



When monetary policy for financial stability backfires: Communication with strategic banks

Giampaolo Bonomi ^{a,*}, Ali Uppal ^b

^a Princeton University, 001 Fisher Hall, Princeton, New Jersey 08544-1012, United States

^b Imperial College London, London, United Kingdom

ARTICLE INFO

Keywords:

Monetary policy
Financial stability
Communication
Forward guidance
Central bank design
Strategic banks
Cheap talk

ABSTRACT

We develop a model in which the central bank incorporates financial stability concerns into its monetary policy and communication decisions, while systemic banks strategically choose portfolios to influence policy. Because fully responding to economic shocks would generate policy surprises that threaten bank stability, the central bank underreacts to those shocks, consistent with the Federal Reserve's response during the 2023 banking crisis. Anticipating this accommodation, banks tilt their portfolios to increase the financial-stability costs of rate hikes, pulling monetary policy toward lower rates while giving the central bank an incentive to announce higher future rates than it expects to set. Financial stability concerns can therefore backfire, creating a novel responsiveness–credibility trade-off: the more the central bank accommodates banks' exposures ex post, the less credible its forward guidance becomes ex ante, making it harder for banks to align their portfolios with future policy and leaving them more exposed to policy surprises. Even when society values financial stability, it benefits from appointing a central banker who places less weight on it. Paradoxically, even banks may prefer such a central banker when the benefits of more credible guidance outweigh the loss of policy influence.

1. Introduction

*How much does it matter if the market faces surprises, and should we care if we surprise them? Well, I think we should care, because we've said a lot about how long we're going to wait until we raise the rate, and a lot of people have built strategies—I'm not sure they've built their futures, their future savings and all of that, on that particular result. For the long run, it doesn't much matter. But it will lead to a view that the FOMC has got things wrong again and didn't send the market a warning until it was much too late. So I think we should be moving a little bit earlier in changing the message we're sending.*¹

The Global Financial Crisis of 2007–08 challenged the consensus that monetary policy should focus primarily on inflation (Smets, 2014). Since then, central banks have increasingly recognized that interest rate decisions can affect bank balance sheets and systemic risk, leading to a broad debate about when and how financial stability should enter the conduct of monetary policy (Woodford, 2012). In practice, many central banks now explicitly or implicitly account for financial stability concerns in their decision-making.² Yet

* Corresponding author.

E-mail addresses: bonomi@princeton.edu (G. Bonomi), a.uppal@imperial.ac.uk (A. Uppal).

¹ Governor Fischer, FOMC Transcript, July 29–30, 2014 (<https://www.federalreserve.gov/monetarypolicy/files/FOMC20140730meeting.pdf>)

² There is anecdotal evidence that central banks are concerned about financial stability (e.g., from transcripts and speeches) as well as empirical evidence. For example, Shapiro and Wilson (2022) find that the Fed is concerned about stock market performance rather than economic slack, while

incorporating such concerns raises two important questions: could financial institutions strategically alter their behavior knowing that the central bank is concerned about their stability? And how does this strategic response affect the central bank's ability to communicate with markets?

Both questions have the same source: once financial stability enters the central bank's objective, banks' portfolio choices become relevant for monetary policy. Large systemic banks may then choose positions not only to prepare for future rates, but also to make adverse policy changes more costly for the central bank. This strategic positioning can make policy less responsive to shocks and tilt it toward the rates preferred by banks. Communication is central because it shapes what banks know when choosing those positions. Guidance helps financial institutions avoid costly surprises, but it also gives them information they can use to position themselves in ways that constrain future policy. Anticipating this response, the central bank cannot credibly provide fully transparent guidance. Financial stability concerns can therefore backfire: the concern that makes guidance valuable also makes it less credible, leaving markets more exposed to costly surprises.

To formalize this tension and isolate the mechanisms at play, we develop a parsimonious game-theoretic framework with a financial-stability-oriented central bank and *strategic* banks (i.e., banks know the central bank cares about their stability). The central bank privately observes an early shock to the policy rate that would stabilize the economy and sends a public message before the final policy decision. Banks then choose portfolio positions based on this message. A second shock is realized before the central bank sets the policy rate. The model has two key ingredients. First, banks face a *policy surprise cost*: in the spirit of the volatility aversion in [Stein and Sunderam \(2018\)](#), unexpected policy rates trigger costly portfolio readjustments, and the central bank partly internalizes these costs in its pursuit of financial stability. Second, banks have a *directional preference* for lower interest rates, akin to the "Fed put" in [Cieslak and Vissing-Jorgensen \(2021\)](#). Unlike the surprise cost, this directional preference does not enter the central bank's loss directly, but it matters in equilibrium because strategic banks tilt their portfolios toward lower-rate positions in anticipation of the central bank's surprise aversion. The key friction is that large systemic banks are non-atomistic: their positions affect the aggregate surprise cost that enters the central bank's objective. We refer to this channel as banks' policy influence.

Our framework links forward guidance, financial stability, and communication credibility within a unified setting and delivers four main insights. First, if financial institutions face costly readjustments when their investment position is not consistent with the policy rate (i.e., the policy surprise cost), and the central bank internalizes part of this cost, then the central bank underreacts to shocks to the economy—consistent with the gradualism approach highlighted by [Stein and Sunderam \(2018\)](#). Our result also echoes the insight of [Caballero and Simsek \(2022\)](#): market expectations pull the policy rate toward the market's view even when the central bank disagrees. A case study of Silicon Valley Bank and Signature Bank illustrates this underreaction ([Appendix C](#)).

Second, when banks are non-atomistic (e.g., "too big to fail" systemic banks), financial stability concerns distort their portfolio choices.³ Each systemic bank understands that its position affects the aggregate surprise cost entering the central bank's objective. Because banks prefer lower rates, they choose positions that make higher rates more costly for the central bank. These positions are biased toward lower-rate portfolios relative to the economy-stabilizing policy. The bias distorts both investment and monetary policy: anticipated policy influence pushes the central bank toward a lower rate than it would choose absent strategic positioning. The distortion is larger when the central bank places more weight on financial stability, when banks' directional preference for lower rates is stronger, and when each bank has greater policy influence.

Third, strategic positioning creates an information loss. Financial stability concerns create a tension between avoiding surprises *ex post* and communicating effectively *ex ante*. If the central bank could commit to transparent communication before the policy decision, it would fully reveal its information and help banks avoid surprise losses. Without commitment, however, communication is cheap talk. Once banks use guidance to choose positions that make higher rates more costly for the central bank, the central bank has an incentive to exaggerate its message in the opposite direction. As in [Crawford and Sobel \(1982\)](#), this conflict prevents full information transmission: equilibrium communication takes the form of a partition equilibrium in which the central bank transmits only coarse information about its policy intentions. Thus, financial stability concerns generate a novel trade-off: they make guidance valuable because banks dislike surprises, but less credible because banks use the information it conveys to influence policy and the policymaker responds by communicating strategically.

Fourth, the model has implications for central bank design. In the spirit of the "conservative central banker" in [Rogoff \(1985\)](#), society benefits from appointing a less *financial-stability-oriented* ("FS-oriented") central banker, i.e., one who places less weight on financial stability than society itself does. Such a central banker reduces banks' incentive to distort their positions and thereby improves both policy and communication. Banks' preferences differ from society's under transparent communication: because a more FS-oriented central banker is more responsive to banks' surprise losses, banks prefer a higher financial stability weight. Under cheap talk, however, banks may paradoxically prefer a central banker who cares less about their stability: weaker policy influence reduces the conflict of interest in communication, and the resulting clearer guidance can outweigh the loss of influence.

A repeated-game extension shows that dynamic incentives need not eliminate these distortions. If transparent guidance is sufficiently valuable to banks, the central bank can threaten to withhold future information after deviations and thereby sustain efficient play. But repeated interaction can also facilitate coordination among banks, allowing them to act collectively as a larger strategic

[Peek et al. \(2015\)](#) provide evidence that the Fed often behaves as if monetary policy has an explicit financial stability mandate, over and above its traditional mandate. See [Appendix D](#) for evidence on central bank mandates across the world.

³ We capture a non-atomistic banking sector by a smaller number of banks: with fewer banks, each bank accounts for a larger share of the financial stability concern internalized by the central bank. This reduced-form interpretation is consistent with the role of size in systemic-risk assessment and with evidence on rising U.S. banking concentration; see, e.g., [Corbae and D'Erasmo \(2021\)](#).

player and increasing policy distortions. Dynamic incentives therefore have ambiguous implications: depending on parameters, they can either discipline banks or amplify their policy influence.

Overall, we contribute to two literatures. First, on central bank communication, there is considerable evidence on both formal and informal communication (see, for instance Hansen and McMahon, 2016; Cieslak et al., 2019; Hansen et al., 2019; Bradley et al., 2020; Morse and Vissing-Jørgensen, 2020). While most existing work focuses on central banks communicating a decision (Blinder et al., 2008), our model explicitly considers communication *prior* to a decision being made given that pre-decision speeches and press releases increasingly “guide” expectations in ways that can affect the decision itself.⁴ To our knowledge, only Vissing-Jørgensen (2020) formalizes *pre-decision* communication—an increasingly prevalent form of communication in practice—but her analysis abstracts from both the benefits of communication and from strategic market behavior. As such, we provide a formal analysis of pre-decision communication when the audience is strategic, non-atomistic, and benefits from guidance. Existing studies of forward guidance or cheap talk between central banks and markets typically abstract from financial stability motives or from the role of banks’ systemic importance (e.g., Blinder et al., 2008; Vissing-Jørgensen, 2020). When the central bank internalizes market fragility (e.g., Woodford, 2012; Stein and Sunderam, 2018), we identify a previously unrecognized trade-off between forward guidance clarity and financial stability, and we show how central bank-type selection can mitigate this trade-off.

Second, and more broadly, the model identifies a responsiveness–credibility trade-off for an informed decision maker. Responding to receivers’ actions is optimal *ex post* because those actions affect the consequences of the final choice. Yet this responsiveness induces each receiver to distort its action; anticipating the distortion, the decision maker has an incentive to shade her message in the opposite direction, weakening the credibility of communication *ex ante*. From the receivers’ perspective, the same mechanism creates an influence–information trade-off: each receiver’s post-message action becomes a lever for steering the final choice, but exploiting this leverage leads the decision maker to provide less informative guidance. Formally, the action distortion pins down the bias in an induced Crawford–Sobel game, linking the tilt in the final choice to the coarseness of prior communication. Whereas in Crawford and Sobel (1982) the sender–receiver bias is a primitive, here the induced bias is an equilibrium object: it grows with individual receiver influence and vanishes as receivers become atomistic, in which case a fully revealing equilibrium attains the first best. Related work derives communication conflict from reputational motives (Morris, 2001; Ottaviani and Sørensen, 2006), stakes chosen before communication (Antić and Persico, 2020), or payoff externalities among atomistic receivers (Lukyanov et al., 2022); here, the conflict arises because receivers’ post-message actions alter the informed party’s subsequent choice.⁵

2. Setting

Game structure. We model the interaction between the central bank (CB) and a set of market players I using the following stylized three-period setting. At time $t = 0$, the CB privately observes a shock $\omega_1 \sim U(-\phi_1, \phi_1)$ to the economy-stabilizing policy (i.e., the interest rate consistent with stabilizing the economy), and sends a public message $m \in M = [-\phi_1, \phi_1]$ potentially informative about ω_1 . At $t = 1$, after having observed m , each market player $i \in I$ simultaneously chooses an investment position $x_i \in X$, based on their expectation of the central bank interest rate. The profile of investment positions $x = (x_i)_{i \in I}$ is observed by the CB.⁶ Finally, at $t = 2$, the CB observes a second shock to the economy-stabilizing policy, $\omega_2 \sim U(-\phi_2, \phi_2)$, independent of ω_1 , and chooses the policy rate $r \in \mathbb{R}$.⁷ The economy-stabilizing policy at the time of the decision is $\omega \in [-\phi, \phi]$, where $\omega = \omega_1 + \omega_2$ and $\phi = \phi_1 + \phi_2$. After the policy decision is made, the game ends and payoffs are realized. Fig. 1 provides a graphical representation of the timing.

Players and payoffs. As mentioned above, there are two types of players: the central bank (CB) and a set of market players (I). We make three main assumptions that shape the strategic incentives of these players in our setting.

First, market players incur a cost when they choose investment positions that are not well-aligned with the policy decision r . Second, *ceteris paribus*, market players have a directional preference for lower rates. Formally, let $X = \mathbb{R}$ and x_i denote the element of X chosen by i . We assume that the payoff function of player $i \in I$ takes the form

$$u_i(x_i, r) = \underbrace{-\frac{1}{2}(x_i - r)^2}_{\text{policy surprise}} \underbrace{- \beta r}_{\text{directional preference}} \quad (1)$$

where $(x_i - r)^2$ is the mismatch between the investment position and the chosen policy rate, and βr captures market players’ directional preference over rates. Without loss of generality, we assume $\beta > 0$, which captures a preference for low rates. The third payoff

⁴ Papers as early as Stein (1989) showed that the Fed can communicate some information about its goals through the cheap talk mechanism of Crawford and Sobel (1982).

⁵ Related feedback models show that policy responsiveness can degrade the informativeness of prices when the authority learns from markets (Bond and Goldstein, 2015; Boleslavsky et al., 2017). There, market behavior affects policy through the information it conveys; here, receivers’ actions directly affect the informed policymaker’s subsequent choice, limiting what she can credibly communicate beforehand.

⁶ The ‘Large Institution Supervision Coordinating Committee’ at the Fed has the ability to observe investment portfolios of systemically important banks. We discuss this observability assumption further in Appendix D. Moreover, in Appendix B.2, we allow for the possibility that the central bank cannot observe investment plans and derive qualitatively similar predictions.

⁷ The uniform distribution assumption is without loss of generality as far as ω_2 is concerned. Assuming that ω_1 is also uniformly distributed simplifies the communication analysis of Section 3.4.




Fig. 1. Timeline.

assumption is that the loss function of the central bank is the sum of a conventional loss term and a market readjustment cost and thus takes the following form. For $\alpha \in (0, 1)$ and $\omega \in [-\phi, \phi]$,

$$L_{CB}(x, r, \omega) = (1 - \alpha) \underbrace{\frac{1}{2}(r - \omega)^2}_{\text{conventional loss}} + \alpha \underbrace{\frac{1}{2N} \sum_{i \in I} (r - x_i)^2}_{\text{market readjustment cost}}, \quad (2)$$

where $N = |I|$ and α captures the welfare weight on market readjustment costs.⁸

Strategies and equilibrium. The solution concept that we adopt for the equilibrium analysis of this simple game is Perfect Bayesian Equilibrium (PBE). We restrict our attention to equilibria in which the central bank choice r at $t = 2$ is not contingent on the announcement $m \in M$ at $t = 0$. Considering equilibria in which the central bank conditions the policy choice r on previous communication would not yield additional insights, in the absence of reputational considerations.⁹ Hence, a strategy for the central bank consists of a communication rule $\sigma_m : [-\phi_1, \phi_1] \rightarrow \Delta(M)$, mapping realizations of ω_1 to a distribution over messages in M , and a policy rule $\sigma_r : [-\phi_1, \phi_1] \times [-\phi_2, \phi_2] \times X^I \rightarrow \Delta(\mathbb{R})$, mapping realizations of (ω_1, ω_2) and market players' positions to distributions over rate decisions. A strategy for market player i is an investment plan $\sigma_{x_i} : M \rightarrow \Delta(\mathbb{R})$, mapping the central bank's messages to investment positions. We write $\sigma_x = (\sigma_{x_i})_{i \in I}$ and denote a generic strategy profile $(\sigma_m, \sigma_x, \sigma_r)$ by $\sigma \in \Sigma$.

2.1. Mapping the structure to central banking

In this section, we briefly map the model's assumptions to the motivating central banking application. Further details on the empirical validity of our assumptions and interpretations for central banking are in [Appendix D](#). While our model is designed to fit the central banking application, the structure may also apply to studying the strategic tensions arising when large firms shape policy, both directly and through anticipation of their influence on policymakers. We discuss this generalizability of our model in [Section 5](#).

Timing. Our game occurs within the Federal Open Market Committee (FOMC) cycle, i.e., the six-week period between two consecutive policy decisions. This feature is distinct from much of the literature that considers decisions at the time of announcement. In our game, $t = 2$ is the end of the cycle, when the policy decision is made and announced. Within the cycle, official communication occurs at fixed dates to increase the transparency of the FOMC's decision process and provide future guidance to markets.¹⁰ We assume that policy-relevant communication occurs only once within the cycle, at $t = 0$, when at the time of communication, the policymaker has some uncertainty about the economy-stabilizing policy (ω_2).¹¹ Market players make changes to their investment positions before the policy decision r is made and as such, at $t = 1$ they choose their investment positions based on their future policy expectations. This timing structure allows for "forward guidance" (e.g., [McKay et al., 2016](#)) and highlights the value that players derive from correctly predicting future policies.

Market player payoffs. Market player i 's payoff is decreasing in the mismatch between her investment position x_i and the policy r . x_i represents a choice from a continuum of portfolios of financial securities, where portfolio $x_i = r$ is the one performing best at $t = 2$ if the policy is r . At $t = 1$, the mismatch is minimized in expectation when i holds a position x_i matching the expected future rate. The further the realized rate is from this expectation, the worse their portfolio performs ex post—our notion of a *policy surprise*.¹² This loss, which we assume to be quadratic, gives market players a clear incentive to forecast monetary policy accurately.¹³ We posit that

⁸ In [Appendix B.1](#), we discuss an alternative payoff specification where we allow market players to disagree with the central bank about the economy-stabilizing rate. The alternative specification delivers qualitative results analogous to those obtained under specification (1).

⁹ Conditional on ω and x , independence of the rate choice r from m is a necessary property of any PBE of this game. Therefore, focusing on policy rules that do not condition on communication is without loss.

¹⁰ See [Faust \(2016\)](#) for an extensive discussion of the role of communication.

¹¹ This simplification is consistent with the fact that the speeches of the FOMC chair and vice-chair occur once per cycle, and that the committee observes a "blackout period" before the end-of-cycle meetings.

¹² We use "policy surprise cost" as shorthand for the readjustment cost bank i faces when the realized rate r differs from x_i . Note that such a mismatch may reflect genuinely unanticipated policy variation, since ω_2 is realized after investment, but a strategic bank may deliberately choose $x_i \neq \mathbb{E}[r | m]$ to influence policy.

¹³ In addition to assuming quadratic loss, we follow [Stein and Sunderam \(2018\)](#) and [Vissing-Jorgensen \(2020\)](#) in also assuming additive separability, which ignores feedback loops but greatly simplifies the analysis.

market players have a structural interest rate exposure, governed by $\beta \neq 0$, related to their business model that makes their profits sensitive to the direction of policy changes r .¹⁴ Our assumption that $\beta > 0$ suggests that profitability is worsened by unexpected rate hikes. Our interpretation of the market player's loss function lends itself to thinking about these players as financial institutions which the central bank believes may pose financial stability risks. For brevity, we refer to these players as banks.¹⁵

Central Bank's payoff. The first term in the central bank's payoff function is akin to the standard objective function of New Keynesian models (e.g., Galí, 2015). We modify this objective function in the spirit of Woodford (2012) by placing a weight α on financial stability that arises due to policies that surprise banks.¹⁶ For simplicity, we assume the N banks carry equal weight in contributing to this instability, so that the relevant welfare inefficiency is a simple average of the readjustment costs of the banks in I . In Appendix D, we clarify that our preferred interpretation treats financial stability as an additional objective, distinct from the central bank's usual focus on the real economy.¹⁷ This interpretation is consistent with the Fed's behavior during the banking panic of 2023 when it underreacted to inflationary pressures due to concerns about banking losses. Appendix C documents this underreaction using a detailed case study on Silicon Valley Bank and Signature Bank.

Communication. We think of the central bank communicating a message m to ensure banks make well-aligned investments. The underlying assumption is that the central bank has private information about its future target rate, possibly because of private economic data or because market players are uncertain about how policymakers interpret the relevant economic data. Importantly, we assume that the information asymmetry cannot be (entirely) resolved via disclosure of hard information. We discuss this assumption and its relation to the cheap talk problem in our application further in Appendix D.

The next section presents the equilibrium analysis and its main implications for policy, investment, and welfare.

3. Main results

3.1. Central bank underreaction

Our first result is a direct consequence of the CB loss function. Specifically, when the CB cares about market surprises, it underreacts to shocks in the economy-stabilizing policy that occur after communication. By doing so, policymakers reduce market surprises and the associated welfare cost.¹⁸ This simple intuition is captured by Proposition 1, which follows immediately from the game structure and payoff assumption(2).

Proposition 1 (Underreaction). *Let σ denote any strategy profile played in a PBE of the game, and let $r^*(\sigma, \omega_1, \omega_2)$ be the corresponding on-path policy rate chosen by the CB, expressed as a function of the realized shocks (ω_1, ω_2) . On the equilibrium path, the central bank underreacts to the shock ω_2 , that is*

$$0 < \frac{\partial r^*(\sigma, \omega_1, \omega_2)}{\partial \omega_2} < 1.$$

3.2. Perfectly competitive benchmark

We begin by characterizing socially optimal communication, investment, and policy for an arbitrary number N of banks.¹⁹ Because L_{CB} is our measure of social welfare loss, the first best is the strategy profile that a social planner would select if it could directly choose banks' investment rule σ_x , as well as communication and policy. We then ask whether this benchmark can arise in equilibrium. Proposition 2 shows that it can: when banks are atomistic, there exists a fully informative PBE that achieves the first best. Before stating this result formally, it is useful to introduce two additional pieces of notation.

First, for each $\sigma \in \Sigma$, let $W(\sigma) = -\mathbb{E}[L_{CB}|\sigma]$ denote the (ex ante) welfare when the strategy profile is σ . Second, let $\hat{W}$ denote the highest achievable welfare in the game if the central bank (or a social planner) could select any strategy profile from Σ , that is

$$\hat{W} = \max_{\sigma \in \Sigma} W(\sigma) \quad (3)$$

and let $\hat{\Sigma}$ be the argmax of the maximization problem (3).

The following proposition describes what such "first" best strategy profiles look like, and it identifies the case where a profile in $\hat{\Sigma}$ arises in a market equilibrium.

¹⁴ We sometimes refer to r as a rate change. This amounts to assuming $r = 0$ at $t = 0$.

¹⁵ We discuss the empirical relevance of this interpretation in Appendix D.

¹⁶ We use the quadratic loss functional form for tractability as well as consistency with the form of central bank loss functions in the literature.

¹⁷ In Appendix D, we also describe an alternative interpretation, in line with Stein and Sunderam (2018), which grounds financial stability concerns in the pursuit of the traditional mandate.

¹⁸ This feature of the policy decision is in line with the gradualism and underreaction that characterize equilibrium policies in Stein and Sunderam (2018) and Vissing-Jorgensen (2020), who also study monetary policy with costly policy innovations.

¹⁹ In what follows, we interpret N as a measure of the competitiveness of the banking sector. For instance, the central bank may care about the largest banks that cumulatively represent a fixed share of the market, so that N is the number of such banks.

Proposition 2 (First Best and Competitive Implementation). $\hat{W}$ does not depend on N . A strategy profile $\sigma \in \Sigma$ achieves $\hat{W}$ if and only if, under the joint distribution induced by σ , the following conditions hold almost surely:

- (i) Communication is fully informative: every on-path message identifies ω_1 almost surely, namely $\text{Var}(\omega_1 | m, \sigma_m) = 0$.
- (ii) Investment is unbiased, namely $x_i = \mathbb{E}[\omega_1 | m, \sigma_m]$ for every $i \in I$.
- (iii) The central bank balances its objectives on the path of play, namely

$$r = (1 - \alpha)\omega + \alpha\bar{x},$$

where $\bar{x}$ denotes the average investment position of banks.

Additionally, when banks are atomistic,²⁰ there exists a PBE satisfying conditions (i)–(iii) and achieving $\hat{W}$.

The proof is in the appendix. The three properties of Proposition 2 have intuitive interpretations: the central bank communicates all available information ω_1 about its expected future policy, investment positions minimize expected readjustments given the expectations about the economy-stabilizing rate ω which are formed after communication, and the final policy is chosen to optimally balance stabilization of the economy and banks. Moreover, (i), (ii), and (iii) taken together imply that in this type of game, $\hat{W}$ does not depend on N . Therefore, looking at the mismatch between $\hat{W}$ and the ex ante expected equilibrium welfare for different levels of concentration (i.e., different levels of N) is informative.

An implication of Proposition 2 is that when there is no concentration (i.e., $N = \infty$), there is an equilibrium σ^{com} that achieves the maximum welfare $\hat{W}$ attainable by a benevolent planner. The intuition is as follows. First, we show that in every PBE strategy profile (including σ^{com}), the central bank uses the same policy rule σ_r^* . In particular, $\sigma_r^*(\omega_1, \omega_2, x) = (1 - \alpha)\omega + \alpha\bar{x}$, which satisfies property (iii) of Proposition 2. The rule selects a rate between ω and $\bar{x}$, optimally balancing the two stability objectives.

Second, atomistic banks behave as policy takers. While they anticipate that policymakers will take into account the aggregate market position $\bar{x}$ in setting the rate, each bank is too small to influence aggregate outcomes. This implies $\bar{x} = \mathbb{E}[\omega | m, \sigma_m^{com}]$. To see why, note that if the average position of the market were different from the expected real economy-stabilizing policy, that is $\bar{x} \neq \mathbb{E}[\omega | m, \sigma_m^{com}]$, then bank i would expect the central bank to follow σ_r^* and choose an intermediate policy, so that $\mathbb{E}[\sigma_r^{com}(\omega_1, \omega_2, x) | m, \sigma_m^{com}, x] \neq \bar{x}$. However, as policy takers, banks would also find it individually rational to change their position to perfectly match this expected r and thus minimize the surprise loss. Hence, $\sigma_{x_i}^{com}(m) = \mathbb{E}[\omega | m, \sigma_m^{com}]$ must hold in equilibrium, satisfying (ii).

Finally, note that this “unbiased” investment plan is also optimal from the point of view of policymakers, who use communication to guide investors to select $x_i = \omega_1$, the investment position consistent with the expected economy-stabilizing policy. Full guidance to $x_i = \omega_1$ is only achieved with fully transparent communication, by choosing $\sigma_m^{com}(\omega_1) = \omega_1$ for all $\omega_1 \in [-\phi_1, \phi_1]$, or any other communication rule that satisfies (i).

When banks are atomistic, there is no conflict of interest between the central bank and banks at the communication stage, so full revelation can be sustained in equilibrium. Cheap talk nevertheless also admits less informative equilibria. Following Crawford and Sobel (1982), we use ex ante Pareto dominance as a selection criterion: before the central bank observes ω_1 , both the central bank and banks prefer the equilibrium that transmits the most information. In the competitive limit, this criterion selects full revelation, hence an equilibrium satisfying (i)–(iii) of Proposition 2. The characteristics of the most informative equilibrium and its comparative statics with respect to N are discussed in Section 3.4.

Policy and payoffs in the competitive benchmark. The on-path policy rate for any competitive strategy profile σ^{com} , and in general for any $\sigma \in \hat{\Sigma}$, is

$$r^*(\sigma, \omega_1, \omega_2) = \omega_1 + (1 - \alpha)\omega_2 = \omega_1 + \underbrace{\omega_2 - \alpha\omega_2}_{\text{underreaction}}. \quad (4)$$

Transparent communication implies that the central bank fully reacts to policy target shocks occurring before communication. However, consistent with Proposition 1, the welfare-optimal policy exhibits underreaction to state innovations that were not previously communicated in order to limit market surprises.

Simple algebra leads to the following expressions for ex ante welfare and banks’ expected payoffs, for each $\sigma \in \hat{\Sigma}$, including the competitive limit case σ^{com} :

$$W^{com} = W(\sigma) = -\frac{1}{2}\alpha(1 - \alpha)\sigma_2^2 \quad (5)$$

$$EU_i^{com} = \mathbb{E}[u_i | \sigma] = -\frac{1}{2}(1 - \alpha)^2\sigma_2^2 \quad (6)$$

where σ_2^2 is the variance of ω_2 . Both expressions are decreasing in σ_2^2 , proportional to the residual policy uncertainty after communication and the resulting market instability. Additionally, the first-best average welfare is the lowest when the welfare weights of real-sector instability and financial instability are close ($\alpha \approx \frac{1}{2}$). Finally, banks are ex ante better off if the central bank puts a high weight on financial stability.²¹

²⁰ Formally, the atomistic case has a unit mass of banks indexed by $i \in [0, 1]$, with averages replaced by integrals and $\bar{x} = \int_0^1 x_i di$. We use $N = \infty$ as shorthand for this economy.

²¹ We do not model the long-term banking losses due to an unstable economy. In this sense, we might think of our banks as “short-term” oriented.

The results of this section are optimistic. Unfortunately, the atomistic-bank assumption is likely violated in reality, where individual systemic banks are large enough to matter for aggregate financial stability, and importantly, receive special attention from policymakers. In the next section, we explore the general solution for $N < \infty$.

3.3. Oligopolistic competition

In our setting, large systemic banks are non-atomistic: their individual positions affect the aggregate surprise cost entering the central bank's objective and therefore influence equilibrium monetary policy. To see this, note that in all PBEs the central bank plays the optimal policy rule $\sigma_r^*(\omega_1, \omega_2, x) = (1 - \alpha)\omega + \alpha\bar{x}$. As shown in the appendix, such a rule provides the unique minimizer of L_{CB} in any possible subgame reached at $t = 2$. Fixing this rule, for each $i \in I$, $(\omega_1, \omega_2) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2]$ and $x \in X^I$, it indeed holds

$$\frac{\partial \sigma_r^*(\omega_1, \omega_2, x)}{\partial x_i} = \frac{\alpha}{N} > 0$$

which is non-negligible when I is finite. Intuitively, because these systemic banks are non-atomistic, their *individual* losses due to financial assets that perform badly under a candidate policy rate matter for aggregate financial stability and consequently for the policy decision.²² Hence a change in the individual position x_i affects the choice of r . We refer to $\frac{\alpha}{N}$ as the *policy influence* of an individual bank.

To better understand how banks' direct policy influence changes the strategic behavior of bank i , let us focus on the simple case in which the central bank has revealed ω_1 at $t = 0$ and $x_j = \omega_1$ for each $j \neq i$. The expected utility of bank i at $t = 1$ given σ_r^* and $\bar{x}_{-i} = \omega_1$ is

$$\mathbb{E}[u_i(x_i, r) | \omega_1, \bar{x}_{-i}, \sigma_r^*] = \underbrace{-\frac{1}{2}(1 - \alpha)^2 \sigma_2^2 - \beta \omega_1}_{\text{expected payoff from unbiased investment}} - \underbrace{\frac{1}{2} \left(1 - \frac{\alpha}{N}\right)^2 (x_i - \omega_1)^2}_{\text{expected readjustment due to investment bias}} - \underbrace{\beta \left[\frac{\alpha}{N}(x_i - \omega_1)\right]}_{\text{policy distortion due to investment bias}}. \quad (7)$$

The first term is the expected payoff under unbiased investment; the remaining two capture the trade-off from changing x_i : a larger portfolio mismatch, but a policy rate pulled toward the lower rate preferred by the bank. We immediately notice two differences relative to the perfectly competitive case ($\frac{\alpha}{N} \approx 0$). First, the expected readjustment due to investment bias is smaller, since the policy influence of i makes the central bank choose a policy rate closer to x_i than if the bank were atomistic and therefore had no policy influence. Second, because of the positive interest rate exposure $\beta > 0$, the investment bias leads to a policy distortion that matters in addition to its effect on the performance of portfolio x_i . This second effect is key for the next result, which shows how biased investment can be used strategically in equilibrium to distort policies in the desired direction.

Proposition 3 (Oligopoly Distortions). *The oligopolistic equilibrium investment plan σ_x^{oli} entails an investment bias. Formally, fix any communication rule σ_m and let σ_x^{oli} be such that each bank best responds to (σ_m, σ_r^*) and to all other banks' strategies. Then,*

- (i) *Investment decisions are consistent with a lower interest rate than the one expected to stabilize the economy. The bias magnitude is strictly increasing in α and β and strictly decreasing in N . For each $i \in I$ and on-path message m ,*

$$\sigma_{x_i}^{oli}(m) - \mathbb{E}[\omega_1 | m, \sigma_m] = -\frac{\alpha\beta}{(1 - \alpha)(N - \alpha)}.$$

- (ii) *If the central bank fully reveals ω_1 , that is, $\sigma_m = \sigma_m^{com}$, then the on-path policy decision exhibits a downward distortion. For each $\omega_1 \in [-\phi_1, \phi_1]$ and $\omega_2 \in [-\phi_2, \phi_2]$,*

$$r^*(\sigma^{oli, tr}, \omega_1, \omega_2) = r^*(\sigma^{com}, \omega_1, \omega_2) - \frac{\alpha^2 \beta}{(1 - \alpha)(N - \alpha)}$$

where $\sigma^{oli, tr} = (\sigma_m^{com}, \sigma_x^{oli}, \sigma_r^*)$, with σ_m^{com} fully informative about ω_1 .

To see why the investment bias arises in equilibrium, consider again the case where the central bank always fully reveals ω_1 and in which all banks other than i have chosen unbiased positions $x_j = \omega_1$. The first-order condition in Eq. (7) is

$$\underbrace{\left(1 - \frac{\alpha}{N}\right)^2 (\omega_1 - x_i)}_{\text{marginal cost from downward bias}} = \underbrace{\frac{\alpha\beta}{N}}_{\text{marginal benefit from downward bias}}$$

which yields the bias of Proposition 3 if $N = 1$.²³ Given that the marginal benefit from a downward investment bias is positive for I finite, $x_i = \omega_1$ is not optimal for an individual bank, even when all other banks are choosing unbiased investment. This effect is amplified when the individual policy influence $\frac{\alpha}{N}$ is larger and when the benefit from a rate cut β is greater. Thus, in equilibrium,

²² This argument is analogous to that in the granularity literature where idiosyncratic firm-level shocks influence aggregate movements given the firm size distribution (see e.g., Gabaix, 2011).

²³ Recall that for the sake of developing intuition we assumed $x_{-i} = \omega_1$. This will not be true in equilibrium, so that the expression of the equilibrium bias for a generic N , reported in Proposition 3, is different from the one we obtain from this special case.

each systemic bank chooses portfolio positions consistent with smaller rate hikes than those really expected, potentially increasing their own losses from a restrictive change in monetary policy. This result, apparently surprising, is consistent with the evidence by [Begenau et al. \(2025\)](#), who find that the largest US banks typically hold net derivative positions reinforcing (instead of hedging) their interest rate exposure.

A consequence of these distortions is that when the central bank commits to full information transmission at $t = 0$ and banks best respond, society incurs an ex ante welfare loss relative to perfect competition. In particular,

$$W(\sigma^{oli,tr}) - W(\sigma^{com}) = -\frac{\alpha^3 \beta^2}{2(N - \alpha)^2(1 - \alpha)} < 0.$$

Banks benefit from their policy influence, achieving a higher ex ante payoff than in the perfectly competitive case. Indeed, for each $i \in I$,

$$EU_i^{oli,tr} - EU_i^{com} = \frac{(2N - 1 - \alpha)\alpha^2 \beta^2}{2(N - \alpha)^2(1 - \alpha)} > 0.$$

As one would expect, both these gaps close if the degree of competition (N) increases, if the welfare weight on financial stability (α) decreases, or if the structural interest rate exposure of systemic banks (β) decreases.

The analysis has shown that when the market is not perfectly competitive, investment is generally biased. Moreover, when the central bank communicates all available information about the economy-stabilizing policy, there is, on average, a welfare loss for society. Can more communication flexibility close this gap relative to perfect competition? Or will it instead lead to further welfare losses? We address this question in the next section.

3.4. Communication

We now relax any assumption about commitment to transparent communication and examine equilibrium strategic communication in our simple model. An unfortunate implication of [Propositions 2](#) and [3](#) is that communication cannot bring us back to the ex ante optimum $W(\sigma^{com})$. In fact, in any oligopolistic PBE, investment is biased downwards relative to $\mathbb{E}[\omega_1 | m, \sigma_m^{oli}]$, the market expectation of the future economy-stabilizing rate given any equilibrium communication rule σ_m^{oli} and the announcement m made at $t = 0$. This means that the unbiasedness requirement of [Proposition 2](#) is violated so that the oligopolistic equilibrium $\sigma^{oli} \notin \hat{\Sigma}$. Importantly, given the cheap talk nature of communication without commitment in this environment, flexibility ends up creating an additional welfare loss relative to commitment to transparent communication.

In this environment, equilibrium communication has the exact same characteristics as the communication strategy of the seminal cheap talk game studied by [Crawford and Sobel \(1982\)](#), when the preference misalignment between players is set to $b = \frac{\alpha\beta}{2\phi_1(N - \alpha)(1 - \alpha)}$.

Proposition 4 (Cheap Talk). *Any communication rule played in an oligopolistic PBE essentially partitions $[-\phi_1, \phi_1]$ into a finite number of intervals, so that equilibrium communication is never fully informative. In addition:*

- (i) A PBE strategy profile $\sigma_P^{oli} = (\sigma_{m,P}^{oli}, \sigma_x^{oli}, \sigma_r^*)$ corresponding to P partition elements, produces the following ex ante welfare:

$$W(\sigma_P^{oli}) = W(\sigma^{oli,tr}) - \frac{1}{2} \alpha(1 - \alpha) \hat{\sigma}_{1,P}^2$$

where $\hat{\sigma}_{1,P}^2 = \mathbb{E}[Var(\omega_1 | m, \sigma_{m,P}^{oli}) | \sigma_{m,P}^{oli}]$ is the residual variance of ω_1 induced by the equilibrium communication strategy.

- (ii) The residual variance $\hat{\sigma}_{1,P}^2$ is increasing in the absolute value of the investment bias, and decreasing in the number of partition elements P . The maximum number of partition elements $\bar{P}$ is the smallest integer greater than or equal to

$$\frac{1}{2} \left(\sqrt{1 + \frac{4\phi_1(N - \alpha)(1 - \alpha)}{\alpha\beta}} - 1 \right)$$

which is non-increasing in the absolute value of the investment bias.

As shown in the appendix, [Proposition 4](#) follows immediately from [Crawford and Sobel \(1982\)](#) and from the game structure once we impose individual rationality of investment and policy decisions.

Fully informative communication about policy intentions would be beneficial for the central bank in this context, since welfare is decreasing in the residual uncertainty about ω_1 , as described in point (i). However, a fully informative equilibrium cannot be sustained. To see this, imagine that the central bank was playing $\sigma_m(\omega_1) = \omega_1$ in a PBE, which is fully informative. Banks would best respond with an investment plan $\sigma_x^{oli}(m) = m - \frac{\alpha\beta}{(N - \alpha)(1 - \alpha)}$ to bias the future rate downwards. But the central bank would then be tempted to announce higher rate hikes at $t = 0$, for instance by deviating to $\sigma_m(\omega_1) = \omega_1 + \frac{\alpha\beta}{(N - \alpha)(1 - \alpha)}$ whenever such a message belongs to the message space. This credibility loss arises because the interests of the central bank and banks are not fully aligned and banks with policy influence no longer behave as policy takers.

The main consequence of the loss in information transmission relative to commitment to informative communication is that the central bank will provide less guidance and banks will be less ready for changes in policies. Therefore, policies will create more instability. These negative effects are minimized in the most informative equilibrium (i.e., when $P = \bar{P}$) as more partitions imply

more informative communication. Not surprisingly, point (ii) of [Proposition 4](#) shows that the informativeness of communication in this society-preferred equilibrium increases the closer the investment distortion is to zero. Hence, the maximum amount of guidance achieved in equilibrium by the central bank increases in the competitiveness of the banking sector and decreases in the welfare weight of market surprises and in the exposures of banks.

Note that banks too are ex ante worse off when the central bank does not communicate clearly. When the equilibrium strategy profile is σ_P^{oli} , banks' loss relative to commitment to full information is

$$EU_i^{oli,tr} - EU_i^{oli,P} = \frac{1}{2}(1 - \alpha)^2 \hat{\sigma}_{1,P}^2 > 0. \quad (8)$$

Intuitively, being less able to predict future policies, banks find it harder to balance the need to follow the central bank while also attempting to influence it.

Finally, notice that for any equilibrium with $P \leq \bar{P}$, the on-path policy distortion will be negative on average,

$$\mathbb{E}[r^*(\sigma_P^{oli}, \omega_1, \omega_2) - r^*(\sigma^{com}, \omega_1, \omega_2)] = -\frac{\alpha^2 \beta}{(N - \alpha)(1 - \alpha)}.$$

Such a distortion can still be positive for some realizations of ω_1 (e.g., those very close to the lower bound of the respective partition element), due to the partial information transmission of the cheap talk equilibrium.

One of the key insights of this section is that the central bank faces a trade-off between avoiding surprises *ex post* and communicating effectively: on the one hand, financial stability concerns ($\alpha > 0$) push policymakers to systematically underreact to state innovations ω_2 that could surprise banks. Such ex post policy adjustments can be desirable for the reasons discussed in our first-best analysis: market surprises generate welfare-detrimental instability. On the other hand, the greater the stability concerns α , the larger the investment bias and the communication loss: the central bank's ex post effort to avoid policy surprises reduces the effectiveness of forward guidance, a force that increases instability.

A remark on transparent communication. Both the central bank and banks are better off, on average, when the central bank commits to transparent communication than when communication is cheap talk. Moreover, no communication rule could do better than full transparency. To see why, consider the strategy profile $\tilde{\sigma} = (\tilde{\sigma}_m, \tilde{\sigma}_x, \tilde{\sigma}_r)$, with $\tilde{\sigma}_m : [-\phi_1, \phi_1] \rightarrow \Delta(M)$, $\tilde{\sigma}_{x_i}(m) = \mathbb{E}[\omega_1 | m, \tilde{\sigma}_m] - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$ for each $i \in I$ and $m \in \text{supp}(\tilde{\sigma}_m)$, and $\tilde{\sigma}_r = \sigma_r^*$. The policy rule σ_r^* solves the central bank loss minimization problem at time $t = 2$, and investment plans maximize banks' payoffs given the policy and communication rules. Following steps analogous to the proof of [Proposition 4.i](#) yields

$$W(\tilde{\sigma}) = W(\sigma^{oli,tr}) - \frac{1}{2}\alpha(1 - \alpha)\mathbb{E}[Var(\omega_1 | m, \tilde{\sigma}_m) | \tilde{\sigma}_m],$$

where $Var(\omega_1 | m, \tilde{\sigma}_m)$ is the residual variance after message realization m when the conditional message distribution is described by $\tilde{\sigma}_m$. First, note that $W(\tilde{\sigma}) \leq W(\sigma^{oli,tr})$, that is, no communication rule can improve on the full-transparency expected welfare. Second, $W(\tilde{\sigma}) = W(\sigma^{oli,tr})$ only if $\mathbb{E}[Var(\omega_1 | m, \tilde{\sigma}_m) | \tilde{\sigma}_m] = 0$, that is, only if communication is fully transparent and leaves no uncertainty about ω_1 . This result captures a central intuition of the paper: any attempt by the central bank to counter market distortions through strategic communication is self-defeating.

3.5. Towards a less financial stability oriented central bank

We have shown that systemic banks, because they are non-atomistic, can influence the policies of the central bank through their portfolio choices, and that this influence creates an inefficiency and a welfare loss, especially when the central bank has private information that cannot be communicated credibly. What remedies could address these issues?

Recall that our measure of policy influence, $\frac{\alpha}{N}$, captures the direct effect of an individual bank's behavior on the central bank's choices. A positive policy influence makes the oligopoly different from perfect competition, and drives the inefficiency due to investment and policy distortions. Therefore, one obvious remedy would be to reduce the scale of systemically important banks. However, such a remedy would likely be a long and difficult process and may itself create inefficiencies due to economies of scale (not modeled here). Moreover, as mentioned earlier, concentration appears to be increasing, rather than decreasing over time ([Corbae and D'Erasmo \(2021\)](#)). An alternative solution would be to reduce α (i.e., the weight placed on financial stability). While it is not possible to reduce the welfare weight of a market disruption, we can consider what would happen if a (representative) central banker with an objective function different from societal welfare was appointed. In particular, in the spirit of [Rogoff \(1985\)](#), consider the following thought experiment: imagine that society could appoint a central banker who has an objective with the same functional form as in payoff specification (2), but with a relative weight $\tilde{\alpha} \in [0, 1]$ on financial stability potentially different from α . What would the optimal central banker look like?

Formally, let us denote by $\sigma_P^{oli}(\tilde{\alpha})$ and $\sigma_P^{oli}(\tilde{\alpha})$, respectively, the oligopolistic strategy profile under transparency and the cheap talk oligopolistic equilibrium when the weight in the loss function of the central bank is $\tilde{\alpha}$.²⁴ For $\tilde{\alpha} > 0$, let $\bar{P}(\tilde{\alpha})$ be the smallest integer greater than or equal to $\frac{1}{2} \left\{ \left[1 + 4\phi_1(N - \tilde{\alpha})(1 - \tilde{\alpha})(\tilde{\alpha}\beta)^{-1} \right]^{1/2} - 1 \right\}$. For each fixed $\tilde{\alpha}$, we apply the same ex ante Pareto selection used above: a finer feasible partition lowers residual uncertainty and weakly raises the ex ante payoffs of both the central banker and

²⁴ When $\tilde{\alpha} = 0$, let $\sigma_P^{oli}(0)$ denote a fully revealing equilibrium. The formula for $\bar{P}(\tilde{\alpha})$ applies only when $\tilde{\alpha} > 0$.

banks, strictly raising both when $\tilde{\alpha} > 0$.²⁵ We ask which $\tilde{\alpha}$ maximizes $W(\sigma_{\tilde{P}}^{oli, tr}(\tilde{\alpha}))$ or $W(\sigma_{\tilde{P}}^{oli}(\tilde{\alpha}))$, where the function W , as in the previous section, maps strategy profiles into ex ante welfare levels (computed based on the welfare weight α).

We denote a central banker with $\tilde{\alpha} < \alpha$ as *financial-stability-agnostic* (hereafter FS-agnostic), with $\tilde{\alpha} > \alpha$ as *financial-stability-oriented* (hereafter FS-oriented), and with $\tilde{\alpha} = \alpha$ as unbiased. An FS-agnostic central banker focuses on the real economy, paying less attention to banks, while an FS-oriented central banker is more accommodating towards banks. Our main result is that it would be optimal for society to appoint an FS-agnostic central banker.

Proposition 5 (Optimal Central Banker). *For any central banker weight $\tilde{\alpha} \in [0, 1]$, assume that when communication is cheap talk the most informative equilibrium $\sigma_{\tilde{P}}^{oli}(\tilde{\alpha})$ is selected, so that the equilibrium communication strategy is fully revealing when $\tilde{\alpha} = 0$ and partitions $[-\phi_1, \phi_1]$ into $\tilde{P}(\tilde{\alpha})$ elements when $\tilde{\alpha} > 0$. Then, the optimal central banker is FS-agnostic. In particular, the following holds:*

- (i) *There exist thresholds $\sigma_2^{oli}, \sigma_2^{oli, tr} > 0$ on the non-communicable uncertainty σ_2 such that: under transparency, maximally FS-agnostic central bankers ($\tilde{\alpha} = 0$) improve welfare over unbiased ones if and only if $\sigma_2 < \sigma_2^{oli, tr}$; under cheap talk, they improve over unbiased ones if and only if $\sigma_2 < \sigma_2^{oli}$. Moreover, $\sigma_2^{oli} > \sigma_2^{oli, tr}$.*
- (ii) *The optimal central bankers are such that $\tilde{\alpha}^{oli, tr} \in (0, \alpha)$ and $\tilde{\alpha}^{oli} \in [0, \alpha)$, where $\tilde{\alpha}^{oli, tr}$ is the central banker under transparency and $\tilde{\alpha}^{oli}$ is the central banker under cheap talk. Welfare under transparency and $\tilde{\alpha}^{oli, tr}$ is higher than under cheap talk and $\tilde{\alpha}^{oli}$. Additionally, $\tilde{\alpha}^{oli, tr}$ is increasing in α , N , and ϕ_2 .*
- (iii) *If the central bank commits to transparency, the market ex ante payoff is always strictly larger when the central banker is unbiased than when she is FS-agnostic. If instead communication is cheap talk, then the market ex ante payoff can be larger under a FS-agnostic central banker than under an unbiased one.*

The intuition for why society might prefer a FS-agnostic central banker is simple. As discussed above, a FS-agnostic central banker reduces systemic banks' influence over policies, and is more welfare-enhancing the greater the policy influence of each systemic bank. If the choice is between an unbiased central banker and one that assigns no weight to financial stability, part (i) tells us that the latter will be desirable when state innovations occurring after communication are of limited magnitude (σ_2^2 low), especially when an FS-agnostic central banker can also improve information transmission (i.e., under cheap talk). If $\tilde{\alpha}$ can be fine-tuned, the optimal central banker is always FS-agnostic, but typically not maximally so: market surprises are costly for society ($\alpha > 0$), so that some degree of market stabilization tends to be optimal. This intuition is captured by part (ii). While the proposition focuses on the most informative equilibria—those partitioning $[-\phi_1, \phi_1]$ into $\tilde{P}(\tilde{\alpha})$ elements—the optimality of a FS-agnostic central banker holds more generally. As we show in [Appendix B.3](#), an analogous result holds if we fix P at any level feasible in equilibrium when $\tilde{\alpha} = \alpha$ and compare the welfare associated with the equilibrium in P partition elements as we let $\tilde{\alpha}$ vary.

A perhaps surprising result of the proposition is part (iii), which states that under cheap talk, an FS-agnostic central banker can paradoxically improve the performance of banks. The key intuition is that when misalignment of incentives compromises information transmission (in the cheap talk equilibrium), a more FS-agnostic central banker improves communication. Specifically, an FS-agnostic central banker reduces the marginal net benefit of a downward investment bias, making it optimal for banks to choose a portfolio position x_i more aligned with their true expectations about ω . Consistent with [Proposition 4.ii](#), this bias reduction means that more information is communicated in equilibrium, and the resulting reduction in policy uncertainty has a stabilizing effect, making *both* society and banks better off. If this gain in information transmission is large enough, banks are willing to give up some of their policy influence to achieve it (as is the case when an FS-agnostic central banker is appointed). The following example provides a stark illustration of this phenomenon.

Babbling monopoly. Let $N = 1$ and $\phi_1 = \frac{1}{2}$. In this monopolistic setting $\frac{\alpha\beta}{(1-\alpha)^2} > \frac{1}{4}$ implies that $\tilde{P} = 1$, so that $\hat{\sigma}_{1, \tilde{P}}^2 = \frac{1}{12}$ in the cheap talk equilibrium. Under an unbiased central banker, the monopolistic bank's ex ante value of the game is $EU^\alpha = -\frac{1}{2}(1 - \alpha)^2 \left[\frac{1}{12} + \sigma_2^2 + \frac{(\alpha\beta)^2}{(1-\alpha)^4} \right] + \left(\frac{\alpha\beta}{1-\alpha} \right)^2$, which simplifies to $-\frac{1}{2}(1 - \alpha)^2 \left[\frac{1}{12} + \sigma_2^2 - \frac{(\alpha\beta)^2}{(1-\alpha)^4} \right]$. With a maximally FS-agnostic central banker, i.e., $\tilde{\alpha} = 0$, the bank's ex ante value of the game is $EU^0 = -\frac{1}{2}\sigma_2^2$ in the fully informative equilibrium. If $\frac{1}{12} - \frac{(\alpha\beta)^2}{(1-\alpha)^4} > \frac{1-(1-\alpha)^2}{(1-\alpha)^2}\sigma_2^2$ then $EU^0 > EU^\alpha$, and our monopolistic bank would give up all its policy influence to improve communication. Note that the condition is satisfied for $\frac{\alpha\beta}{(1-\alpha)^2} \in (\frac{1}{4}, \frac{1}{\sqrt{12}})$ and σ_2^2 small enough.

We have shown that an FS-agnostic central bank improves welfare and potentially makes systemic banks better off as the central bank becomes more predictable. An FS-agnostic central bank, however, does not get us back to the first-best level of welfare, as it underweights policy surprises relative to the $N = \infty$ case. These welfare gains raise a further question: even if such policies improve outcomes, can a central bank credibly commit to being FS-agnostic in practice? In our context, even if a commitment to a formal rule is not feasible or credible, repeated interactions can allow an unbiased central banker to build an FS-agnostic reputation. In other words, the gain from future market discipline can encourage an unbiased central bank to choose FS-agnostic policies in the present, provided that (i) banks observe information about ω at the end of each repetition and (ii) the central bank values future welfare sufficiently. In the next section, we show that repeated interactions open up different possibilities, even when there is no feedback about ω . In some cases, FS-agnostic monetary policies may not be necessary.

²⁵ A reduction $\Delta > 0$ in residual variance raises the central banker's expected payoff by $\frac{1}{2}\tilde{\alpha}(1 - \tilde{\alpha})\Delta$ and each bank's expected payoff by $\frac{1}{2}(1 - \tilde{\alpha})^2\Delta$. At $\tilde{\alpha} = 0$, the central banker is indifferent, whereas banks strictly prefer full revelation.

4. Game repetition

Assuming that $\tilde{\alpha} = \alpha$, we now consider the implications of infinite-horizon repetition. Let $\tau = 0, 1, \dots$ index the stage game repetition, and let $\omega^\tau, m^\tau, x^\tau, r^\tau$ denote the respective quantities in the stage game τ .

In our discussion of repeated interactions, we focus on PBEs such that any investment decision x_i^τ is contingent only on past messages $(m^0, \dots, m^\tau)$ and investment decisions $(x_i^0, \dots, x_i^{\tau-1})_{i \in I}$,²⁶ and therefore does not condition directly on past central bank rate choices or shock realizations. We find this selection criterion particularly meaningful because the corresponding equilibria exist regardless of whether ω_1^τ and ω_2^τ are observed by market players at the end of the stage game. Additionally, for simplicity, we focus on equilibria such that $r^\tau = (1 - \alpha)\omega^\tau + \alpha\bar{x}^\tau$, so that the central bank still follows the same rate rule described in the previous sections.²⁷

We address the following question: can an infinitely repeated interaction allow the central banker to better discipline markets, communicate more transparently, and achieve the first best? As we will see in the next propositions, repeated interaction can indeed achieve the first best on the path of play with a very simple type of PBE strategy profile, provided that forward guidance is valued enough by markets (i.e., ϕ_1 is high). At the same time, repeated interaction can facilitate collusion between large banks, sometimes with perverse consequences on welfare.

Proposition 6 (CB Disciplines Markets). *Consider the infinitely repeated game with discount factor $\delta \in (0, 1)$. If $\phi_1 > \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}\sqrt{3\frac{2N-\alpha-1}{1-\alpha}}$ then there exists a threshold $\delta^* \in (0, 1)$ such that if $\delta \geq \delta^*$ the ex ante stage welfare is at its first-best level on the equilibrium path of some PBEs. Such PBE strategy profiles are qualitatively equivalent to the following simple profile.*

- (i) At $\tau = 0$, the CB selects $m^\tau = \omega_1^\tau$. At $\tau > 0$, the CB selects $m^\tau = \omega_1^\tau$ if $x_i^s = m^s$ for each $s < \tau$ and each $i \in I$, while the CB draws $m^\tau \sim U(-\phi_1, \phi_1)$ after any other history.
- (ii) At $\tau = 0$, each $i \in I$ selects $x_i^\tau = m^\tau$. At $\tau > 0$ each $i \in I$ selects $x_i^\tau = m^\tau$ if $x_j^s = m^s$ for each $s < \tau$ and each $j \in I$, while each $i \in I$ selects $x_i^\tau = -\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$ after any other history.
- (iii) At each $\tau = 0, 1, \dots$, the CB sets rate $r^\tau = (1 - \alpha)\omega^\tau + \alpha\bar{x}^\tau$.

The proposition has a very simple interpretation. If the variance of the shock ω_1 is large enough (ϕ_1 high), the value of forward guidance is high for banks. The central bank can leverage its informational advantage by conditioning fully informative guidance on market discipline and threatening to revert to meaningless communication if the directives are not followed. The threat is credible because when banks expect communication to be uninformative and stop listening, no deviation from the central bank can restore confidence in its announcements. In this equilibrium, policymakers play an active role in disciplining large banks. After banks' deviations, the central bank does not simply revert to the most efficient stage PBE communication rule presented in Section 3.4: market discipline follows from a threat to revert to maximally inefficient communication. This threat allows for market discipline, even in cases where banks would be better off under repetition of the stage game equilibrium σ_P^{oli} .

A natural question is whether the first best can also be sustained by spontaneous coordination of banks, without policymakers threatening a reversal to uninformative communication. To address this question, we focus on PBE profiles σ satisfying the following additional selection criteria.

First, in every repetition τ of the stage game, each bank i chooses an investment plan of the type $x_i^\tau(h^{\tau-1}, m^\tau) = \mathbb{E}_\sigma[\omega_1^\tau | m^\tau] - b(h^{\tau-1})$, where $h^\tau = (m^s, x^s)_{s=0}^\tau$ is the corresponding (on-path or off-path) history of past communication and investment decisions. Second, in every τ , the central bank plays the stage-game communication rule of the type $\sigma_m^\tau(\cdot | \omega_1^\tau, b(h^{\tau-1})) \in \Delta(M)$ minimizing its stage-game expected loss function and partitioning $[-\phi_1, \phi_1]$ into the maximal number of elements $\bar{P}$ given banks' stage game investment bias $b(h^{\tau-1})$.²⁸ Third, in every stage τ , the ex-ante expected stage game payoff of each bank $i \in I$ on the equilibrium path is weakly greater than the expected payoff in the equilibrium σ_P^{oli} of the (non-repeated) stage game.

We call these equilibria *collusive*, as they represent a Pareto improvement for large banks relative to the one-stage setting of the previous sections, and such improvement comes from banks' coordination effects on investment bias $b(h^{\tau-1})$. Unlike the case of the previous proposition, the central bank plays a passive role: it always communicates as transparently as possible, given the investment bias of banks.²⁹ The next proposition suggests that under certain conditions, the repeated game has collusive PBEs where the oligopolists coordinate on the competitive equilibrium. At the same time, repeated interaction can also lead to collusive PBEs that are detrimental to financial stability as well as to the real economy. To fix notation, for each $n = 1, \dots, N$, and fixed parameters $\alpha, \beta, \phi_1, \phi_2$, let $\hat{\sigma}_{1,\bar{P}}^2(n)$ denote the residual variance of ω_1 after communication in the most informative stage-game equilibrium when the market consists of n large banks.

Proposition 7 (Market Collusion). *Consider the infinitely repeated game with discount factor $\delta \in (0, 1)$. The following holds:*

- (i) If $\hat{\sigma}_{1,\bar{P}}^2(N) > \left[\frac{\alpha\beta}{(N-\alpha)(1-\alpha)} \right]^2 \frac{2N-\alpha-1}{1-\alpha}$, then there exists $\delta_1^* \in (0, 1)$ such that, if $\delta \geq \delta_1^*$, the first-best ex ante stage welfare is the ex ante stage welfare of some collusive PBE of the infinitely repeated game.

²⁶ In the repeated-game extension, each period's investment profile is publicly observed before the next period begins.

²⁷ In these equilibria, following such a rate rule is indeed optimal for the central bank.

²⁸ In this context, $\bar{P}$ is the smallest integer greater than or equal to $\frac{1}{2} \left\{ \left[1 + 4\phi_1 b(h^{\tau-1})^{-1} \right]^{1/2} - 1 \right\}$. When $b(h^{\tau-1}) = 0$, the corresponding most informative communication rule is fully revealing.

²⁹ If collusive equilibria exist, they are a subset of the PBEs of the dynamic game (given other players' strategies, no player has a strict incentive to deviate to other collusive or non-collusive strategies).

(ii) If $\hat{\sigma}_{1,\bar{p}}^2(1) - \hat{\sigma}_{1,\bar{p}}^2(N) < \left[\frac{\alpha\beta}{(1-\alpha)^2} \right]^2 - \left[\frac{\alpha\beta}{(N-\alpha)(1-\alpha)} \right]^2 \frac{2N-\alpha-1}{1-\alpha}$, then there exists $\delta_2^* \in (0, 1)$ such that, if $\delta \geq \delta_2^*$, the ex ante welfare of the monopolistic static game is the ex ante stage welfare of some collusive PBE of the infinitely repeated game.

Banks are ex ante better off in the monopolistic equilibrium (ii) relative to the first best (i) if $\hat{\sigma}_{1,\bar{p}}^2(1) < \left[\frac{\alpha\beta}{(1-\alpha)^2} \right]^2$, while they are better off at the first best if $\hat{\sigma}_{1,\bar{p}}^2(1) > \left[\frac{\alpha\beta}{(1-\alpha)^2} \right]^2$.

The first part of the proposition suggests that if the policy uncertainty that can be resolved by forward guidance is sufficiently large (i.e., payoff-relevant for banks), then banks will benefit from self-discipline. Specifically, banks refrain from using their policy influence to bias future policies, which allows the central bank to communicate all its information on future policy changes, achieving the first best.

The second part of the proposition serves as a caveat. Collusion between banks might instead increase their collective policy influence, allowing them to act as a single large player.

It is possible that large banks might try to coordinate on an investment strategy that maximizes their aggregate expected payoffs. The type of collusion that is optimal for large banks need not be efficient from the central bank's perspective, nor need it be as inefficient as the monopolistic one. However, the two equilibria highlighted in Proposition 7 are useful focal points from the point of view of the analyst and, potentially, the policymaker. The end of Proposition 7 provides a sufficient condition for the monopolistic collusive equilibrium to benefit large banks more than the first best. If the gains from monopolistic policy influence are large, banks should not be expected to self-coordinate on the first best, even when the first best is among the feasible collusive PBE.

Overall, the (partial) analysis of the repeated game suggests that dynamic incentives could lead to either higher or lower central bank losses. To maximize the chances of achieving efficient outcomes, policymakers might have to play an active role by threatening to withhold future guidance if large banks refuse to cooperate.³⁰

5. Discussion and policy implications

We have shown that when the central bank cares about financial stability and banks are asymmetrically affected by rate hikes and rate cuts, systemic banks can use their policy influence to make adverse rate changes less likely. Specifically, they can choose investment positions that would entail costly readjustment if the central bank were to choose a high policy rate. This policy distortion depends on banks' policy influence $\frac{\alpha}{N}$ which decreases with the degree of competition N and increases with the weight the central bank places on financial stability α . The resulting welfare loss is larger if the central bank cannot guide banks credibly, which is reasonable when misaligned objectives lead banks to distrust policy announcements. Increasing competition in the financial sector (i.e., a larger number of smaller and less systemic banks) and appointing a central banker who puts less weight on financial stability might increase both information transmission and welfare in equilibrium. Repeated interaction is not guaranteed to resolve the conflict of interest: the central bank can try to discipline banks under the threat of withholding future guidance if banks exercise their policy influence. But, when forward guidance is not valuable enough, market discipline might not be feasible, and large systemic banks could instead increase their collective policy influence through collusion.

Our analysis yields a number of predictions. It suggests that a central bank that cares about large banks is expected to tailor policies to the interest rate exposures of those banks, with this tailoring increasing in the concentration of the sector. As a consequence, net interest rate exposures of the largest systemic banks should systematically predict the variation in interest rate choices not explained by economic fundamentals. Moreover, the model would also predict that the rising banking concentration over time (see e.g., Corbae and D'Erasmus, 2021), by increasing each systemic bank's policy influence, should result in greater policy distortions by the central bank and, all else equal, a declining trend in the interest rate. Indeed, while not a causal claim, rising U.S. bank concentration since the 1980s has coincided with a steadily declining interest rate. While the effectiveness and informativeness of forward guidance are in general expected to be limited when individual systemic banks have substantial policy influence, market discipline and transparent communication should be more common when the uncertainty regarding economic fundamentals is sufficiently high.

While we do not explicitly model alternative policy instruments, one might expect that financial regulations could potentially help control banks' incentives to influence the central bank: deposit insurance and bailouts on the one hand, and specific requirements on investments x_i for systemic banks on the other hand. The former two practices are common, though each comes with its own trade-offs (e.g., they raise the issue of moral hazard). Nonetheless, such regulations could relieve monetary policy committees from the task of addressing financial stability with interest rate decisions. This independence from financial stability concerns might allow them to better achieve their traditional mandate and is consistent with the Tinbergen rule (Tinbergen, 1952). Constraints and requirements on x_i can also reduce the inefficiencies addressed in our model, especially if, like stress tests, they dynamically adapt to changing economic conditions. It seems optimistic, however, to expect capital requirements and stress tests to amount to regulators choosing x_i , as banks typically have some room to choose their risk models and liability structures strategically (Sundaresan and Wang, 2023). Thus, reducing policy influence by intervening on the central bank's side may provide a valuable complementary strategy. A possible

³⁰ Extending the analysis to a broader set of equilibria, including those where market strategies are contingent on previous monetary policy, would likely yield additional insights. Large banks, not only policymakers, could strategically exploit communication to further bias policy or increase the informativeness of announcements. For instance, they could threaten to ignore the central bank's future forward guidance if previous announcements are not precise enough or previous rate changes are not accommodating enough.

avenue of future research could be extending the current model to explicitly study the interaction between communication, central bank design, and alternative financial stability policies.

The simplicity of our model also makes it tractable for studying a number of different extensions which we leave for future research. Within the context of central banking, future work could study banks with heterogeneous policy influence or rate preferences, and ask how these types of heterogeneity affect financial stability and communication. Letting banks have private information and communicate privately with each other and the central bank could also deliver additional insights and policy implications. A closer look at the strategic complementarities in banks' positions, which are present in this paper but not its focus, would be another avenue for further investigation that could be particularly relevant in more complex settings. Moreover, extending the model to account for the effect of uncertainty, (mis)communication, and surprises on risk premia (e.g., Cieslak and McMahon, 2023), and the effects of the latter on economic stabilization, could highlight additional forces to consider when evaluating the optimal weight the central bank assigns to financial stability.

Our model may also generate insights relevant to time-consistency problems in policy areas beyond central banking. Such problems arise whenever policymakers face large, strategic actors capable of imposing meaningful economic costs through their actions. In these settings, policymakers anticipating such actions may need to deviate from otherwise optimal policies—analogueous to the underreaction and investment bias in our central bank setting. For example, a policymaker concerned with both stable energy output and environmental sustainability may wish to tax fossil fuel use for environmental reasons. However, the energy sector is typically concentrated, and large strategic energy companies could increase reliance on fossil fuels *ex ante* to heighten their exposure to potential losses from such taxes. If these firms become sufficiently exposed, the policymaker might refrain from enacting the tax in the first place to preserve the stability of the sector. Moreover, communication may also play a similar role as in our central bank setup, where it helps prevent policy surprises. Policymakers often engage with industry through formal consultations, policy white papers, and industry roundtables, partly in order to guide market expectations. These parallels indicate that similar mechanisms may operate in other strategic policy settings, which future research could investigate.

CRediT authorship contribution statement

Giampaolo Bonomi: Writing – review & editing, Writing – original draft, Conceptualization; **Ali Uppal:** Writing – review & editing, Writing – original draft, Conceptualization.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used ChatGPT (OpenAI) to assist in improving the language, clarity, and organization of portions of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank the editor (Marco Ottaviani), the associate editor and the referees for their careful reading and constructive comments, which greatly improved the paper. We are also grateful to Francesco Amodeo, Renee Bowen, Simone Galperti, Germán Gieczewski, Aram Grygorian, James Hamilton, Juan Herreño, Michael McMahon, Giacomo Rondina, Denis Shishkin, Domenico Siniscalco, Joel Sobel, Pietro Spini, Savitar Sundaresan, Johannes Wieland, and seminar participants at the CEPR Central Bank Communication Seminar Series for helpful comments and suggestions.

Data availability

Data will be made available on request.

Appendix A. Proofs

Proof of Proposition 1. Assume that a PBE of the game exists and let $\sigma = (\sigma_m, \sigma_x, \sigma_r)$ be the corresponding strategy profile, consisting of a communication rule $\sigma_m : [-\phi_1, \phi_1] \rightarrow \Delta(M)$, a rate rule $\sigma_r : [-\phi_1, \phi_1] \times [-\phi_2, \phi_2] \times X^I \rightarrow \Delta(\mathbb{R})$ for the *CB*, and an investment plan

$\sigma_{x_i} : M \rightarrow \Delta(X)$ for each $i \in I$.³¹ Fix any $(\omega_1, \omega_2, m, x) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2] \times M \times X^I$, and consider the problem of the CB at $t = 2$ after history $(\omega_1, \omega_2, m, x)$. In any PBE, the CB must pick $\hat{r}$ solving

$$\min_{r \in \mathbb{R}} L_{CB}(x, r, \omega) \quad (\text{A.1})$$

for $\omega = \omega_1 + \omega_2$. Taking the derivative of the objective function with respect to r one obtains

$$\frac{\partial L_{CB}}{\partial r} = (1 - \alpha)(r - \omega) + \alpha(r - \bar{x})$$

where $\bar{x}$ is the average of $x_1, \dots, x_N$. Given the strict convexity of the objective function, the first-order condition is necessary and sufficient for optimality, and gives the unique PBE rate choice $\sigma_r^*(\omega_1, \omega_2, x)$ made by the CB at any history $(\omega_1, \omega_2, m, x)$ reached at $t = 2$ as a function of the corresponding (ω_1, ω_2, x) . From the first-order condition, $\sigma_r^*(\omega_1, \omega_2, x)$ must satisfy

$$\sigma_r^*(\omega_1, \omega_2, x) = (1 - \alpha)\omega + \alpha\bar{x}. \quad (\text{A.2})$$

Thus, any PBE must prescribe $\sigma_r = \sigma_r^*$. Note that communication and investment decisions are not functions of ω_2 , since they are taken before ω_2 is realized. Hence for each $(\omega_1, \omega_2) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2]$ and $m \in \text{supp}(\sigma_m(\omega_1))$, and $x_i \in \text{supp}(\sigma_{x_i}(m))$,

$$\frac{\partial \sigma_r(\omega_1, \omega_2, x)}{\partial \omega_2} = \frac{\partial \sigma_r^*(\omega_1, \omega_2, x)}{\partial \omega_2} = 1 - \alpha,$$

so that $0 < \frac{\partial \sigma_r^*}{\partial \omega_2} < 1$ follows from $\alpha \in (0, 1)$. Thus, for each $(\omega_1, \omega_2) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2]$, $0 < \frac{\partial r^*(\sigma, \omega_1, \omega_2)}{\partial \omega_2} < 1$, which proves [Proposition 1](#). $\square$

Proof of Proposition 2.

For any $N \geq 1$, completing the square in (2) yields

$$L_{CB}(x, r, \omega) = \frac{1}{2}[r - ((1 - \alpha)\omega + \alpha\bar{x})]^2 + \frac{\alpha(1 - \alpha)}{2}(\omega - \bar{x})^2 + \frac{\alpha}{2N} \sum_{i \in I} (x_i - \bar{x})^2.$$

Because investment positions are chosen before the independent, mean-zero shock ω_2 , taking expectations under σ gives

$$\mathbb{E}[L_{CB} \mid \sigma] = \frac{\alpha(1 - \alpha)}{2}\sigma_2^2 + \frac{1}{2}\mathbb{E}[(r - ((1 - \alpha)\omega + \alpha\bar{x}))^2 \mid \sigma] + \frac{\alpha(1 - \alpha)}{2}\mathbb{E}[(\omega_1 - \bar{x})^2 \mid \sigma] + \frac{\alpha}{2N}\mathbb{E}\left[\sum_{i \in I} (x_i - \bar{x})^2 \mid \sigma\right],$$

where $\sigma_2^2 = \text{Var}(\omega_2)$. The last three terms are nonnegative. Hence, $\mathbb{E}[L_{CB} \mid \sigma] \geq \frac{1}{2}\alpha(1 - \alpha)\sigma_2^2$, with equality if and only if $r = (1 - \alpha)\omega + \alpha\bar{x}$, $\bar{x} = \omega_1$, and $x_i = \bar{x}$ for every $i \in I$, almost surely under the joint distribution induced by σ .

It remains to relate these equality conditions to conditions (i) and (ii). Since σ_{x_i} conditions only on m , x_i and ω_1 are conditionally independent given m . Consequently,

$$\mathbb{E}[(\omega_1 - x_i)^2 \mid \sigma] = \mathbb{E}[\text{Var}(\omega_1 \mid m, \sigma_m) \mid \sigma] + \mathbb{E}[(x_i - \mathbb{E}[\omega_1 \mid m, \sigma_m])^2 \mid \sigma].$$

Both terms on the right-hand side are nonnegative. Therefore, $x_i = \omega_1$ for every $i \in I$ almost surely if and only if conditions (i) and (ii) hold almost surely. Condition (iii) is exactly the remaining equality condition for the policy rate. Thus, a strategy profile achieves $\hat{W}$ if and only if conditions (i)–(iii) hold almost surely.

The lower bound is attainable for every finite N : take the communication rule $m = \omega_1$, let every bank choose $x_i = m$, and let the policy rule assign probability one to $r = (1 - \alpha)\omega + \alpha\bar{x}$. It follows that $\hat{W} = -\frac{1}{2}\alpha(1 - \alpha)\sigma_2^2$, which does not depend on N .

It remains to establish equilibrium existence when banks are atomistic. At the policy stage, the central bank's unique best response after every history is the policy rule σ_r^* in (A.2). Let $\mu(m) = \mathbb{E}[\omega_1 \mid m, \sigma_m]$. Given σ_r^* , an atomistic bank takes $\bar{x}$, and hence the policy rate, as unaffected by its own investment. The term $-\beta r$ in its payoff is therefore independent of its choice, and its unique best response after message m is $x_i = \mathbb{E}[r \mid m] = (1 - \alpha)\mu(m) + \alpha\bar{x}$. Aggregate consistency then requires $\bar{x} = (1 - \alpha)\mu(m) + \alpha\bar{x}$, and hence $x_i = \bar{x} = \mu(m)$ for every bank.

Consider the strategy profile $\sigma^{com} = (\sigma_m^{com}, \sigma_x^{com}, \sigma_r^*)$, where σ_m^{com} assigns probability one to $m = \omega_1$ and every bank's investment rule assigns probability one to $x_i = m$. Let beliefs assign probability one to $\omega_1 = m$ after every $m \in M = [-\phi_1, \phi_1]$. These beliefs are consistent with σ_m^{com} , and the investment and policy rules are sequentially optimal by the preceding arguments. If the central bank observes ω_1 and instead sends message m , its expected continuation loss is $\frac{1}{2}\alpha(1 - \alpha)[(\omega_1 - m)^2 + \sigma_2^2]$. This expression is minimized at $m = \omega_1$, so the central bank has no profitable communication deviation. Therefore, σ^{com} is a PBE satisfying conditions (i)–(iii) and achieving $\hat{W}$. $\square$

Proof of Proposition 3. We start from part (i). Fix a communication rule $\sigma_m \in \Sigma_m$. We now derive the oligopolistic best responses $\sigma_{x_i}^{oli}$ and σ_r^{oli} , where σ_r^{oli} is a best response to $(\sigma_x^{oli}, \sigma_m)$ and, for each $i \in I$, $\sigma_{x_i}^{oli}$ is a best response to $(\sigma_m, \sigma_r^{oli})$.

First, $\sigma_r^{oli} = \sigma_r^*$ holds pointwise, where σ_r^* satisfies Eq. (A.2). Hence,

$$\sigma_r^{oli}(\omega_1, \omega_2, x) = \sigma_r^*(\omega_1, \omega_2, x) = (1 - \alpha)\omega + \alpha\bar{x} \quad (\text{A.3})$$

³¹ We only look at rate rules which are not contingent on messages $m \in M$, as doing otherwise is not meaningful in this static setting.

for each $(\omega_1, \omega_2, x) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2] \times X^I$, where $\bar{x} = \frac{1}{N} \sum_{i \in I} x_i$.

Plugging (A.3) in the market player payoff specification (1), the payoff of i as a function of $(\omega_1, \omega_2, x) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2] \times X^I$ becomes

$$u_i(\sigma_r^*(\omega_1, \omega_2, x), x_i) = -\frac{1}{2} [x_i - (1 - \alpha)\omega - \alpha\bar{x}]^2 - \beta(1 - \alpha)\omega - \alpha\beta\bar{x}.$$

At $t = 1$ and after any message $m \in M$, each i chooses a $\sigma_{x_i}^{oli}(m)$ supported on the set of maximizers of $\mathbb{E}[u_i(\sigma_r^*(\omega_1, \omega_2, x), x_i) | m, \sigma_{x_{-i}}(m)]$ where the expectation is based on a posterior belief on ω_1 that satisfies Bayes rule whenever possible, and on the strategy profile $\sigma_{x_{-i}}$ played by all other investors. As the objective function is strictly concave, the unique maximizer x_i^* satisfies the FOC,

$$[x_i^* - (1 - \alpha)\mathbb{E}[\omega_1 | m, \sigma_m] - \alpha\bar{x}^e] \left(1 - \frac{\alpha}{N}\right) + \frac{\alpha\beta}{N} = 0 \quad (\text{A.4})$$

where $\bar{x}^e = \frac{x_i^*}{N} + \mathbb{E}[\sum_{j \in I \setminus \{i\}} \frac{x_j}{N} | m, \sigma_{x_{-i}}(m)]$. First, note that the uniqueness of the maximizer implies that no mixed strategies are played. Second, the fact that x_i^* depends on $x_j, j \neq i$, only through the average investment $\bar{x}$ implies symmetry of investment whenever m is in the support of σ_m .³² Hence, setting $\bar{x}^e = x_i^*$ yields the optimal strategy of each $i \in I$ after message m in the support of σ_m . In particular, the profile best responses $\sigma_{x_i}^{oli}$ must satisfy

$$\sigma_{x_i}^{oli}(m) - \mathbb{E}[\omega_1 | m, \sigma_m] = -\frac{\alpha\beta}{(N - \alpha)(1 - \alpha)} \quad (\text{A.5})$$

which proves part (i).

For part (ii), let $\sigma_m = \sigma_m^{com}$ for some σ_m^{com} fully revealing the state and let σ_x^{oli} satisfy condition ((A.5)) with respect to σ_m^{com} . On the path of play we have that for each $\omega_1 \in [-\phi_1, \phi_1]$, $m \in \text{supp}(\sigma_m^{com}(\omega_1))$ and each $i \in I$, $\sigma_{x_i}^{oli}(m) = \omega_1 - \frac{\alpha\beta}{(N - \alpha)(1 - \alpha)}$, so that

$$\bar{x} = \omega_1 - \frac{\alpha\beta}{(N - \alpha)(1 - \alpha)}. \quad (\text{A.6})$$

Substituting (A.6) into (A.3) gives the expression of $r^*(\sigma^{oli, tr}, \omega_1, \omega_2)$, for $\sigma^{oli, tr} = (\sigma_m^{com}, \sigma_x^{oli}, \sigma_r^*)$,

$$r^*(\sigma^{oli, tr}, \omega_1, \omega_2) = \omega_1 + (1 - \alpha)\omega_2 - \frac{\alpha^2\beta}{(N - \alpha)(1 - \alpha)},$$

which proves part (ii). $\square$

Proof of Proposition 4. First, in any PBE the CB optimally plays the strategy given by (A.3) at $t = 2$. Fixing this rate rule, the expected loss of the CB given state $\omega_1 \in [-\phi_1, \phi_1]$ as a function of $x \in X^I$ is

$$\mathbb{E}[L_{CB}(x, \omega) | \omega_1] = \frac{1}{2} \alpha^2 (1 - \alpha) \mathbb{E}[(\bar{x} - \omega)^2 | \omega_1] + \alpha \frac{1}{2N} \sum_{i \in I} \mathbb{E}[(1 - \alpha)\omega + \alpha\bar{x} - x_i]^2 | \omega_1] \quad (\text{A.7})$$

For any on-path message m , Bayes' rule gives all banks the same posterior belief about ω_1 . Moreover, their continuation investment problems are strictly concave. Subtracting the first-order condition of bank j from that of bank i therefore gives

$$\left(1 - \frac{\alpha}{N}\right) [\sigma_{x_i}^{oli}(m) - \sigma_{x_j}^{oli}(m)] = 0.$$

Hence, investment is symmetric after every on-path message: $\sigma_{x_i}^{oli}(m) = \sigma_{x_j}^{oli}(m)$ for every $i, j \in I$. Let $\mathcal{M}^{on}$ denote the set of on-path messages.

Accordingly, after any $m \in \mathcal{M}^{on}$, the induced investment profile belongs to

$$\bar{X} = \{x \in X^I : x_i = x_j \text{ for every } i, j \in I\}.$$

For $x \in \bar{X}$, Eq. (A.7) simplifies to

$$\mathbb{E}[L_{CB}(x, \omega) | \omega_1] = \frac{1}{2} \alpha (1 - \alpha) \mathbb{E}[(\bar{x} - \omega)^2 | \omega_1] = \frac{1}{2} \alpha (1 - \alpha) (\bar{x} - \omega_1)^2 + \frac{1}{2} \alpha (1 - \alpha) \sigma_2^2.$$

In a PBE, each type ω_1 must prefer every message used by that type to every other on-path message. From (A.5), the common investment following each $m \in \mathcal{M}^{on}$ is

$$\sigma_{x_i}^{oli}(m) = \mathbb{E}[\omega_1 | m, \sigma_m^{oli}] - \frac{\alpha\beta}{(N - \alpha)(1 - \alpha)}.$$

Consequently, it is necessary that, for every $\omega_1 \in [-\phi_1, \phi_1]$, every $m \in \text{supp}(\sigma_m^{oli}(\omega_1))$, and every $m' \in \mathcal{M}^{on}$,

$$(\sigma_{x_i}^{oli}(m) - \omega_1)^2 \leq (\sigma_{x_i}^{oli}(m') - \omega_1)^2. \quad (\text{A.8})$$

Together with Bayes' rule and banks' optimal investment responses, these are precisely the on-path incentive conditions of the Crawford–Sobel game characterized below. They therefore imply the partition restrictions stated in the proposition.

³² For m in the support of σ_m the expectation term in Eq. (A.4) takes the same value for all investors, but after a surprising m the expectation term might differ across investors since Bayes rule does not apply.

For completeness, off-path messages impose no additional restriction on the existence of these partition equilibria. Consider any Crawford–Sobel partition satisfying the preceding conditions and fix an arbitrary on-path message $m_0 \in \mathcal{M}^{\text{on}}$. Following every off-path message $m' \notin \mathcal{M}^{\text{on}}$, assign all banks the same belief they hold following m_0 , and let them choose the corresponding optimal investment. The message m' then induces exactly the same investment profile, and hence the same continuation outcome for the central bank, as m_0 . Because (A.8) already rules out a profitable deviation to m_0 , it also rules out a profitable deviation to m' . Combining these beliefs and investment plans with the policy rule σ_r^* therefore supports the partition as a PBE. The general solution of this cheap-talk game has been characterized by Crawford and Sobel (1982), hereafter CS.

The mapping to CS is particularly evident by noticing that we can transform the variables and parameters of interest in our setting to match a specific case of their formulation. First, define the following transformation: $t: \mathbb{R} \rightarrow \mathbb{R}$, $t(x) = (2\phi_1)^{-1} \left(x + \phi_1 + \frac{\alpha\beta}{(N-\alpha)(1-\alpha)} \right)$. For each $\omega_1 \in [-\phi_1, \phi_1]$ we can apply the transformation t to the perfect information ideal investment levels of the CB and the representative market player i , denoted as $x_{CB}^*(\omega_1) = \omega_1$ and $x_i^*(\omega_1) = \omega_1 - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$, respectively. The transformation gives $t(x_i^*(\omega_1)) = (2\phi_1)^{-1}(\omega_1 + \phi_1)$ and $t(x_{CB}^*(\omega_1)) = t(x_i^*(\omega_1)) + b$ for $b = \frac{\alpha\beta}{2\phi_1(N-\alpha)(1-\alpha)}$. Denote by $t_{\omega_1}^*$ the random variable equal to $(2\phi_1)^{-1}(\omega_1 + \phi_1)$, the transformed ideal point of i . First, note that $t_{\omega_1}^* \sim U[0, 1]$, supported on the unit interval as in CS. Second, for each $\omega_1 \in [-\phi_1, \phi_1]$ and $x_i \in \mathbb{R}$, the following proportionality holds:

$$\left(\omega_1 - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)} - x_i \right) \propto (t_{\omega_1}^* - t(x_i)).$$

Similarly,

$$(\omega_1 - x_i) \propto (t_{\omega_1}^* + b - t(x_i)).$$

Third, note that t is a bijection between X and X and is therefore invertible.

Substantially, t creates a “bridge” between the communication subgame characterized by (A.8) and the seminal example studied by CS (Section 4). Consider that example, and set $b = \frac{\alpha\beta}{2\phi_1(N-\alpha)(1-\alpha)}$. The above relations imply that (i) for every partition PBE of Example 4 where communication partitions $[0, 1]$ into P intervals with cutoffs $a_0 < a_1 < \dots < a_P$, there exists a PBE of our game where σ_m^{oli} partitions $[-\phi_1, \phi_1]$ into P intervals, with cutoffs $t^{-1}(a_0 + b) < t^{-1}(a_1 + b) < \dots < t^{-1}(a_P + b)$; and (ii) for every partition PBE of our game where σ_m^{oli} partitions $[-\phi_1, \phi_1]$ into P intervals with cutoffs $a_0 < a_1 < \dots < a_P$, there exists a PBE of CS’s seminal example, where communication partitions $[0, 1]$ into P intervals, with cutoffs $t(a_0) - b < t(a_1) - b < \dots < t(a_P) - b$.

Finally, for any equilibrium with P partition elements, the induced residual variance of ω_1 is $\hat{\sigma}_{1,P}^2 = \mathbb{E}[Var(\omega_1 | m, \sigma_{m,P}^{\text{oli}}) | \sigma_{m,P}^{\text{oli}}]$. Since, after any on-path message, given

$$\bar{x} = \mathbb{E}[\omega_1 | m, \sigma_{m,P}^{\text{oli}}] - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)},$$

we have

$$\mathbb{E}[(\bar{x} - \omega)^2 | \sigma_{m,P}^{\text{oli}}] = \sigma_2^2 + \left[\frac{\alpha\beta}{(N-\alpha)(1-\alpha)} \right]^2 + \hat{\sigma}_{1,P}^2.$$

The first two terms are exactly the terms present under transparent communication. Therefore,

$$W(\sigma_P^{\text{oli}}) = W(\sigma^{\text{oli},tr}) - \frac{1}{2}\alpha(1-\alpha)\hat{\sigma}_{1,P}^2,$$

which proves part (i).

This correspondence with CS also implies that the strong comparative statics in Section 4 of CS apply to our setting for $b = \frac{\alpha\beta}{2\phi_1(N-\alpha)(1-\alpha)}$, proving Proposition 4.i-ii. $\square$

Proof of Proposition 5. Let $b: [0, 1] \rightarrow \mathbb{R}$, $b(\tilde{\alpha}) = -\frac{\tilde{\alpha}\beta}{(N-\tilde{\alpha})(1-\tilde{\alpha})}$, mapping central banker types to the corresponding investment bias in equilibrium. The case $\tilde{\alpha} = 1$ generates no finite equilibrium investment bias and leads welfare to $-\infty$, so it is without loss to restrict attention to $\tilde{\alpha} \in [0, 1)$. Repeating the steps of the proof of Proposition 3 and replacing α with $\tilde{\alpha}$, yields that, for each $m \in M$, $(\omega_1, \omega_2) \in [-\phi_1, \phi_1] \times [-\phi_2, \phi_2]$, and $x \in X^I$,

$$\sigma_r^*(\omega_1, \omega_2, x; \tilde{\alpha}) = (1 - \tilde{\alpha})\omega + \tilde{\alpha}\bar{x}, \quad \sigma_{x_i}^{\text{oli}}(m; \tilde{\alpha}) = \mathbb{E}[\omega_1 | m, \sigma_m] + b(\tilde{\alpha}).$$

For compactness, write $Q(\tilde{\alpha}) = \tilde{\alpha}^2(1 - \alpha) + \alpha(1 - \tilde{\alpha})^2$. It follows that, under transparency and under cheap talk with the most informative equilibrium selected,

$$W(\sigma^{\text{oli},tr}(\tilde{\alpha})) = -\frac{1}{2}Q(\tilde{\alpha})(\sigma_2^2 + b(\tilde{\alpha})^2), \quad (\text{A.9})$$

$$W(\sigma_P^{\text{oli}}(\tilde{\alpha})) = -\frac{1}{2}Q(\tilde{\alpha})(\hat{\sigma}_{1,P}^2(\tilde{\alpha}) + \sigma_2^2 + b(\tilde{\alpha})^2). \quad (\text{A.10})$$

When $\tilde{\alpha} = 0$, the investment bias is zero and the most informative cheap-talk equilibrium is fully revealing; hence we use the convention $\hat{\sigma}_{1,P}^2(0) = 0$.

First, consider part (i) of the proposition. Comparing the welfare level under a maximally FS-agnostic central banker with the welfare level under an unbiased central banker gives

$$W(\sigma^{\text{oli},tr}(0)) > W(\sigma^{\text{oli},tr}(\alpha)) \iff \sigma_2 < \sqrt{\frac{\alpha}{1-\alpha}} \left(\frac{\beta}{N-\alpha} \right) \equiv \sigma_2^{\text{oli},tr},$$

$$W(\sigma_P^{oli}(0)) > W(\sigma_P^{oli}(\alpha)) \iff \sigma_2 < \sqrt{\frac{\alpha}{1-\alpha} \left(\frac{\beta}{N-\alpha} \right)^2 + \frac{1-\alpha}{\alpha} \hat{\sigma}_{1,P}^2(\alpha)} \equiv \sigma_2^{oli}.$$

Since $\hat{\sigma}_{1,P}^2(\alpha) > 0$, it follows that $\sigma_2^{oli} > \sigma_2^{oli,tr}$, proving part (i).

Next, we prove part (ii). We first consider transparent communication. Differentiating (A.9) with respect to $\tilde{\alpha}$ yields

$$\frac{\partial W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial \tilde{\alpha}} = (\alpha - \tilde{\alpha})(\sigma_2^2 + b(\tilde{\alpha})^2) - Q(\tilde{\alpha}) \frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}} b(\tilde{\alpha}).$$

Because $b(\tilde{\alpha}) < 0$ and $\frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}} < 0$ for $\tilde{\alpha} > 0$, the product $\frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}} b(\tilde{\alpha})$ is positive. Hence the derivative above is strictly negative for all $\tilde{\alpha} \geq \alpha$. At $\tilde{\alpha} = 0$, the derivative is $\alpha\sigma_2^2 > 0$.

We now show that $W(\sigma^{oli,tr}(\tilde{\alpha}))$ is strictly concave in $\tilde{\alpha}$. From (A.9),

$$-2W(\sigma^{oli,tr}(\tilde{\alpha})) = \sigma_2^2 Q(\tilde{\alpha}) + \beta^2 \alpha \left(\frac{\tilde{\alpha}}{N-\tilde{\alpha}} \right)^2 + \beta^2 (1-\alpha) \left(\frac{\tilde{\alpha}}{N-\tilde{\alpha}} \right)^2 \left(\frac{\tilde{\alpha}}{1-\tilde{\alpha}} \right)^2.$$

For any $c \geq 1$, $\left(\frac{\tilde{\alpha}}{c-\tilde{\alpha}} \right)^2$ is nonnegative, increasing, and strictly convex on $[0, 1)$, since

$$\frac{\partial}{\partial \tilde{\alpha}} \left(\frac{\tilde{\alpha}}{c-\tilde{\alpha}} \right)^2 = \frac{2c\tilde{\alpha}}{(c-\tilde{\alpha})^3} \geq 0, \quad \frac{\partial^2}{\partial \tilde{\alpha}^2} \left(\frac{\tilde{\alpha}}{c-\tilde{\alpha}} \right)^2 = \frac{2c(c+2\tilde{\alpha})}{(c-\tilde{\alpha})^4} > 0.$$

Moreover, $Q''(\tilde{\alpha}) = 2$. The product of two nonnegative, increasing, convex functions is convex, since $(fg)'' = f''g + 2f'g' + fg'' \geq 0$. Therefore, $-2W(\sigma^{oli,tr}(\tilde{\alpha}))$ is strictly convex, and $W(\sigma^{oli,tr}(\tilde{\alpha}))$ is strictly concave. Combining strict concavity with the sign of the derivative at 0 and on $[\alpha, 1)$, the transparent-communication objective has a unique maximizer $\tilde{\alpha}^{oli,tr} \in (0, \alpha)$.

We now establish the comparative statics of $\tilde{\alpha}^{oli,tr}$. Define $G(\tilde{\alpha}; \alpha, N, \sigma_2^2) = \frac{\partial W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial \tilde{\alpha}}$. At the optimum, $G = 0$ and $G_{\tilde{\alpha}} < 0$. First, $G_{\sigma_2^2} = \alpha - \tilde{\alpha}^{oli,tr} > 0$, so the implicit function theorem gives $\frac{\partial \tilde{\alpha}^{oli,tr}}{\partial \sigma_2^2} > 0$. Since $\sigma_2^2 = \frac{\phi_2^2}{3}$, it follows that $\tilde{\alpha}^{oli,tr}$ is increasing in ϕ_2 .

Second, since $\frac{\partial b}{\partial N} = -\frac{b}{N-\tilde{\alpha}}$ and $\frac{\partial^2 b}{\partial \tilde{\alpha} \partial N} = -\frac{1}{N-\tilde{\alpha}} \frac{\partial b}{\partial \tilde{\alpha}} - \frac{b}{(N-\tilde{\alpha})^2}$, differentiating G with respect to N gives

$$G_N = -\frac{2(\alpha - \tilde{\alpha})b(\tilde{\alpha})^2}{N-\tilde{\alpha}} + \frac{2Q(\tilde{\alpha})b(\tilde{\alpha})}{N-\tilde{\alpha}} \frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}} + \frac{Q(\tilde{\alpha})b(\tilde{\alpha})^2}{(N-\tilde{\alpha})^2}.$$

Using the first-order condition $Q(\tilde{\alpha})b(\tilde{\alpha}) \frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}} = (\alpha - \tilde{\alpha})(\sigma_2^2 + b(\tilde{\alpha})^2)$, evaluated at $\tilde{\alpha}^{oli,tr}$, we obtain

$$G_N = \frac{2(\alpha - \tilde{\alpha})\sigma_2^2}{N-\tilde{\alpha}} + \frac{Q(\tilde{\alpha})b(\tilde{\alpha})^2}{(N-\tilde{\alpha})^2} > 0.$$

Hence $\frac{\partial \tilde{\alpha}^{oli,tr}}{\partial N} > 0$.

Third, differentiating G with respect to α gives

$$G_{\alpha} = (\sigma_2^2 + b(\tilde{\alpha})^2) - (1 - 2\tilde{\alpha}) \frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}} b(\tilde{\alpha}).$$

Using again the first-order condition,

$$G_{\alpha} = \left[\frac{Q(\tilde{\alpha})}{\alpha - \tilde{\alpha}} - (1 - 2\tilde{\alpha}) \right] \frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}} b(\tilde{\alpha}) = \frac{\tilde{\alpha}(1 - \tilde{\alpha})}{\alpha - \tilde{\alpha}} \frac{\partial b(\tilde{\alpha})}{\partial \tilde{\alpha}} b(\tilde{\alpha}) > 0,$$

so that $\frac{\partial \tilde{\alpha}^{oli,tr}}{\partial \alpha} > 0$.

We now turn to cheap talk. By Crawford and Sobel (1982), the residual variance in the most informative equilibrium is weakly increasing in the preference misalignment. In the present notation, the relevant misalignment is $|b(\tilde{\alpha})|$, which is increasing in $\tilde{\alpha}$. Hence $\hat{\sigma}_{1,P}^2(\tilde{\alpha})$ is weakly increasing in $\tilde{\alpha}$. The selected residual variance is also continuous: within a fixed number of partition elements this follows from the Crawford-Sobel characterization, and at a cutoff where the maximal feasible number of partition elements falls from P to $P-1$, the first interval of the P -partition shrinks to zero and the remaining $P-1$ intervals coincide with the $(P-1)$ -partition at that cutoff. Therefore $W(\sigma_P^{oli}(\tilde{\alpha}))$ is continuous on $[0, 1)$, with $W(\sigma_P^{oli}(\tilde{\alpha})) \rightarrow -\infty$ as $\tilde{\alpha} \uparrow 1$, and a maximizer exists.

For $\tilde{\alpha} \geq \alpha$, $W(\sigma^{oli,tr}(\tilde{\alpha}))$ is strictly decreasing. Moreover, $Q(\tilde{\alpha})$ and $\hat{\sigma}_{1,P}^2(\tilde{\alpha})$ are weakly increasing on $[\alpha, 1)$, so the additional cheap-talk term $-\frac{1}{2}Q(\tilde{\alpha})\hat{\sigma}_{1,P}^2(\tilde{\alpha})$ is weakly decreasing on $[\alpha, 1)$. Thus $W(\sigma_P^{oli}(\tilde{\alpha}))$ is strictly decreasing on $[\alpha, 1)$, and no maximizer lies in $(\alpha, 1)$.

To rule out $\tilde{\alpha} = \alpha$, note that $\frac{\partial W(\sigma^{oli,tr}(\tilde{\alpha}))}{\partial \tilde{\alpha}} \Big|_{\tilde{\alpha}=\alpha} < 0$, so for small enough $\varepsilon > 0$ there is $c > 0$ such that

$$W(\sigma^{oli,tr}(\alpha - \varepsilon)) - W(\sigma^{oli,tr}(\alpha)) > c\varepsilon.$$

Also, $Q(\alpha - \varepsilon) = Q(\alpha) + \varepsilon^2$ and $\hat{\sigma}_{1,P}^2(\alpha - \varepsilon) \leq \hat{\sigma}_{1,P}^2(\alpha)$. Hence

$$Q(\alpha - \varepsilon)\hat{\sigma}_{1,P}^2(\alpha - \varepsilon) - Q(\alpha)\hat{\sigma}_{1,P}^2(\alpha) \leq \varepsilon^2\hat{\sigma}_{1,P}^2(\alpha).$$

It follows that, for $\varepsilon > 0$ small enough,

$$W(\sigma_P^{oli}(\alpha - \varepsilon)) - W(\sigma_P^{oli}(\alpha)) \geq c\varepsilon - \frac{1}{2}\varepsilon^2\hat{\sigma}_{1,P}^2(\alpha) > 0.$$

Therefore any cheap-talk maximizer satisfies $\tilde{\alpha}^{oli} \in [0, \alpha)$.

It remains to compare the welfare losses under the two optimal central bankers. If $\tilde{\alpha}^{oli} = 0$, then $W(\sigma_{\tilde{P}}^{oli}(\tilde{\alpha}^{oli})) = W(\sigma_{\tilde{P}}^{oli, tr}(0)) < W(\sigma_{\tilde{P}}^{oli, tr}(\tilde{\alpha}^{oli, tr}))$, because the transparent optimum is unique and interior. If instead $\tilde{\alpha}^{oli} > 0$, then $\hat{\sigma}_{1, \tilde{P}}^2(\tilde{\alpha}^{oli}) > 0$, and therefore

$$W(\sigma_{\tilde{P}}^{oli}(\tilde{\alpha}^{oli})) < W(\sigma_{\tilde{P}}^{oli, tr}(\tilde{\alpha}^{oli})) \leq W(\sigma_{\tilde{P}}^{oli, tr}(\tilde{\alpha}^{oli, tr})).$$

Thus the welfare loss under transparency and $\tilde{\alpha}^{oli, tr}$ is smaller than the welfare loss under cheap talk and $\tilde{\alpha}^{oli}$. This completes the proof of part (ii).

Finally, we prove part (iii). Under transparent communication, banks' payoff as a function of $\tilde{\alpha}$ is

$$\begin{aligned} EU_i^{oli, tr}(\tilde{\alpha}) &= -\frac{1}{2}(1-\tilde{\alpha})^2(\sigma_2^2 + b(\tilde{\alpha})^2) + \frac{(\tilde{\alpha}\beta)^2}{(N-\tilde{\alpha})(1-\tilde{\alpha})} \\ &= -\frac{1}{2}(1-\tilde{\alpha})^2\sigma_2^2 + \frac{(2N-1-\tilde{\alpha})(\tilde{\alpha}\beta)^2}{2(N-\tilde{\alpha})^2(1-\tilde{\alpha})} \\ &= -\frac{1}{2}(1-\tilde{\alpha})^2\sigma_2^2 + \left(\frac{\tilde{\alpha}\beta}{N-\tilde{\alpha}}\right)^2 \left(\frac{N-1}{1-\tilde{\alpha}} + \frac{1}{2}\right), \end{aligned}$$

which is strictly increasing in $\tilde{\alpha}$. Hence, under transparent communication, banks are strictly better off with the unbiased central banker than with any FS-agnostic central banker. To see that under cheap talk banks can be better off with an FS-agnostic central banker relative to an unbiased one, consider the example at the end of Section 3.5. There, for $N = 1$, $\phi_1 = \frac{1}{2}$, $\frac{\alpha\beta}{(1-\alpha)^2} \in (\frac{1}{4}, \frac{1}{\sqrt{12}})$, and σ_2^2 small enough, the bank's payoff under $\tilde{\alpha} = 0$ is strictly larger than under $\tilde{\alpha} = \alpha$. This example proves part (iii), and completes the proof of Proposition 5. $\square$

Proof of Proposition 6. We want to show that there exists a $\delta^* \in (0, 1)$ such that the proposed strategy is a PBE of the repeated game with discount factor $\delta \geq \delta^*$. Note that in the punishment phase players revert to a fixed stage game PBE, so only on-path deviations need to be checked.

Consider the *CB* first. The *CB* actions at time $s < \tau$ do not influence period τ play of any player. Hence, the only possible profitable deviation by the *CB* would increase its current stage payoff. Such a deviation is not profitable because, on the equilibrium path, stage-game play satisfies the efficiency conditions of Proposition 2.

Turning to the deviation incentives of player $i \in I$, the no profitable deviation condition requires

$$\frac{1}{2} \left[\frac{\alpha\beta}{N-\alpha} \right]^2 + \frac{\delta}{1-\delta} \left\{ \frac{1}{2} \left[\frac{\alpha\beta}{N-\alpha} \right]^2 \frac{2N-1-\alpha}{(1-\alpha)} - \frac{1}{2}(1-\alpha)^2 Var(\omega_1) \right\} \leq 0,$$

where the left hand side is the net benefit that i obtains if she deviates to the stage best response once and then follows the equilibrium strategy in all future stages. Using $Var(\omega_1) = \frac{\phi_1^2}{3}$, solving the above condition yields

$$\delta^* \equiv \frac{\left[\frac{\alpha\beta}{N-\alpha} \right]^2}{\left[\frac{\alpha\beta}{N-\alpha} \right]^2 \left[1 - \frac{2N-1-\alpha}{1-\alpha} \right] + \frac{\phi_1^2}{3}(1-\alpha)^2}.$$

The equilibrium exists only if $\delta^* < 1$, which is true if and only if $\phi_1 > \frac{\alpha\beta}{(N-\alpha)(1-\alpha)} \sqrt{3 \frac{2N-1-\alpha}{1-\alpha}}$. $\square$

Proof of Proposition 7. The proof of part (i) only requires constructing a collusive PBE strategy profile that implements efficient play on path in every stage, so that stage play on the equilibrium path satisfies the conditions of Proposition 2. We then verify that such equilibrium exists if $\hat{\sigma}_{1, \tilde{P}}^2(N) > \left[\frac{\alpha\beta}{(N-\alpha)(1-\alpha)} \right]^2 \frac{2N-1-\alpha}{1-\alpha}$.

Consider the following profile. At $\tau = 0$, *CB* chooses $m^\tau = \omega_1^\tau$; at $\tau > 0$, the *CB* selects $m^\tau = \omega_1^\tau$ if $x_i^s = m^s$ for each $s < \tau$ and $i \in I$; if there exists $s < \tau$ and $i \in I$ such that $x_i^s \neq m^s$, then the *CB* uses the stage communication rule $\sigma_{m, \tilde{P}}$ corresponding to the most informative stage game PBE when the bias is at the no-coordination level $-\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$. Each investor $i \in I$ plays the following strategy. At $\tau = 0$, $x_i^\tau = m^\tau$. At $\tau > 0$, $i \in I$ selects $x_i^\tau = m^\tau$ if $x_j^s = m^s$ for each $s < \tau$ and $j \in I$; if there exists $s < \tau$ and $j \in I$ such that $x_j^s \neq m^s$, then $x_i^\tau = \mathbb{E}[\omega_1^\tau | m^\tau, \sigma_{m, \tilde{P}}] - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$. At each $\tau \geq 0$, *CB* uses the policy rule $r^\tau = (1-\alpha)\omega^\tau + \alpha\bar{x}^\tau$.

By construction, in every stage the *CB* communicates as much as possible given the equilibrium expected investment bias of that period. As in the case of the previous proposition, punishment is carried out by playing a fixed stage game PBE in any period after a deviation (regardless of how the history of play evolves after the deviation). Hence deviations are not profitable during the punishment phase. Moreover, the *CB* has no incentive to deviate from the equilibrium path, since efficient play is implemented on the equilibrium path in each stage. Hence, the proposed strategy profile is a PBE of the repeated game if no $i \in I$ has an incentive to make a one-shot deviation from the equilibrium path, which requires

$$\frac{1}{2} \left[\frac{\alpha\beta}{N-\alpha} \right]^2 + \frac{\delta}{1-\delta} \left\{ \frac{1}{2} \left[\frac{\alpha\beta}{N-\alpha} \right]^2 \frac{2N-1-\alpha}{1-\alpha} - \frac{1}{2}(1-\alpha)^2 \hat{\sigma}_{1, \tilde{P}}^2(N) \right\} \leq 0,$$

where the left hand side is the net benefit that i obtains if she does the one-shot deviation at time τ that maximizes her expected payoff at τ given the other players' strategies and then follows the equilibrium strategy in all future stages. Note that at $\delta = \delta_1^*$ the

previous relation must hold with equality. Simple algebra yields

$$\delta_1^* \equiv \frac{\left[\frac{\alpha\beta}{N-\alpha}\right]^2}{\left[\frac{\alpha\beta}{N-\alpha}\right]^2 \left[1 - \frac{2N-1-\alpha}{1-\alpha}\right] + (1-\alpha)^2 \hat{\sigma}_{1,\bar{P}}^2(N)}.$$

Existence requires $\delta_1^* < 1$, which holds if and only if $\hat{\sigma}_{1,\bar{P}}^2(N) > \left[\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}\right]^2 \frac{2N-1-\alpha}{1-\alpha}$. Note that the latter inequality also guarantees that i 's stage-game payoff under the equilibrium considered is greater than under the most informative stage-game equilibrium for the parameters considered.

The proof of part (ii) only requires constructing a collusive PBE strategy profile that mimics the stage-game equilibrium when $N = 1$ on the path of play in every stage. We will then show that existence holds if $\hat{\sigma}_{1,\bar{P}}^2(1) - \hat{\sigma}_{1,\bar{P}}^2(N) < \left[\frac{\alpha\beta}{(1-\alpha)^2}\right]^2 - \left[\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}\right]^2 \frac{2N-1-\alpha}{1-\alpha}$. Consider the following profile. At $\tau = 0$, CB uses a stage communication rule $\sigma_{m,\bar{P}}^1$ corresponding to the most informative stage game PBE when market players' bias is at the monopolistic level $-\frac{\alpha\beta}{(1-\alpha)^2}$; at $\tau > 0$, the CB keeps using the same communication rule if $x_i^s = \mathbb{E}[\omega_1^s | m^s, \sigma_{m,\bar{P}}^1] - \frac{\alpha\beta}{(1-\alpha)^2}$ for each $s < \tau$ and $i \in I$; if there exists $s < \tau$ and $i \in I$ such that $x_i^s \neq \mathbb{E}[\omega_1^s | m^s, \sigma_{m,\bar{P}}^1] - \frac{\alpha\beta}{(1-\alpha)^2}$, then CB uses a stage communication rule $\sigma_{m,\bar{P}}^N$ corresponding to the most informative stage game PBE when the bias is at the no-coordination level $-\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$. Each market player $i \in I$ plays the following strategy. At $\tau = 0$, $x_i^\tau = \mathbb{E}[\omega_1^\tau | m^\tau, \sigma_{m,\bar{P}}^1] - \frac{\alpha\beta}{(1-\alpha)^2}$. At $\tau > 0$, $i \in I$ selects $x_i^\tau = \mathbb{E}[\omega_1^\tau | m^\tau, \sigma_{m,\bar{P}}^1] - \frac{\alpha\beta}{(1-\alpha)^2}$ if $x_j^s = \mathbb{E}[\omega_1^s | m^s, \sigma_{m,\bar{P}}^1] - \frac{\alpha\beta}{(1-\alpha)^2}$ for each $s < \tau$ and $j \in I$; if there exists $s < \tau$ and $j \in I$ such that $x_j^s \neq \mathbb{E}[\omega_1^s | m^s, \sigma_{m,\bar{P}}^1] - \frac{\alpha\beta}{(1-\alpha)^2}$, then $x_i^\tau = \mathbb{E}[\omega_1^\tau | m^\tau, \sigma_{m,\bar{P}}^N] - \frac{\alpha\beta}{(N-\alpha)(1-\alpha)}$. At each $\tau \geq 0$, the CB uses the policy rule $r^\tau = (1-\alpha)\omega^\tau + \alpha\bar{x}^\tau$.

By construction, in every stage the CB communicates as much as possible given the equilibrium expected investment bias of that period. As in the previous cases, punishment is carried out by playing a stage game PBE in any period after a deviation (regardless of how the history of play evolves after the deviation). Hence deviations are not profitable during the punishment phase. The CB has no incentive to deviate from the equilibrium path, since it is maximizing its stage game expected payoff given the market players' strategy, and CB deviations have no influence on subsequent play. Hence, the proposed strategy profile is a PBE of the repeated game if no $i \in I$ has an incentive to make a one-shot deviation from the equilibrium path. Following the same procedure as in the previous proofs yields that δ_2^* exists in $(0, 1)$ as long as the punishment stage payoff of market player $i \in I$ is strictly lower than her on path stage payoff. The condition requires

$$-\frac{1}{2}(1-\alpha)^2(\hat{\sigma}_{1,\bar{P}}^2(1) + \sigma_2^2) + \frac{1}{2}\left[\frac{\alpha\beta}{1-\alpha}\right]^2 > -\frac{1}{2}(1-\alpha)^2(\hat{\sigma}_{1,\bar{P}}^2(N) + \sigma_2^2) + \frac{1}{2}\left[\frac{\alpha\beta}{N-\alpha}\right]^2 \frac{2N-1-\alpha}{1-\alpha},$$

or equivalently,

$$\hat{\sigma}_{1,\bar{P}}^2(1) - \hat{\sigma}_{1,\bar{P}}^2(N) < \left[\frac{\alpha\beta}{(1-\alpha)^2}\right]^2 - \left[\frac{\alpha\beta}{(N-\alpha)(1-\alpha)}\right]^2 \frac{2N-1-\alpha}{1-\alpha}.$$

which simply amounts to requiring that the gain from policy influence in the single-bank case exceeds the surprise loss due to less transparent communication than in the no-coordination case. Letting $N \rightarrow \infty$ in the previous expression, banks are better off in the monopolistic equilibrium relative to the efficient equilibrium if and only if $\hat{\sigma}_{1,\bar{P}}^2(1) < \left[\frac{\alpha\beta}{(1-\alpha)^2}\right]^2$, which completes the proof. $\square$

Appendix B. Model extensions

B.1. An alternative payoff specification for market players

Our specification of market players' payoffs follows three criteria. First, it captures the benefits of correctly reading policymakers' intentions; intuitively, for any central bank rate choice, a financial institution is better off if it forecasts the policy correctly, while it incurs a costly readjustment otherwise. Second, the specification easily accommodates the idea that rate cuts boost financial institutions' profits (for $\beta > 0$), either directly or because they stimulate the economy. Third, it achieves these goals in a simple, parsimonious way.

The simplicity of our model comes with the benefit of making the mechanisms at play as transparent as possible but at the risk of overlooking relevant aspects. An important assumption, which naturally comes in handy, is additive separability: we abstract from cases where the gain from learning the future policy depends on whether such policy is a hike or a cut. Doing so greatly simplifies the analysis of the communication stage. At the same time, profit functions that are additively separable in a "surprise loss" and a "directional preference" are not uncommon in (simple) economic decision problems under uncertainty.³³

³³ One example is the classical case of a monopolist who faces a linear demand $p = a - x$, has a constant marginal cost of production c , and chooses output x . As is standard $a > c > 0$, but the monopolist is uncertain about the precise level of $\gamma = a - c$, because a, c , or both depend on economy or policy shocks. Then, the monopolist's profits can be decomposed as the sum of two terms: one representing full-information maximized profits, monotonic (increasing) in $\frac{\gamma}{2}$, and an additive "surprise loss," proportional to $(\frac{\gamma}{2} - x)^2$.

Importantly, bank profits might depend on where r lies relative to ω , rather than on the level of r . Naturally, one could just replace the term βr with $\beta(r - \omega)$ in the market player payoff specification (1), without affecting any of our results. Still, assuming linearity of the directional preference in r can be seen as a strong assumption. An alternative—still parsimonious—specification addressing such concern could be the following,

$$u_i(x_i, r, \omega) = -\frac{1}{2}(x_i - r)^2 - \frac{\beta}{2}(r - \omega + \Delta)^2,$$

where $\beta \in (0, 1]$ and $\Delta \neq 0$. Drawing a parallel with the market player payoff specification (1), one can interpret $\Delta > 0$ as picking up banks' preference for a monetary policy with a “Fed put” component: conditional on $x_i = r$, bank profits will be maximized by $r = \omega - \Delta < \omega$. At the same time, the specification allows for a reversal interest rate, since when $r < \omega - \Delta$ banks are instead hurt by further rate cuts (see e.g., Abadi et al., 2023). Finally, under an alternative interpretation, the second squared loss can capture systematic disagreement between banks and the central bank regarding the best way to resolve the inflation vs unemployment trade-off, for instance, when the central bank is perceived to put too little weight on one of its traditional objectives.

Following the same steps of our main analysis using the above alternative payoff specification instead of specification (1) delivers very similar functional forms—and indeed the same qualitative results—as the ones presented in Section 3.³⁴

B.2. Partial investment plan observability

In this section, we study an extension of the game where we allow for the possibility that the CB cannot observe investment plans. For simplicity, we focus on the case $N = 1$, but the intuition is similar for $N > 1$. We assume that the game has the same structure presented in Section 2, except that at $t = 2$, the CB observes x with probability $p \in [0, 1]$ and fails to observe it with probability $1 - p$. The observability of x is learned at $t = 2$.

We look for an equilibrium σ where the market player uses a “pure” investment plan, $\sigma_x : M \rightarrow \mathbb{R}$. Given σ_x , m and ω , at $t = 2$, there are two cases. If CB cannot observe x , it minimizes $L_{CB}(\sigma_x(m), \omega)$, setting the rate equal to $(1 - \alpha)\omega + \alpha\sigma_x(m)$. If x is observed, the central bank will set a rate equal to $(1 - \alpha)\omega + \alpha x$. Then, at $t = 1$, σ_x is an equilibrium investment plan if, for each m in the support of σ_m , $\sigma_x(m)$ solves the minimization of

$$\frac{1}{2} \mathbb{E} [p(1 - \alpha)^2(x - \omega)^2 + (1 - p)[x - (1 - \alpha)\omega - \alpha\sigma_x(m)]^2 | m, \sigma_m] + \mathbb{E} [(1 - \alpha)\beta\omega | m, \sigma_m] + \alpha\beta[p\alpha x + (1 - p)\sigma_x(m)]$$

with respect to $x \in \mathbb{R}$. The FOC is necessary and sufficient for x^* to be a minimizer, hence we have

$$x^* = \frac{[p(1 - \alpha)^2 + (1 - p)(1 - \alpha)]\mathbb{E}[\omega_1 | m, \sigma_m] + (1 - p)\alpha\sigma_x(m) - p\alpha\beta}{p(1 - \alpha)^2 + 1 - p},$$

and, setting $\sigma_x = x^*$ we obtain

$$\sigma_x(m) = \mathbb{E}[\omega_1 | m, \sigma_m] - \frac{p\alpha\beta}{(1 - \alpha)[p(1 - \alpha) + (1 - p)]},$$

which gives us the expression for the equilibrium investment plan of the monopolist market player. Hence, even with partial observability, investment bias arises in equilibrium, provided that $p > 0$, and still takes an additive form.

It is then clear that the equilibrium communication will have the same main characteristics as those highlighted in Section 3.4, except that investment bias is now smaller in absolute value. Hence, our main results carry through when $p > 0$. Finally, note that the absolute value of the investment bias is increasing in p . This suggests that the first-best expected welfare could be achieved in a game where the CB can commit to being completely uninformed about x , and, in general, a commitment to low information about x could improve welfare.

B.3. Optimal central banker for fixed partition elements

In this section we prove that, when we let $\tilde{\alpha}$ vary and compare the equilibria that induce the same number of partition elements, ex ante welfare is maximized by some FS-agnostic central banker rather than by an unbiased or FS-oriented central banker. Let $P \in \mathbb{N}$ be such that, when $\tilde{\alpha} = \alpha$, an equilibrium partitioning $[-\phi_1, \phi_1]$ in P partition elements exists. Let $\tilde{\alpha}$ be in the interval $[0, \tilde{\alpha}]$, where $\tilde{\alpha}$ is any weight greater than or equal to α such that when $\tilde{\alpha} = \tilde{\alpha}$ an equilibrium in P partitions still exists. Consider the ex ante welfare when the equilibrium strategy profile is $\sigma_P^{oli}(\tilde{\alpha})$. Existence follows by the continuity of

$$W(\sigma_P^{oli}(\tilde{\alpha})) = W(\sigma_P^{oli, tr}(\tilde{\alpha})) - \frac{1}{2}[\tilde{\alpha}^2(1 - \alpha) + \alpha(1 - \tilde{\alpha})^2]\hat{\sigma}_{1, P}^2(\tilde{\alpha}).$$

In the proof of Proposition 5 we showed that $W(\sigma_P^{oli, tr}(\tilde{\alpha}))$ is strictly decreasing in $\tilde{\alpha}$ for $\tilde{\alpha} \geq \alpha$. It is therefore sufficient to show that $\hat{\sigma}_{1, P}^2(\tilde{\alpha})$ is non-decreasing in $\tilde{\alpha}$. This result follows immediately from Theorem 4 of Crawford and Sobel (1982) who provide an explicit

³⁴ The implications for investment and policy distortions, welfare, and communication remain the ones presented in the paper. Provided that the welfare weight on financial stability α is not much larger than the weight on the traditional monetary policy objectives (something that we deem reasonable), central bank design will also necessarily result in a FS-agnostic central bank, which can benefit all parties due to more transparent communication. The communication stage yields the same insights because even under the alternative specification, the investment bias discussed in Section 3.3 takes an additive form.

expression for the residual variance after communication in Eq. (25) of their paper. In the language of the current paper, they show that, for fixed P , the residual variance after communication $\hat{\sigma}_{1,P}^2$ is continuous and increasing in the preference misalignment between the sender and the receiver, $|b(\tilde{\alpha})|$. Since $|b(\tilde{\alpha})|$ is increasing in $\tilde{\alpha}$, $\hat{\sigma}_{1,P}^2(\tilde{\alpha})$ is increasing in $\tilde{\alpha}$. This guarantees that (i) when $\tilde{\alpha} \geq \alpha$, a marginal decrease in $\tilde{\alpha}$ increases $W(\sigma_P^{oli}(\tilde{\alpha}))$; (ii) any maximizer of $W(\sigma_P^{oli}(\tilde{\alpha}))$, if it exists, must be strictly less than α . Existence follows from the continuity of $W(\sigma_P^{oli}(\tilde{\alpha}))$ in $\tilde{\alpha}$. Finally, it follows immediately from Crawford and Sobel (1982), Eq. (25), that the derivative of $\hat{\sigma}_{1,P}^2(\tilde{\alpha})$ with respect to $\tilde{\alpha}$ is equal to 0 at $\tilde{\alpha} = 0$. Hence, $W(\sigma_P^{oli}(\tilde{\alpha}))$ is maximized at some $\tilde{\alpha} > 0$.

Appendix C. Silicon Valley Bank and Signature Bank case study

In this Appendix, we provide a case study of Silicon Valley Bank (SVB) and Signature Bank (SB) to help illuminate a key idea in our theory: when the central bank is concerned about financial sector losses, it will underreact to inflationary pressure or output gap concerns (see Proposition 1). Showing this underreaction also provides supportive evidence that the central bank is directly concerned about financial stability, over and above its traditional goals. This underreaction will take the form of interest rates being lower than they would have been had the central bank not been concerned by such losses. In the language of our model, the interest rate will be lower than the economy-stabilizing rate (i.e., the rate needed for inflation/output gap stability).³⁵

SVB was a bank with over \$200 billion in assets and SB was a bank with over \$100 billion in assets. In March 2023, SVB faced severe problems as a result of interest rate risk, which it had not actively hedged (Metrick, 2024). The increases in interest rates combined with inadequate hedging negatively impacted its asset valuations which led to its failure when depositors began to withdraw their deposits. The failure of SVB led to widespread concerns and similarly rapid deposit outflows at SB which subsequently failed. These failures ultimately led to the '2023 Banking Crisis' and broader concerns around systemic risk.³⁶

We do not claim that SVB had taken too much risk to push the Fed to keep rates low. The more pertinent question from the perspective of our theoretical model is instead whether these concerns around systemically risky banking sector losses induced the Fed to underreact to the high inflation during that period. Specifically, our theory predicts that the Fed set the fed funds rate (FFR) lower than it would have done in the counterfactual world of no banking crisis. However, this is difficult to prove given the empirical challenge of knowing that counterfactual.

We use three different methods to consider the counterfactual world: (1) Taylor rule predictions; (2) Fed Funds futures implied probabilities; and (3) FOMC minutes. Each of these supports the idea that the Fed kept rates lower than it would have done in the absence of banking sector concerns and therefore is consistent with the predictions from our theoretical model. Below, we describe each method in more detail.

Taylor rule predictions. The Taylor rule is an equation that prescribes a value for the FFR based on the values of inflation relative to target and the output gap (Taylor, 1993). In a sense, it captures the typical central bank loss function (the conventional loss term in payoff specification (2)) and can be thought of as a theory-based prediction. The Taylor rule (and modifications of it) have been used consistently as a 'benchmark' for the FFR. Therefore, one can compare the actual FFR to that predicted by the Taylor rule to see whether the Fed is keeping rates too low or too high.

We use the following Taylor rule³⁷:

$$\hat{r}_t = \bar{r}_t + \pi_t + \theta_\pi(\pi_t - \pi_t^*) + \theta_Y \tilde{Y}_t \quad (\text{C.1})$$

where $\hat{r}_t$ is the predicted FFR, $\bar{r}_t$ is the long-run real interest rate, π_t is inflation, π_t^* is the target inflation rate, $\tilde{Y}_t$ is the output gap, θ_π is the weight the central bank puts on inflation, and θ_Y is the weight on the output gap. We use commonly used values for the different variables. Specifically, we set the long-run real interest rate to 2%, the target inflation rate to 2%, θ_π to 0.5, and θ_Y to 0.5.

Given that the underlying data are quarterly, we can look at the predicted rate versus the actual rate in the second quarter of 2023 (the first data point after the Fed's rate-setting meeting). Our theory suggests that the Taylor rule predicted rate should be higher than the Fed's chosen rate because the former does not account for banking sector losses. Indeed, (C.1) predicts an FFR of 6.17% while the actual FFR was only 4.83%. While this is consistent with our theory, the Taylor rule is not always an accurate predictor of the FFR and therefore may not serve as an effective counterfactual.

Fed funds futures implied probabilities. Next, we consider a market-based prediction of the FFR. Specifically, the probabilities of different rate decisions are computed using 30-day Fed Funds futures pricing data.³⁸ The useful feature of this method is that we can see the distribution of expectations rather than a single outcome.

Fig. C.2 below shows the probabilities of the different actions that the Fed would take according to market expectations. We look at the different probabilities from Friday 10 March to Friday 17 March. The important events took place between Friday 10 March and Sunday 12 March, when both SVB and SB failed. Therefore, our theory would predict that FFR hikes should be less likely after

³⁵ Our theory also predicts that knowing the central bank will underreact, financial institutions will take greater risk through misaligned portfolio positions. This case study is only intended to show an example of the central bank's underreaction rather than the consequent risk-taking decisions of financial institutions.

³⁶ <https://www.fdic.gov/news/press-releases/2023/pr23017.html>

³⁷ Note that we use a simple Taylor rule. One can use a number of different modified rules. For the purposes of our analysis, the different rules overall do not make a material difference.

³⁸ Our data is from the CME FedWatch Tool.




Fig. C.2. Probability of different FFR actions.

that weekend than before it. This is precisely what we see. On Friday 10 March, the market priced in a 60% chance of a 25 basis point hike, a 40% chance of a 50 basis point hike, and a 0% chance of no change. These were driven by the very high inflation at the time. However, just a few days later, there was a significant change in these probabilities. The probability of a 50 basis point hike fell to zero percent and stayed there while the probability of no change rose to over a third.

Figure C.2 is consistent with our theory as inflation concerns did not suddenly disappear over that weekend. Indeed, inflation expectations for the year ahead rose substantially from March 2023 to April 2023 (3.6% versus 4.7%) which further suggests that the Fed's concern for financial stability was unlikely to be driven by concerns that the banking panic would materially reduce inflation and may instead have led to increased inflation expectations.³⁹

FOMC minutes. Finally, to show that the Fed was explicitly concerned about banking sector developments and this played a role in its decision not to raise rates as aggressively, we examine the minutes of the FOMC meeting in March 2023. The following excerpt from the minutes provides much clearer evidence of a counterfactual:

"Some participants noted that given persistently high inflation and the strength of recent economic data, they would have considered a 50 basis point increase in the target range to have been appropriate at this meeting *in the absence of recent developments in the banking sector* [emphasis added]."

The Fed ultimately decided on a 25 basis point hike which is consistent with our theory that the Fed did not raise rates as much as it would have done had there been no banking sector concerns.

Taken together, these three pieces of evidence suggest, through a leading example, (i) that the Fed is directly concerned about banking sector stability and (ii) systemic stresses in the banking sector can lead the Fed to underreact to inflation or the output gap, consistent with our theoretical model.⁴⁰

Appendix D. Empirical relevance of core assumptions

At the end of Section 2, we briefly discussed how to interpret our model in terms of a central bank with a financial stability component to its monetary policy rule and the implications of this adjusted rule when banks behave strategically. The applicability and insights of our model depend crucially on whether the core assumptions are empirically relevant. In this appendix, we discuss this empirical relevance in greater detail.

³⁹ The inflation expectations data is available from the University of Michigan Surveys of Consumers.

⁴⁰ Note that to have rigorous empirical evidence, one would need to conduct a systematic study of the data rather than a single case study.

Central bank behavior. The most important assumptions of our framework are embedded in payoff specification (2). The assumptions that influence the behavior of the central bank are mostly contained within the second term of the central bank loss function. First, we assume that the central bank is concerned with financial stability *in addition to* its typical concerns about the real economy (e.g., inflation and growth). This assumption is consistent with important institutional developments. Since the Global Financial Crisis, central banks around the world have adopted explicit financial stability objectives that go beyond traditional inflation and growth mandates. These objectives typically include monitoring the profitability of banks and the volatility of asset prices (Calvo et al., 2018).⁴¹ These changes in legal and operational frameworks suggest that financial stability is now recognized as a goal in its own right. If stability mattered only through its effect on inflation and growth, one might reasonably argue that such new mandates are redundant. Theoretically, one can interpret this additional loss as a social cost of disorderly rebalancing: even holding output and inflation constant, society would prefer a smoother path with less volatility.⁴²

An alternative interpretation of our loss function is that the central bank only cares about financial stability through its effect on growth. This perspective is close in spirit to Stein and Sunderam (2018), who model the Fed as averse to bond market volatility as a way to better deliver on its dual mandate. Like Stein and Sunderam (2018), we do not formally separate these channels. Nonetheless, one could in principle represent the central bank's growth objective as having two components: one not linked to financial instability, and one that is. Our results would then depend on the relative weight placed on the stability-related component.

While both interpretations above are consistent with our results, the empirical evidence supporting each interpretation is mixed. Some studies, such as Cieslak and Vissing-Jørgensen (2021), highlight the way financial stability matters is through its effects on growth expectations. Others, including Shapiro and Wilson (2022), document more direct attention by the Fed to stock market performance. Most relevant for our purposes is Peek et al. (2015), who show that financial variables enter the FOMC's loss function and significantly improve the fit of an estimated Taylor rule. They also provide suggestive evidence that the FOMC often behaves as if monetary policy has an explicit financial stability mandate, over and above its traditional inflation and growth objectives. This latter finding aligns with our case study in Appendix C, which shows that during the banking panic of 2023 the Fed deviated substantially from a standard Taylor rule by not hiking despite high inflation, contributing to a notable rise in inflation expectations for the year ahead. While the empirical evidence is not yet conclusive, it is nonetheless important to understand the implications of a monetary policy framework that places direct weight on financial stability. In that vein, one can think of our model as shedding light on the following question: *if the central bank cares directly about financial stability, what are the implications for communication and policymaking?*

Another component of the second term of the central bank loss function is the extent to which individual banks are non-atomistic, as captured by $\frac{1}{N}$. In our banking interpretation, these non-atomistic banks are systemic: we interpret $\frac{1}{N}$ as the policy influence of an individual bank which is decreasing in N . In our context, a bank has policy influence if (i) its stability bears implications for systemic risk; and (ii) it receives special attention from policymakers concerned with financial stability. Systemically important banks appear to satisfy both requirements. Naturally, these banks tend to be considerably larger than others, as size matters for systemic risk.⁴³

While it is difficult to measure policy influence directly, central banks engage much more with financial institutions than non-financial institutions. Analyzing entries in the calendars of Federal Reserve governors between 2007 and 2018, Morse and Vissing-Jørgensen (2020) document a striking finding: governors met with financial institutions and financial interest groups more than three times as much as with non-financial institutions and non-financial interest groups (1,573 interactions versus 507 interactions). This evidence suggests there is at least a forum for possible influence, but it does not conclusively prove policy influence. However, there are a number of additional factors that point towards the potential for policy influence. First, central banks with a regulatory mandate *must* account for the stability of systemic banks. We can think of systemic banks as representing our model when N is small. Given such banks receive enhanced monitoring (e.g., through stress tests) and are subject to additional regulations (e.g., the systemic risk buffer), we expect this additional engagement raises the potential for policy influence. Second, and perhaps more directly, central banks specifically monitor the profits of systemic banks which allows the investments of these banks to influence the central bank decision. A recent example that highlights the pressure the Federal Reserve faces in relation to raising interest rates was during the banking panic of 2023. During this time, the former chair of the Federal Deposit Insurance Corporation (one of the institutions with a financial stability mandate in the US) publicly called on the Federal Reserve to stop raising interest rates.⁴⁴ Indeed, during this banking panic, both the Fed's behavior and market expectations were consistent with the underreaction prediction of our model (see Appendix C for details).

Our results are similar in spirit to some of the findings in the literature around risk-taking of 'too-big-to-fail' banks (which we refer to as systemic banks in this paper). For example, Afonso et al. (2014) show empirically that banks deemed too-big-to-fail engage in greater risk-taking because they believe they will be rescued if they fail. However, taking greater risk, all else equal, increases the

⁴¹ One example from the Federal Reserve comes from its May 2023 Financial Stability Report. The Fed highlighted that higher interest rates could lead to credit losses for banks exposed to commercial real estate and that as a result it would expand its monitoring of banks with such exposures. See <https://www.federalreserve.gov/publications/files/financial-stability-report-20230508.pdf>. Another example includes the Bank of England which specifically publishes the indicators it uses for macroprudential purposes. These include banking sector profitability and the VIX index (a measure of volatility). See <https://www.bankofengland.co.uk/-/media/boe/files/core-indicators/countercyclical-capital-buffer.xlsx>.

⁴² While we do not microfound this volatility aversion explicitly, the intuition is straightforward under standard models of risk-averse consumers, where consumers would prefer to smooth consumption over different states.

⁴³ The Basel Committee's assessment methodology as well as that of US regulators for determining systemic risk specifically includes size as one of its core categories (<https://www.financialresearch.gov/bank-systemic-risk-monitor/>).

⁴⁴ See <https://www.cnn.com/2023/03/13/investing/sheila-bair-svb-fed-rates/index.html>

likelihood of their failure and so pushes the regulatory authority to provide more protective measures (e.g., asset guarantee programs) than it would have otherwise. Similarly, in our paper, systemic banks increase their exposure to losses from higher rates which pushes the central bank to be more accommodative by cutting rates or raising rates less than it would have otherwise.

An additional feature of our central bank loss function is our choice of quadratic losses. We use this functional form largely for tractability and consistency with the general form of central bank loss functions in the literature. However, one implication of this choice is that we have a continuously increasing penalty, which one could interpret as the central bank caring about bank profits generally, not just avoiding crashes. A potentially more realistic assumption would be that the central bank is only concerned about tail-risk-type events and thus is not concerned about small losses, but only large losses. The predictions of our model would likely continue to hold qualitatively, but may weaken quantitatively. Specifically, a tail-risk-type assumption would likely imply market distortions, policy distortions, and our underreaction mechanism are quantitatively weaker (on average) as the mechanism would only apply in some states of the world (when losses are beyond a specific threshold), while our current functional form assumes the mechanism applies at all times. Therefore, while our quadratic form choice is undoubtedly a simplification, it captures the essence of the mechanism in a tractable way.

Implicit in the behavior of the central bank is that it can fully observe the investments of each bank. Full observability is unlikely to hold for all banks at all times. However, central banks typically conduct stress tests of systemic banks to ensure they are sufficiently stable. Specifically, such tests seek to ensure that banks' investment positions are able to withstand adverse shocks (including unexpected changes in interest rates). Indeed, as part of their supervisory role, central banks directly monitor the investment positions of systemic banks.⁴⁵

If the central bank can observe investment positions and thus predict which systemic banks would suffer large losses in response to rate changes, then one might ask why it does not simply design appropriate interventions such as the provision of special liquidity lines or other financial support to deal with any potential consequences of monetary policy on financial stability. Indeed, such actions would directly offset our underreaction results. However, in reality, such targeted interventions can be difficult to implement, especially in the US, due to political or legal factors. For example, Section 13(3) of the Federal Reserve Act explicitly states that the Fed cannot 'aid a failing financial company.' Moreover, broad-based emergency support from the Fed requires Treasury (and thus political) approval.⁴⁶ Caballero and Simsek (2020) argue that monetary policy should have a financial stability goal precisely because macroprudential tools face these limitations. This limited nature of macroprudential tools is supported by the fact that the Federal Reserve has never enacted the countercyclical capital buffer (a core macroprudential tool). Therefore, while an important assumption in our framework is that there is no bank-specific tool that the central bank can use to offset the impact of monetary policy, this assumption is empirically relevant in the context of the US. Nonetheless, if we instead assumed that the central bank could not fully observe banks' investments then, as shown formally in Section B.2, our results are qualitatively the same, albeit the investment bias is quantitatively smaller.⁴⁷

While most of our examples and our case study in Appendix C focus on the Federal Reserve, the framework we develop applies to central banks more broadly. In particular, our model insights are most relevant to any jurisdiction in which the central bank has a financial stability mandate and the domestic financial sector is relatively concentrated. These conditions are not restrictive. For example, according to a survey of 114 countries by the Bank for International Settlements, over 80% of central banks have some form of explicit financial stability objective (Jeanneau, 2014). Moreover, World Bank data show that as of 2021, across 123 countries, the top five banks in a given country hold on average over 80% of assets, indicating that banking sectors are highly concentrated.⁴⁸

Bank behavior. The assumptions that influence the behavior of the market players (i.e., banks) are implied by the market player loss function (payoff specification (1)). The core assumption here is that banks have a directional preference for interest rates which implies that $\beta > 0$. We interpret $\beta > 0$ as banks having a *structural* positive interest rate exposure deriving from their primary business, meaning that their profits are typically negatively impacted by rate hikes.⁴⁹ In principle, the sign of β bears no importance for the theoretical results of this paper, which only require $\beta \neq 0$.⁵⁰ From a high-level perspective, $\beta \neq 0$ simply means that holding constant how surprising a policy change is, the direction of the change still matters.

For the central banking application, $\beta > 0$ is a natural benchmark. For instance, English et al. (2018) show that bank stock prices and profits respond negatively to an interest rate hike. Uppal (2025) finds a similar result and provides evidence that the overall decline in profits is due to rising delinquencies, in line with cross-country evidence by Zimmermann (2019), who finds profit reductions due to loan losses and lower credit growth.⁵¹ Hence, a positive β could capture both stock price reactions and banks' exposure to credit

⁴⁵ The 'Large Institution Supervision Coordinating Committee' at the Federal Reserve actively supervises and monitors the detailed investment portfolios of systemically important banks. See <https://www.federalreserve.gov/publications/files/liscc-program-manual-202302.pdf>.

⁴⁶ See <https://www.federalreserve.gov/aboutthefed/section13.htm>.

⁴⁷ We thank an anonymous referee for this insightful observation.

⁴⁸ We use the 5-bank asset concentration (series GFDD.OI.06) across all countries for which it is available.

⁴⁹ Begeau et al. (2025) show that the US banking sector typically carries a positive interest rate exposure.

⁵⁰ More generally, we could also allow for heterogeneity in β across institutions in terms of sign and magnitude of β . In this richer setup, the rate distortions of Section 3.3 would only require the *average* β to be different from 0, while our results on investment distortions would hold unless $\beta = 0$ for all institutions. The main results for central bank design would still hold under these types of investors' heterogeneity.

⁵¹ Much of the literature finds that unexpected rate hikes lead to bank stock price declines and short-term bank profit declines (e.g. Flannery and James, 1984; Alessandri and Nelson, 2015; Altavilla et al., 2019; Jiang et al., 2024)

risk via unexpected changes to interest rates, among other factors.⁵² A positive β is also consistent with the fact that banking crises are typically followed by sharp reductions in the interest rate set by the central bank to, among other things, support the banking sector.

Why would banks take risky positions that create even more losses in case of adverse rate shocks? The key is that, by behaving as if the rates were to remain lower than optimal for the economy, the systemic banks in our game make it harder for the central bank to make large and undesired policy changes. This result is shown in the second part of [Proposition 3](#) if we assume commitment to fully informative communication: the on-path rate decision is systematically lower than the one that maximizes welfare. As such, our results provide one possible microfoundation for the empirical evidence on interest rate exposure.

Communication. The discussion on central bank behavior above provides evidence that central banks, explicitly or implicitly, pursue a financial stability objective alongside their traditional mandate. Therefore, if central banks can communicate, one would expect that at least one purpose of this communication would be to support the financial stability objective. One way the central bank can support this objective is to help ensure banks make investment decisions aligned with future policy and thus are not exposed to surprise losses. Such communication and “forward guidance” is consistent with the behavior of the Fed as discussed in previous sections. Speeches, transcripts (e.g., Governor Fischer’s 2014 remarks quoted at the start of the paper), and official communications such as financial stability reports repeatedly emphasize the risks of market positions being misaligned with policy decisions. The example discussed earlier illustrates this precisely: the Fed explicitly warned that higher interest rates could cause losses for banks exposed to commercial real estate.

Beyond such qualitative evidence, there is direct empirical research consistent with the role of communication in our model. [Cieslak et al. \(2019\)](#) document systematic informal communication between the Fed and the financial sector, and show that this communication conveys news about monetary policy *prior* to a decision being made. This kind of pre-decision guidance, which shapes markets before the policy choice itself, is the behavior our model is designed to capture.

Given the central role of such pre-decision communication, it is worth noting that our cheap-talk results rely on an important assumption about its nature. If communication merely disclosed hard economic data, the central bank could overcome credibility problems by publishing this information directly. However, in practice, forward guidance often involves interpretive or judgmental information that cannot be fully verified. This inability to fully verify may stem from differences in interpretation of the same data, from preference shocks among policymakers, or from the complexity of institutional decision-making. In such cases, information asymmetries cannot be eliminated through disclosure, making communication inherently imperfect and prone to cheap-talk frictions.

References

- Abadi, J., Brunnermeier, M., Koby, Y., 2023. The reversal interest rate. *Am. Econ. Rev.* 113 (8), 2084–2120. <https://doi.org/10.1257/aer.20190150>
- Afonso, G., Santos, J. A.C., Traina, J., 2014. Do “too-big-to-fail” banks take on more risk? *Econ. Policy Rev.* 20 (2), 41–58.
- Alessandri, P., Nelson, B.D., 2015. Simple banking: profitability and the yield curve. *J. Money Credit Bank.* 47 (1), 143–175. <https://doi.org/10.1111/jmcb.12172>
- Altavilla, C., Brugnolini, L., Gürkaynak, R.S., Motto, R., Ragusa, G., 2019. Measuring Euro area monetary policy. *J. Monet. Econ.* 108, 162–179. <https://doi.org/10.1016/j.jmoneco.2019.08.016>
- Begenau, J., Piazzesi, M., Schneider, M., 2025. Banks’ Risk Exposure. Unpublished working paper.
- Blinder, A.S., Ehrmann, M., Fratzscher, M., De Haan, J., Jansen, D.-J., 2008. Central bank communication and monetary policy: a survey of theory and evidence. *J. Econ. Lit.* 46 (4), 910–945. <https://doi.org/10.1257/jel.46.4.910>
- Boleslavsky, R., Kelly, D.L., Taylor, C.R., 2017. Selloffs, bailouts, and feedback: can asset markets inform policy? *J. Econ. Theory* 169, 294–343.
- Bond, P., Goldstein, I., 2015. Government intervention and information aggregation by prices. *J. Finance* 70 (6), 2777–2812.
- Bradley, D., Finer, D.A., Gustafson, M., Williams, J., 2020. When Bankers Go to Jail: Insights into Fed-Bank Interactions from Taxi Data. Unpublished working paper.
- Caballero, R., Simsek, A., 2020. Prudential Monetary Policy. Working Paper Series No. 25977. National Bureau of Economic Research.
- Caballero, R., Simsek, A., 2022. A Monetary Policy Asset Pricing Model. Working Paper Series No. 30132. National Bureau of Economic Research.
- Calvo, D., Crisanto, J.C., Hohl, S., Gutiérrez, O.P., 2018. Financial supervisory architecture: what has changed after the crisis? *FSI Insights Policy Implement.* 8, 1–33.
- Cieslak, A., McMahon, M., 2023. Tough Talk: The Fed and the Risk Premium. Unpublished working paper <https://api.semanticscholar.org/CorpusID:261622385>.
- Cieslak, A., Morse, A., Vissing-Jørgensen, A., 2019. Stock returns over the FOMC cycle. *J. Finance* 74 (5), 2201–2248.
- Cieslak, A., Vissing-Jørgensen, A., 2021. The economics of the fed put. *Rev. Financ. Stud.* 34 (9), 4045–4089. <https://doi.org/10.1093/rfs/hhaa116>
- Corbae, D., D’Erasmo, P., 2021. Capital buffers in a quantitative model of banking industry dynamics. *Econometrica* 89 (6), 2975–3023. <https://doi.org/10.3982/ecta16930>
- Crawford, V.P., Sobel, J., 1982. Strategic information transmission. *Econometrica* 50 (6), 1431–1451. <http://www.jstor.org/stable/1913390>.
- English, W.B., Van den Heuvel, S.J., Zakrajšek, E., 2018. Interest rate risk and bank equity valuations. *J. Monet. Econ.* 98, 80–97. <https://doi.org/10.1016/j.jmoneco.2018.04.010>
- Faust, J., 2016. Oh What a Tangled Web We Have: Monetary Policy Transparency in Divisive Times. Unpublished working paper.
- Flannery, M.J., James, C.M., 1984. The effect of interest rate changes on the common stock returns of financial institutions. *J. Finance* 39 (4), 1141–1153. <https://doi.org/10.1111/j.1540-6261.1984.tb03898.x>
- Gabaix, X., 2011. The granular origins of aggregate fluctuations. *Econometrica* 79 (3), 733–772.
- Galí, J., 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton University Press.
- Hansen, S., McMahon, M., 2016. Shocking language: understanding the macroeconomic effects of central bank communication. *J. Int. Econ.* 99, 114–133.
- Hansen, S., McMahon, M., Tong, M., 2019. The long-run information effect of central bank communication. *J. Monet. Econ.* 108, 185–202. <https://doi.org/10.1016/j.jmoneco.2019.09.002>
- Jeanneau, S., 2014. Financial stability objectives and arrangements—What’s new? BIS Papers No 76.

⁵² One can consider the exposure arising from maturity transformation as being determined by β while the exposure from all other investments as being determined by the portfolio choice x_i . Alternatively, one can consider β as capturing a bank’s exposure to credit risk via changes in interest rates while x_i captures its direct exposure to interest rate risk. In both cases, $\beta > 0$ implies higher rates negatively impact banks.

- Jiang, E.X., Matvos, G., Piskorski, T., Seru, A., 2024. Monetary tightening and U.S. bank fragility in 2023: mark-to-market losses and uninsured depositor runs? *J. Financ. Econ.* 159, 103899.
- Lukyanov, G., Shamruk, K., Su, T., Wakrim, A., 2022. Public communication with externalities. *Games Econ. Behav.* 136, 177–196. <https://doi.org/10.1016/j.geb.2022.09.003>
- McKay, A., Nakamura, E., Steinsson, J., 2016. The power of forward guidance revisited. *Am. Econ. Rev.* 106 (10), 3133–58. <https://doi.org/10.1257/aer.20150063>
- Metrick, A., 2024. The failure of silicon valley bank and the panic of 2023. *J. Econ. Perspect.* 38 (1), 133–152.
- Morris, S., 2001. Political correctness. *J. Polit. Econ.* 109 (2), 231–265.
- Morse, A., Vissing-Jørgensen, A., 2020. Information Transmission from the Federal Reserve to the Stock Market: Evidence from Governors' Calendars. Unpublished working paper.
- Ottaviani, M., Sørensen, P.N., 2006. Reputational cheap talk. *RAND J. Econ.* 37 (1), 155–175. <https://doi.org/10.1111/j.1756-2171.2006.tb00010.x>
- Peek, J., Rosengren, E.S., Tootell, G. M.B., 2015. Should U.S. monetary policy have a ternary mandate? In: Federal Reserve Bank of Boston 59th Economic Conference: Macroprudential Monetary Policy.
- Antić, N., Persico, N., 2020. Cheap talk with endogenous conflict of interest. *Econometrica* 88 (6), 2663–2695. <https://doi.org/10.3982/ECTA14953>
- Rogoff, K., 1985. The optimal degree of commitment to an intermediate monetary target. *Q. J. Econ.* 100 (4), 1169–1189.
- Shapiro, A.H., Wilson, D.J., 2022. Taking the fed at its word: a new approach to estimating central bank objectives using text analysis. *Rev. Econ. Stud.* 89 (5), 2768–2805.
- Smets, F., 2014. Financial stability and monetary policy: how closely interlinked? *Int. J. Cent. Bank.* 10 (2), 263–300.
- Stein, J.C., 1989. Cheap talk and the fed: a theory of imprecise policy announcements. *Am. Econ. Rev.* 79 (1), 32–42. <http://www.jstor.org/stable/1804771>.
- Stein, J.C., Sunderam, A., 2018. The fed, the bond market, and gradualism in monetary policy. *J. Finance* 73 (3), 1015–1060.
- Sundaresan, S., Wang, Z., 2023. Strategic bank liability structure under capital requirements. *Manag. Sci.* 69 (10), 6349–6368. <https://doi.org/10.1287/mnsc.2022.4570>
- Taylor, J.B., 1993. Discretion versus policy rules in practice. In: *Carnegie-Rochester Conference Series on Public Policy*. vol. 39. Elsevier, pp. 195–214.
- Tinbergen, J., 1952. *On the Theory of Economic Policy*. Amsterdam: North-Holland.
- Uppal, A., 2025. Do Higher Interest Rates Make The Banking System Safer? Evidence From Bank Leverage. Unpublished working paper.
- Vissing-Jørgensen, A., 2020. Central Banking with Many Voices: The Communications Arms Race. Unpublished working paper.
- Woodford, M., 2012. Inflation Targeting and Financial Stability. Working Paper Series No. 17967. National Bureau of Economic Research.
- Zimmermann, K., 2019. Monetary Policy and Bank Profitability, 1870–2015. Unpublished working paper.