraint, an increase in output x_t induced by expansionary monetary policy reduces leverage within the period: higher current demand raises entrepreneurial net worth, lowers the equity premium, and dampens the concentration of risk. Once we account for the endogenous accumulation of equity, the sign of the leverage response to demand stimulus is reversed relative to the static constraint (6.1).

Anticipated monetary accommodation of uncertainty shocks can therefore reduce the volatility of leverage and monitoring costs over the cycle, even though it would raise leverage on impact for a given q_t^e. In equilibrium, entrepreneurs choose higher q_t^e precisely because they expect the central bank to lean against uncertainty-driven downturns.

Monetary accommodation and moral hazard

In the absence of nominal rigidities, our flexible-price model is an Arnott–Greenwald–Stiglitz environment (Arnott and Stiglitz, 1986; Greenwald and Stiglitz, 1986): there is an information asymmetry between borrowers and lenders, while trade in other markets is competitive and anonymous. The standard lesson from this literature is that policy should target the *complements* of moral hazard, rather than attempting to observe effort directly.

In our setting, higher uncertainty raises the private payoff to opportunistic risk-taking and makes monitoring more costly. If firms entered uncertain business cycles with more equity — more skin in the game — the marginal benefit of taking excessive risk would fall, and the social cost of moral hazard would be smaller. Monetary stimulus in high-uncertainty states raises the value of entrepreneurial equity. By improving entrepreneurial balance sheets in precisely those states where risk-bearing is most costly, it discourages excessive risk-taking and mitigates its effects on employment and output.

Uncertainty shocks are therefore real shocks to the technology process, but they still warrant an accommodative monetary response alongside macroprudential policy. Macroprudential tools can shape the ex ante distribution of leverage across states; monetary policy improves the initial conditions — the strength of entrepreneurial balance sheets — so that prudential policy does not have to bear the full burden of containing leverage and financial stress when the next bout of uncertainty arrives.

7 Discussion

In the presence of idiosyncratic risk and contracting frictions, aggregate risk markets do not suppress the financial amplification and the competitive equilibrium is constrained inefficient.

Two macroprudential externalities follow. First, excessive financial stress builds in downturns not only in response to risk shocks, but also in response to technology and monetary shocks: leverage, monitoring costs and the equity premium all move in ways that private contracts do not internalise. Second, changes in the distribution of wealth between households and entrepreneurs generate aggregate demand externalities that are not fully priced by market participants, even when aggregate risk can be traded. In both cases, the wedge between private and social valuations runs through balance-sheet exposures and the joint behaviour of output, leverage and the equity premium.

In this setting, welfare-enhancing monetary and macroprudential policies typically need to be coordinated. Each instrument has a comparative advantage: other things equal, monetary policy is more powerful when disturbances are transitory, while macroprudential tools are better suited to persistent changes in risk and technology. Within bounds, the two can substitute for one another when one set of instruments is constrained, but the substitution is partial and state-dependent. A monetary policy that focuses too narrowly on financial stability risks losing control of the price level: in our model, keeping financial conditions "too safe" at all times can entail permanently higher inflation in response to temporary recessionary shocks, even when prudential policy leans in the right direction.

Taken together, these results suggest a 'modest' but durable role for monetary policy in financial stability: macroprudential tools should bear the primary burden of managing leverage and risk, but monetary policy cannot ignore financial conditions, nor can financial-stability objectives be pursued without regard to prices. Any practical division of labour between the two will have to respect this interaction between balance sheets, risk sharing and the price level.

References

Ozge Akinci, Gianluca Benigno, Marco Del Negro, and Albert Queraltó. The Financial (In)Stability Real Interest Rate, R*. International Finance Discussion Papers 1308, Board of Governors of the Federal Reserve System (U.S.), January 2021.

Franklin Allen and Kenneth Rogoff. Asset prices, financial stability and monetary policy. *The Riksbank's inquiry into the risks in the Swedish housing market*, pages 189–218, 2011.

Cristina Arellano, Yan Bai, and Patrick J. Kehoe. Financial frictions and fluctuations in volatility. *Journal of Political Economy*, 127(5):2049–2103, 2019. doi: 10.1086/701792.

Richard Arnott and Joseph E. Stiglitz. Moral hazard and optimal commodity taxation. *Journal of Public Economics*, 29(1):1–24, February 1986.

Anmol Bhandari, David Evans, Mikhail Golosov, and Thomas J. Sargent. Inequality, business cycles, and monetary-fiscal policy. *Econometrica*, 89(6):2559–2599, 2021. doi: https://doi.org/10.3982/ECTA16414.

Nicholas Bloom, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. Really Uncertain Business Cycles. *Econometrica*, 86(3):1031–1065, May 2018.

Kim C. Border and Joel Sobel. Samurai accountant: A theory of auditing and plunder. *The Review of Economic Studies*, 54(4):525–540, 1987. ISSN 00346527, 1467937X. URL `http://www.jstor.org/stable/2297481`.

Ricardo J Caballero and Emmanuel Farhi. The safety trap. *Review of Economic Studies*, 85(1):223–274, 2017.

Ricardo J Caballero and Alp Simsek. Prudential monetary policy. Technical report, National Bureau of Economic Research, 2019.

Xiaoshan Chen, Tatiana Kirsanova, and Campbell Leith. How optimal is us monetary policy? *Journal of Monetary Economics*, 92:96–111, 2017. ISSN 0304-3932.

Lawrence J. Christiano, Roberto Motto, and Massimo Rostagno. Risk shocks. *American Economic Review*, 104(1):27–65, January 2014. doi: 10.1257/aer.104.1.27.

Sebastian Di Tella. Uncertainty shocks and balance sheet recessions. *Journal of Political Economy*, 125(6):2038–2081, 2017. doi: 10.1086/694290.

Alfred Duncan and Charles Nolan. Disputes, debt and equity. *Theoretical Economics*, 14(3), September 2019.

Alfred Duncan and Charles Nolan. Destabilizing insurance. Technical report, 2021.

Emmanuel Farhi and Ivan Werning. A theory of macroprudential policies in the presence of nominal rigidities. *Econometrica*, 84(5):1645–1704, 2016. Lead article.

Mark Gertler and Nobuhiro Kiyotaki. Chapter 11 - financial intermediation and credit policy in business cycle analysis. volume 3 of *Handbook of Monetary Economics*, pages 547–599. Elsevier, 2010.

Bruce C Greenwald and Joseph E Stiglitz. Externalities in Economies with Imperfect Information and Incomplete Markets. *The Quarterly Journal of Economics*, 101(2):229–64, May 1986.

Anton Korinek and Alp Simsek. Liquidity trap and excessive leverage. *American Economic Review*, 106(3):699–738, 2016.

Luc Laeven, Angela Maddaloni, and Caterina Mendicino. Monetary policy, macroprudential policy and financial stability. Technical report, European Central Bank, 2022.

Alberto Martin, Caterina Mendicino, and Alejandro Van der Ghote. On the interaction between monetary and macroprudential policies. Technical report, European Central Bank, 2021.

Dilip Mookherjee and Ivan Png. Optimal auditing, insurance, and redistribution. *The Quarterly Journal of Economics*, 104(2):399–415, May 1989.

Stephanie Schmitt-Grohe and Martin Uribe. Prudential policy for peggers. Working Paper 18031, National Bureau of Economic Research, May 2012.

Kevin D. Sheedy. Debt and Incomplete Financial Markets: A Case for Nominal GDP Targeting. *Brookings Papers on Economic Activity*, 45(1 (Spring):301–373, 2014.

Jeremy C Stein. Monetary policy as financial stability regulation. *The Quarterly Journal of Economics*, 127(1):57–95, 2012.

Jeremy C Stein. Overheating in credit markets: origins, measurement, and policy responses. In *Speech given to the symposium on Restoring Household Financial Stability After the Great Recession, Federal Reserve Bank of St. Louis, St. Louis, Missouri, February*, volume 7, 2013.

Robert M. Townsend. Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory*, 21(2):265–293, October 1979.

A From the optimal contract to aggregate wedges (derivation sketch)

Wedges and their transformed version

The complete derivation of the optimal contract is in the online appendix. Here we summarise the key steps going from the optimal contract solution to the aggregate equations for leverage and the equity premium. Recall that firms hire factors before they know how productive they will be. Because of their risk aversion, there is at the firm level a labour wedge which is a proportional gap between the expected marginal product of labour and the wage:

$$\tau_{Nt} := \frac{\bar{Y}_{Nt} - W_t}{\bar{Y}_{Nt}} \qquad \Longleftrightarrow \qquad W_t = (1 - \tau_{Nt})\,\bar{Y}_{Nt},$$

where $\bar{Y}_{Nt}$ is the expected marginal product of labour for the representative entrepreneur.

For the contracting algebra it is convenient to work with a transformed wedge T_t defined by

$$T_t := \frac{\tau_{Nt}}{1 - \tau_{Nt}} \qquad \Longleftrightarrow \qquad \tau_{Nt} = \frac{T_t}{1 + T_t}, \tag{A.1}$$

so that

$$W_t = \frac{\bar{Y}_{Nt}}{1 + T_t}.$$

When wedges are small, (A.1) implies

$$T_t = \tau_{Nt} + \tau_{Nt}^2 + \mathcal{O}(\tau_{Nt}^3), \qquad \tau_{Nt} = T_t - T_t^2 + \mathcal{O}(T_t^3),$$

so to first order T_t and τ_{Nt} coincide. In the main text we therefore denote by τ_t the linearised labour wedge (equivalently, the linearised T_t) and reserve T_t for the level object used in the closed-form contract.

From the contract to $\tau_t \approx \tau_l l_t + \tau_\xi \xi_t$

In the imperfect state-verification problem, the optimal static contract implies a closed-form leverage ratio

$$L_t = \frac{\bar{\pi}}{\bar{\pi}\,\Xi_t - T_t}, \tag{A.2}$$

in the limit of vanishing type-I audit errors, where $\bar{\pi}$ is the probability of the high outcome, Ξ_t is the spread between high and low idiosyncratic outcomes, and T_t is the transformed labour wedge defined above.

Let (L_0, T_0, Ξ_0) denote the steady state, and define (log) deviations

$$l_t := \log \frac{L_t}{L_0}, \qquad \tau_t := \log \frac{T_t}{T_0}, \qquad \xi_t := \log \frac{\Xi_t}{\Xi_0}.$$

Taking a second-order Taylor expansion of $\log L_t$ around (L_0, T_0, Ξ_0) and solving for τ_t yields (see on-line Appendix A for the full algebra)

$$\tau_t = \tau_l \Big(l_t - \tfrac{1}{2} l_t^2 \Big) + \tau_\xi \Big(\xi_t + \tfrac{1}{2} \xi_t^2 \Big) + \mathcal{O}(z^3), \tag{A.3}$$

where $z := (l_t, \xi_t)$ and

$$\tau_l := \bar{\pi}\,\Xi_0 - T_0 > 0, \qquad \tau_\xi := \bar{\pi}\,\Xi_0 > 0$$

are functions of steady-state contract primitives. Dropping quadratic and higher-order terms in (A.3) gives the linear reduced form used in the main text:

$$\tau_t \approx \tau_l\, l_t + \tau_\xi\, \xi_t. \tag{A.4}$$

From the contract to $\rho_t \approx \psi(l_t + \xi_t)$

The same contract implies that the gross equity return satisfies

$$R_t = 1 + L_t T_t, \tag{A.5}$$

so the (log) equity premium ρ_t can be written as a function of (L_t, T_t). A second-order expansion of ρ_t around the steady state, combined with (A.3), yields

$$\rho_t = \psi\Big(l_t + \tfrac{1}{2}l_t^2 + \xi_t + \tfrac{1}{2}\xi_t^2 + l_t\xi_t\Big) + \mathcal{O}(z^3), \tag{A.6}$$

where

$$\psi := \frac{\bar{\pi} + L_0 T_0}{1 + L_0 T_0} > 0.$$

To first order in (l_t, ξ_t), (A.6) reduces to

$$\rho_t \approx \psi\,(l_t + \xi_t), \tag{A.7}$$

which is the reduced-form relationship used in the main text.

Equations (A.4) and (A.7) summarise the only channels through which the underlying contracting problem enters the IS and Phillips curves: a factor wedge increasing in leverage and idiosyncratic risk, and an equity premium with the same dependence.

B From static leverage to the dynamic leverage curve

This section sketches how the static leverage equation just derived gives rise to a persistent law of motion for leverage. The key point is that the contract is static *conditional on beginning-of-period net worth*, but entrepreneurial net worth itself is a slow-moving state variable. As a result, optimal leverage inherits persistence.

Static leverage and the role of net worth

For a representative entrepreneur, let Q_t denote the scale of the risky project (assets), N_t^e entrepreneurial net worth at the start of period t, and define leverage as

$$L_t := \frac{Q_t}{N_t^e}.$$

The imperfect state-verification problem yields a static contract that pins down the *ratio* L_t as a function of the labour wedge T_t and idiosyncratic risk Ξ_t. In the limit $\eta \to 0^+$ (no type-I audit errors), the optimal leverage ratio is

$$L_t = \frac{\bar{\pi}}{\bar{\pi}\,\Xi_t - T_t}, \tag{B.1}$$

where $\bar{\pi}$ is the probability of the high outcome, Ξ_t is the spread between high and low project outcomes, and T_t is the transformed labour wedge introduced just above.

Equation (B.1) is purely static: given (T_t, Ξ_t) and N_t^e, the entrepreneur chooses Q_t such that $Q_t/N_t^e = L_t(T_t, \Xi_t)$.

Net worth accumulation and persistence

Persistence arises because N_t^e is predetermined and evolves with past leverage and returns. Let R_t^E be the (gross) return on entrepreneurial equity between $t-1$ and t. A simple log-linear law of motion for entrepreneurial net worth around the steady state can be written as

$$n_t^e = \phi\, n_{t-1}^e + (1-\phi)\Big[\omega\sigma\Delta\xi_t - \xi_{t-1}\Big] - (1-\phi)\frac{\tilde{\sigma}}{\psi}\,\Delta x_t + \mathcal{O}(z^2), \tag{B.2}$$

where: $n_t^e := \log(N_t^e/N_0^e)$ is net worth in deviations from steady state, x_t is the (log) output gap, ξ_t is the (log) idiosyncratic risk process, ω is the steady-state household/entrepreneur consumption ratio, σ is household risk aversion, $\tilde{\sigma}$ is the composite risk-aversion parameter that appears in the IS curve, $\psi > 0$ is the elasticity of the equity premium with respect to leverage and risk (see Appendix A), $z := (x_t, \xi_t)$ collects small deviations, $\phi \in (0,1)$ summarises the persistence of entrepreneurial net worth induced by savings behaviour and contract structure.

Equation (B.2) is obtained by combining: (i) the entrepreneur's intertemporal Euler equation, (ii) the definition of R_t^E in terms of the contract and wedges, and (iii) the aggregate risk-sharing condition

$$\sigma\Delta c_t = \Delta c_t^e - \rho_t - \big(1 + \sigma\omega(1-\psi)\big)\delta_t,$$

with the reduced-form relationships

$$\tau_t \approx \tau_l l_t + \tau_\xi \xi_t, \qquad \rho_t \approx \psi(l_t + \xi_t),$$

derived in Appendix A. The term with Δx_t in (B.2) captures how higher current output raises leverage in the short run but also affects future net worth and hence future leverage; the terms in $(\Delta\xi_t, \xi_{t-1})$ capture the impact of risk shocks on net worth.

From (L_t, N_t^e) to the leverage curve

By definition,

$$l_t := \log \frac{L_t}{L_0} = \log Q_t - \log N_t^e - \big(\log Q_0 - \log N_0^e\big).$$

Using the fact that Q_t is chosen each period to satisfy the static relation (B.1), and linearising around steady state, we can write l_t as a linear function of n_t^e, ξ_t and x_t. Eliminating n_t^e between this relation and (B.2) yields

$$l_t = \phi\, l_{t-1} + (1-\phi)\Big[\omega\sigma\Delta\xi_t - \xi_{t-1}\Big] - (1-\phi)\frac{\tilde{\sigma}}{\psi}\,\Delta x_t - (1-\phi)\,\delta_t + \mathcal{O}(z^2), \quad \text{(B.3)}$$

which is the leverage curve used in the main text, up to notation.

Equation (B.3) makes the source of persistence explicit: The static contract pins down leverage as a function of wedges and risk *within* each period. On the other hand, entrepreneurial net worth is a slow-moving stock that carries the effects of past leverage decisions into the future. And so, combining these two elements yields an AR(1)-type law of motion for leverage with coefficient $\phi \in (0, 1)$.

In other words, leverage is chosen each period *given* net worth, but net worth evolves endogenously with past leverage and returns. This is why a seemingly static contract generates a persistent leverage process.

C Technology shocks, leverage dynamics, and policy instruments

This appendix derives the results used in Section 4 for technology shocks, risk aversion, and the interaction between monetary and macroprudential policy. We start from the simplified system under Assumption 2 in the text, derive the financial amplification factor, and obtain the leverage dynamics underlying the charts and results reported in Section 4.

C.1 Baseline system under Assumption 2

Under Assumption 2 we (i) shut down uncertainty shocks and (ii) restrict monetary policy to rules that deliver zero expected inflation. Technology follows

$$a_t = \rho_a a_{t-1} + \varepsilon_{at}, \qquad \varepsilon_{at} \sim \text{i.i.d.}(0, \sigma_a^2). \tag{C.1}$$

The flexible-price Phillips curve and leverage curve reduce to

$$(\tilde{\sigma} + \chi) x_t = -\mu_l l_t + \chi a_t + \gamma \varepsilon_{at}, \tag{C.2}$$

$$l_t = \phi l_{t-1} - (1 - \phi) \frac{\tilde{\sigma}}{\psi} (x_t - x_{t-1}) - \delta \varepsilon_{at}, \tag{C.3}$$

where: $\tilde{\sigma} = \sigma - 1$ is household risk aversion minus one; $\chi = \frac{1+\phi}{1-\alpha}$ is the usual composite parameter from labour supply and production; $\mu_l > 0$ is the elasticity of marginal cost with respect to leverage (from the financial friction); $\phi \in (0, 1)$ measures the persistence of leverage; $\psi > 0$ is the elasticity of the equity premium with respect to leverage and risk; and γ, δ are the monetary and macroprudential policy parameters, respectively.

Equation (C.2) says that marginal cost (hence inflation, in the sticky-price version) depends on output, leverage and technology, plus the direct monetary-policy term $\gamma \varepsilon_{at}$. Equation (C.3) says that leverage is persistent, reacts to changes in output, and is directly affected by macroprudential policy $\delta \varepsilon_{at}$.

C.2 Financial amplification factor

To characterise the instantaneous amplification of technology shocks, it is useful to combine (C.2) and (C.3). Substitute (C.3) into (C.2) and collect terms in x_t and l_{t-1}:

$$(\tilde{\sigma} + \chi)x_t + \mu_l \phi l_{t-1} - \mu_l(1-\phi)\frac{\tilde{\sigma}}{\psi}(x_t - x_{t-1}) + \mu_l \delta \varepsilon_{at} = \chi a_t + \gamma \varepsilon_{at}.$$

Rearranging gives

$$\left[(\tilde{\sigma}+\chi) - \mu_l(1-\phi)\frac{\tilde{\sigma}}{\psi}\right] x_t = \chi a_t + \gamma \varepsilon_{at} - \mu_l \phi l_{t-1} + \mu_l(1-\phi)\frac{\tilde{\sigma}}{\psi} x_{t-1} - \mu_l \delta \varepsilon_{at}. \quad \text{(C.4)}$$

Define

$$m := \frac{\tilde{\sigma} + \chi}{(\tilde{\sigma} + \chi) - \mu_l(1-\phi)\frac{\tilde{\sigma}}{\psi}}. \quad \text{(C.5)}$$

Then we can rewrite (C.4) as

$$x_t = m\left[\frac{\chi}{\tilde{\sigma}+\chi} a_t + \frac{\gamma - \mu_l \delta}{\tilde{\sigma}+\chi}\varepsilon_{at} - \frac{\mu_l \phi}{\tilde{\sigma}+\chi} l_{t-1} + \frac{\mu_l(1-\phi)\tilde{\sigma}}{\psi(\tilde{\sigma}+\chi)} x_{t-1}\right]. \quad \text{(C.6)}$$

The factor $m > 1$ is the *financial amplification factor*. In the absence of leverage feedback ($\mu_l = 0$), we have $m = 1$ and (C.6) collapses to the standard flexible-price response to technology shocks. When $\mu_l > 0$, a fall in output raises leverage, which in turn raises marginal costs and further depresses output. The term

$$\frac{\mu_l}{\tilde{\sigma}+\chi} \cdot \frac{(1-\phi)\tilde{\sigma}}{\psi}$$

is the product of (i) the sensitivity of marginal cost to leverage relative to output, and (ii) the sensitivity of leverage to changes in output. The closer this product is to one, the larger the amplification factor m.

C.3 Leverage dynamics under technology and policy shocks

Substituting (C.6) back into the leverage equation (C.3) and iterating forward yields a closed-form expression for leverage as a function of current and past technology shocks $\{\varepsilon_{at-\tau}\}$ and the policy parameters γ, δ. After some algebra, the solution for

l_t can be written as

$$l_t = -\frac{m-1}{\mu_l}\,\chi \left[\sum_{\tau=0}^{\infty} \phi^\tau \varepsilon_{a,t-\tau} - (1-\rho_a)\sum_{\tau=1}^{\infty}\sum_{j=1}^{\tau} \phi^{j-1}\rho_a^{\tau-j}\varepsilon_{a,t-\tau}\right] \qquad \text{(shock)}$$

$$-\frac{m-1}{\mu_l}\,\gamma \left[\sum_{\tau=0}^{\infty} \phi^\tau \varepsilon_{a,t-\tau} - \sum_{\tau=1}^{\infty} \phi^{\tau-1}\varepsilon_{a,t-\tau}\right] \qquad \text{(monetary policy)}$$

$$-\,m\,\delta \sum_{\tau=0}^{\infty} \phi^\tau \varepsilon_{a,t-\tau}. \qquad \text{(prudential policy)}$$

This is the key analytic equation underlying the results of Section 4. It is decomposed into three contributions. There is the effect of the technology shock itself (first line); the additional effect coming from the monetary policy response γ (second line); and the additional effect of the macroprudential policy response δ (third line).

The double sum in (shock) captures the predictable reversal of a transitory technology shock: a one-period negative innovation $\varepsilon_{at} < 0$ reduces output and raises leverage on impact, but—because technology is mean-reverting—raises output growth and reduces leverage in subsequent periods. The same logic applies to the self-reversing component of the monetary policy term: under the zero-expected-inflation restriction, any expansionary monetary policy today must be followed by expected tightening tomorrow. The macroprudential policy contribution in (prudential policy) is purely forward-looking: a change in δ today alters the impact of ε_{at} on leverage and, through persistence, on future leverage as well.

C.4 Persistence and relative merits of monetary vs prudential tools

Equations (monetary policy) and (prudential policy) clarify the different persistence properties of the two instruments.

When technology shocks are highly persistent ($\rho_a \to 1$), the ϕ-geometric sums multiplying δ in (prudential policy) closely track the natural leverage response to the shock. A prudential intervention therefore dampens leverage for the *full duration* of a persistent technology shock. By contrast, the monetary-policy term (monetary policy) is largely self-reversing under the zero-expected-inflation constraint: expansionary policy today is offset by expected tightening tomorrow, so the leverage impact of γ

is strongest on impact and weaker thereafter.

When technology shocks are instead transitory ($\rho_a \to 0$), the double-sum term in (shock) shrinks and the self-reversing component of (monetary policy) becomes more closely aligned with the natural reversal of the shock. In that case, a monetary intervention that leans against leverage on impact also dampens leverage in subsequent periods, while a prudential intervention tends to overshoot—reducing leverage too strongly on impact and raising leverage relative to the no-prudential-policy benchmark later on.

These patterns are illustrated in Figure 4 in the main text, where we choose γ and δ so that both instruments deliver the same on-impact response of leverage to a unit recessionary technology shock. When shocks are persistent, prudential policy dominates in terms of leverage stabilisation over the life of the shock; when shocks are i.i.d., monetary policy does.

Taken together, the derivations above justify the qualitative conclusions in Section 4: both instruments should generally be used; monetary policy is relatively better suited to addressing the financial amplification of short-lived technology shocks, while macroprudential policy is relatively better suited to persistent shocks; and the strength of these conclusions depends on the amplification factor m, which itself increases with risk aversion and the leverage sensitivity of marginal costs.

D Financial-stability-focused monetary policy: derivations

This appendix provides the formal derivations underlying Section 5. We characterise the interest-rate rule that stabilises the equity risk premium, discuss conditions for its existence, and derive the properties stated in Propositions 2 and 3. We also contrast the financial-stability real rate r_t^{**} with the standard New Keynesian natural rate $r^*_{NK,t}$.

D.1 Stabilising the equity premium and the condition $l_t = -\xi_t$

Recall from Appendix A that the equity risk premium admits the second-order approximation

$$\rho_t = \psi \left(l_t + \tfrac{1}{2} l_t^2 + \xi_t + \tfrac{1}{2}\xi_t^2 + l_t \xi_t \right) + \mathcal{O}(z^3),$$

where l_t is log leverage, ξ_t is the log deviation of idiosyncratic risk from steady state, and $\psi > 0$ collects the dependence of the premium on leverage and risk.[7]

Remark 1 *The equity risk premium ρ_t is stabilised when $l_t = -\xi_t$.*

Proof. Setting $\rho_t = 0$ in the above approximation and neglecting higher-order terms yields

$$l_t + \xi_t + \tfrac{1}{2}(l_t^2 + \xi_t^2 + 2l_t\xi_t) \approx 0,$$

which is satisfied when $l_t = -\xi_t$ to second order. ■

Remark 1 is the basis for our definition of the financial-stability interest rate r_t^{**}: it is the (real) rate that, together with macroprudential policy, implements $l_t = -\xi_t$ for all t in equilibrium.

Combining $l_t = -\xi_t$ with the linearised leverage curve (B.6) yields the condition that the financial-stability interest rate must satisfy:

$$(1-\phi)\tfrac{\tilde{\sigma}}{\psi}\Delta x_t = (1 + (1-\phi)\sigma\omega)\,\Delta\xi_t - \delta_t, \tag{D.1}$$

where δ_t is the macroprudential wedge.

D.2 Existence of a financial-stability interest rate

The ability of monetary policy to stabilise ρ_t depends on financial amplification, which in turn requires households to be more risk-averse than entrepreneurs.

Remark 2 *When the representative household has log utility ($\sigma = 1$), no financial-stability interest-rate policy exists.*

Proof. When $\sigma = 1$, there is no amplification of technology and uncertainty shocks through the leverage curve: the parameter combination $(1-\phi)\tilde{\sigma}/\psi$ in (D.1) collapses to zero. Then l_t does not respond to ξ_t in a way that can be controlled by monetary policy, so there is no interest-rate rule that can implement $l_t = -\xi_t$ for all t. ■

[7] See Appendix A for the derivation of this expression.

D.3 Characterising responses to technology and uncertainty shocks

We now characterise the dynamics of output, interest rates and leverage under the financial-stability interest-rate policy, abstracting from prudential intervention ($\delta_t = 0$).

Proposition 2 *(Characterisation of the financial-stability interest rate in the absence of prudential policy.)*

a. *In response to technology shocks, a financial-stability interest-rate policy holds the real interest rate and output constant.*

b. *In response to an increase in uncertainty, the financial-stability interest-rate policy allows output to increase and the real interest rate to fall. The effect on the nominal interest rate is ambiguous.*

Sketch of proof. Part (a) follows from imposing $l_t = -\xi_t$ and $\delta_t = 0$ in the log-linearised system and using the fact that technology shocks affect leverage only through their impact on output and consumption. A financial-stability rule must therefore offset the output response to technology shocks, keeping x_t at zero and stabilising real rates.

For part (b), an increase in ξ_t raises the equity premium and would normally contract activity. Under the constraint $l_t = -\xi_t$, monetary policy must instead expand demand so that leverage falls one-for-one with the rise in uncertainty. This requires a lower real interest rate. Because uncertainty shocks also enter the Phillips curve as cost-push shocks, the nominal interest rate may either rise or fall depending on whether the inflation effect dominates the real-rate effect. ■

D.4 Random-walk property and interaction with prudential policy

The key time-series property of the financial-stability interest-rate rule is summarised in the following proposition.

Proposition 3 *(Random-walk property of the financial-stability interest rate.)*

a. *Temporary departures from the financial-stability interest rate (i.e. monetary shocks) result in permanent output and inflation gaps.*

b. Prudential policies, in combination with the financial-stability interest rate, result in permanent output and inflation gaps.

Sketch of proof. Under the rule that keeps ρ_t constant via $l_t = -\xi_t$, the equilibrium paths of output, inflation and leverage are pinned down jointly. A one-off deviation from r_t^{**} alters the level of output and inflation but, because the rule subsequently re-imposes $l_t = -\xi_t$, the system converges to a new stochastic steady state with a different output and inflation level. Analogous reasoning applies when prudential policy alters the on-impact response of leverage to uncertainty shocks: the financial-stability rule forces leverage back to its original path while leaving the output gap and inflation at a shifted level. ■

This random-walk property underlies the inflationary bias discussed in Section 5: once the economy is in a high-inflation, positive-output-gap equilibrium compatible with financial stability, returning to target inflation would require the central bank to tolerate a period of financial stress.

D.5 The three interest rates: r_t^*, r_t^{**}, and the profit rate

Finally, it is useful to contrast the real interest rate that stabilises inflation under strict inflation targeting with the financial-stability interest rate and the implied profit rate.

Under strict inflation targeting, the real interest rate r_t^* and equity risk premium ρ_t in our model can be expressed as

$$r_t^* = -\frac{m\sigma'\chi}{\tilde{\sigma}+\chi}\left((1-\rho_a) + \frac{(1-\phi)^2\tilde{\sigma}}{\psi}\left(\underbrace{\frac{m\mu_l}{\tilde{\sigma}+\chi}}_{\text{cost push}} - \underbrace{\sigma\omega(1-\psi)}_{\text{demand pull}}\right)\right)a_t + \mathcal{F}(\Omega_{t-1}),$$

$$\rho_t = -\frac{(m-1)\chi\psi}{\mu_l}a_t + \mathcal{F}'(\Omega_{t-1}),$$

where m is the amplification factor from Appendix C, a_t is the technology process, and $\mathcal{F}(\Omega_{t-1})$ and $\mathcal{F}'(\Omega_{t-1})$ collect terms measurable at $t-1$. The term $1/\sigma'$ is the population intertemporal elasticity of substitution,

$$\frac{1}{\sigma'} = \frac{1}{1+\omega(1-\psi)}\left(\frac{1}{\sigma} + \omega(1-\psi)\cdot 1\right).$$

For comparison, in the standard three-equation New Keynesian model the natural real rate is

$$r^*_{NK,t} = -\frac{\sigma\chi}{\tilde{\sigma} + \chi}(1 - \rho_a)a_t.$$

Financial amplification ($m > 1$), the passthrough of output to leverage, and the demand effects of leverage all modify the required real-rate response to stabilise inflation in our model relative to the standard benchmark.

The profit rate $r^*_t + \rho_t$ therefore responds much more strongly to contractionary technology shocks under strict inflation targeting than under the financial-stability interest-rate policy, where both the real rate and equity premium are stabilised by construction.