

Talking Over Time

Dynamic Central Bank Communication

Laura Gáti*

Abstract

This paper studies the optimal dynamic communication strategy of central banks using a Bayesian persuasion game framework. In a dynamic environment, financial market participants and the general public have misaligned interests because the present and future have different relevance in their optimization problems, leading to a novel tradeoff for the monetary authority. Compared to the static benchmark, I show that the central bank's optimal dynamic communication policy should put a higher weight on talking about the present state than the future. In addition, the central bank should strategically send more noisy signals than in the static benchmark.

Keywords: Dynamic communication, Bayesian persuasion, Delphic forward guidance

JEL Codes: E58, E71, D83, C73

*European Central Bank, Directorate General Research, Sonnemannstraße 20, 60314 Frankfurt am Main, Germany. Email: laura.veronika.gati@ecb.europa.eu. I thank the editor, Sanjay Chugh, and one anonymous referee for helpful advice and suggestions. I also thank Ryan Chahrour, Susanto Basu and Peter Ireland, as well as Sylvérie Herbert, Bartosz Maćkowiak and Francesco Zanetti. The views expressed herein are my own and do not necessarily reflect those of the ECB or the Eurosystem.

1 INTRODUCTION

While financial market participants listen intently to central bank announcements, this cannot be said of the general public. As [Coibion et al. \(2020\)](#), [Binder and Kim \(Forthcoming\)](#) and many others show, most firms and households do not incorporate information provided by the central bank into their expectations. A dynamic economic environment, however, introduces a discrepancy between the interests of Wall Street and Main Street: the future economy is relevant for the returns of the former, while financial constraints and discounting of the future limit the latter's focus to the present. Talking over time thus presents a central bank with a novel tradeoff: how should it communicate so as to serve the interest of the general public while knowing that its messages only get through to investors?

The novelty of this paper is to characterize a dynamic central bank communication policy that responds optimally to this tradeoff. To highlight the novel tradeoffs coming from dynamics, I contrast a dynamic communication problem with its static analogue. The static and dynamic models involve a Bayesian persuasion-type communication game between a central bank (CB) and the financial market (FM) about two economic fundamentals that represent the current and future stance of the business cycle. Main Street cares about current employment and is thus concerned with today's business cycle, while Wall Street sets today's investment choices with an eye toward future profits. Capturing the idea that the central bank has a mandate to stabilize current inflation and employment, I assume that the central bank is concerned with the current business cycle, as opposed to the financial market that is maximizing future profits. At the same time, the monetary authority takes into account that a only subset of the public listens to its messages - Wall Street.

In this environment, the central bank sends the financial market a noisy signal which is a weighted sum of the present and the future stance of the business cycle. Importantly, the static and dynamic models are identical up to the correlation structure between the two states. This allows me to isolate the role of dynamics for optimal communication design. In particular, I investigate two ways in which the central bank can address the tradeoff between the interests of Wall Street and Main Street. I first ask how strongly the central bank should weight the current against the future in its signal - a dimension of communication I refer to as “targetedness.” Secondly, I explore whether the optimal precision of the dynamic signal differs from that of the static one. In other words, does dynamic communication involve a different amount of noise than static communication?

There are two key findings. First, the central bank’s signal is always more targeted toward current conditions in the dynamic model than in the static one. In other words, in a dynamic world, the central bank prefers to talk more about the present than the future relative to the static environment. On the one hand, this is because a *temporal* correlation between output today and output tomorrow renders the two states more distinct than variables that are correlated in the cross-section. On the other, in a dynamic world, the financial market uses all the available information to learn about the evolution of the business cycle. Therefore the central bank has to weight the signal more heavily towards the present to skew the available information over time in the right direction.

Second, in the dynamic model, the central bank optimally communicates more noisily than in the static model. The reason behind this is that because information is carried over from one period to the next, providing too much information today low-

ers the central bank's ability to persuade the financial market tomorrow. Intuitively, communicating very precisely today makes the financial market too confident in its beliefs the next period, making it very hard for the central bank to convince it otherwise. Optimal precision is thus lower in the dynamic problem than in the static one as the central bank smoothes the information it provides the financial market over time.

The dynamic communication policy studied in this paper yields testable implications regarding the everyday communication decisions central bankers make. First, the paper provides conditions for when the central bank should give investors precise information, and when it is preferable to communicate in noisy "Fedspeak" that leaves some residual uncertainty around the current and future economic outlook. Second, it also suggests that it is in the interest of the general public that the central bank should talk more about the current business cycle in order to align the actions of Wall Street more closely with the needs of Main Street.

The paper is outlined as follows. Section 2 develops the dynamic communication game, while Section 3 outlines the static analogue. Sections 4 and 5 describe the targetness and precision dimensions of the optimal communication policy respectively. Section 6 concludes.

1.1 Related literature

The paper is related to three strands of literature. First, my main point of reference is the global games literature in the vein of [Morris and Shin \(2002\)](#), [Svensson \(2006\)](#), [Angeletos and Pavan \(2007\)](#) and [Hellwig and Veldkamp \(2009\)](#). Several papers have used this literature as a starting point to study particular dimensions of communication. [Chahrour \(2014\)](#), for example, uses the rational inattention literature à la [Sims](#)

(2003) to investigate the optimal amount of central bank communication.

To my knowledge, only few papers consider the time dimension in some form. One is [Gaballo \(2016\)](#), who analyzes Delphic forward guidance, the communication of the central bank about its own information set, in an overlapping generations (OLG) global games model. However, Gaballo only introduces dynamics in the evolution of the fundamental; central bank communication in his model is simply revealing information about the central bank's one-period-ahead forecast. In this sense, central bank communication in [Gaballo \(2016\)](#) maps one-to-one to the standard static problem of [Morris and Shin \(2002\)](#). [Reis \(2011\)](#) analyzes the optimal timing decision of an authority that knows about a future policy change the public is unaware of. The problem of the authority is to decide when to publicly announce the future policy change. [Hansen and McMahon \(2016\)](#) instead examine how the voting behavior of monetary policy members changes over time as a result of dynamically changing signaling incentives.

Secondly, my work is related to the Bayesian persuasion literature in the wake of [Kamenica and Gentzkow \(2011\)](#). Bayesian persuasion is a signaling game where a sender designs his communication so as to persuade a receiver to take a sender-preferred action. Such a setting has been widely adopted in many applications such as stress tests ([Goldstein and Leitner 2018](#) and [Inostroza and Pavan 2022](#)) or even central bank communication ([Ko 2019](#) and [Herbert 2021](#)). Importantly, the sender has access to a full-commitment technology, and thus persuasion in this setting works through signal design, not through untruthfulness. To the best of my knowledge, my paper is the first to present a dynamic extension of the Bayesian persuasion setting, applied to central bank communication.

Viewed from the lens of commitment, my work ties in closer with the Bayesian persuasion literature than with the macroeconomic literature on discretionary monetary policy and cheap talk such as [Barro and Gordon \(1983\)](#), [Moscarini \(2007\)](#) and [Frankel and Kartik \(2018\)](#). While the discretionary monetary policy literature has an explicit concern for dynamics, the central question in this literature is not how to provide information to the public, but how to design central bank action over time.

My paper highlights the fact that dynamics introduces a distinction between various subsets of the private sector: those who listen to central bank announcements versus those whose welfare monetary policy is mandated to support. This is a point of connection with a third literature: the literature on whether firms and households pay attention to monetary policy communication. In most empirical papers, the answer is in general no ([Binder 2017](#), [Coibion et al. 2018](#), [Candia et al. 2021](#) and [Pfäuti 2021](#)). In emphasizing the varying levels of attentiveness among the public, both this literature and my paper motivate the question whether it is desirable from a welfare perspective to increase the public's attentiveness to central bank statements. Put more starkly, this raises the possibility that in recent decades, central banks might have been too successful at anchoring the private sector's expectations.

Lastly, I focus on central bank communication about economic fundamentals, not about future policy. To use the terminology of [Campbell et al. \(2012\)](#), I thus model Delphic, and not Odyssean, forward guidance. This way I can isolate the effect of pure communication instead of considering the implications of the central bank tying its hands concerning the evolution of future policy.

2 THE DYNAMIC MODEL

Consider an economy where an economic fundamental, θ , evolves dynamically over time. Suppose the fundamental evolves according to

$$\theta_{t+1} = \rho\theta_t + \varepsilon_{t+1}, \quad \text{with } \varepsilon_{t+1} \sim \mathcal{N}(0, \sigma_\varepsilon^2), \quad \text{and } \sigma_\varepsilon^2 = 1 - \rho^2. \quad (1)$$

The interpretation for the fundamental I have in mind is the stance of the business cycle. For example, think of θ_t as the output gap today, and of θ_{t+1} as the output gap tomorrow. To keep things simple, I will thus refer to θ as “output.” Thus ρ captures temporal correlation between today’s output and tomorrow’s output, and I have set the variance of the innovation, σ_ε^2 , so that output has unit variance.

In this setting, consider the problem of a central bank that seeks to communicate with the public about the current and future business cycle. First of all, the central bank needs to come to terms with the idea that its communication will be received by financial market participants only (Wall Street), as households living on Main Street tend not to participate in financial markets and not to pay attention to central bank communication (Van Rooij et al. 2011, Kumar et al. 2015, Candia et al. 2021, Binder 2017, Coibion et al. 2018). Second, the central bank needs to communicate with the financial market, while keeping its mandate of price stability and full employment in mind. In other words, the central bank faces the problem of talking to Wall Street, yet representing Main Street.

The key point of the paper is to suggest that in a dynamic environment, this creates a novel tradeoff for the central bank. This tradeoff stems from the fact that while the central bank’s mandate leads it to want to stabilize the *current* business cycle, financial market participants are inherently concerned with the *future* business cycle, as that is

what is relevant for their profits.

To appreciate the latter point, think of an investor's optimal choice of investment. As spelled out in detail in Appendix A, a standard Q-theory of optimal investment à la [Tobin \(1969\)](#) suggests that current investment is chosen to maximize the future stream of profits. To select the right level of investment, investors thus need to glean information about future profitability.

The central bank, instead, manages demand to stabilize the current business cycle. Since investment is a component of aggregate demand, high current investment boosts demand, leading to higher production and thus higher employment today. The dynamic nature of the problem thus introduces a misalignment in preferences between the central bank and the financial sector.

I model this dynamic communication problem as a dynamic extension to a Bayesian persuasion game à la [Kamenica and Gentzkow \(2011\)](#). In particular, this is a communication game between a central bank (CB) and the financial market (FM) in which the FM chooses an action that I will think of as investment, I_t , and the CB provides information about the current and future business cycle to the FM. Like [Kamenica and Gentzkow \(2011\)](#), I assume that lying is not allowed, so that all signals are truthful.¹

To capture the preference misalignment inherent to the dynamic setup, the payoffs of the two players are

$$\mathcal{L}^{FM,dynamic} = \mathbb{E}_t^{FM} (I_t - \theta_{t+1})^2, \quad (2)$$

$$\mathcal{L}^{CB,dynamic} = \mathbb{E}_0^{CB} \sum_{t=0}^{\infty} \beta^t (I_t - b\theta_t)^2, \quad (3)$$

with $b \in [0, \infty)$ denoting the weight which the CB places on tracking θ_t , and $\beta \in (0, 1)$ being the central bank's discount factor.

The “persuasion” element of the model, following [Kamenica and Gentzkow \(2011\)](#), is that the CB designs its signals in a way to induce the FM to choose the action that maximizes the CB’s expected payoff. This is the “talk to Wall Street, represent Main Street” element of the model. It captures the idea that the CB needs to communicate with the FM so as to make it choose investment to the benefit of the general public, that is, to stabilize the current business cycle. At the same time, the CB talks to the FM with the understanding that the FM cares about future profitability.

For simplicity, I assume the CB observes the full history of states perfectly, including the one-period-ahead state, while the FM’s information set only includes the history of the signal. Formally, the information sets of the two players in a particular period are

$$\mathcal{I}_t^{CB} = \{\theta_{t+1}, \theta_t, \dots, \theta_0\}, \quad \mathcal{I}_t^{FM} = \{s_t, s_{t-1}, \dots, s_0\}. \quad (4)$$

To capture communication simultaneously about the present and the future, I assume that the central bank sends a signal to the financial market of the following form

$$s_t = \theta_t + \frac{1}{\psi} \theta_{t+1} + v_t, \quad v_t \sim \mathcal{N}(0, \sigma_v^2), \quad (5)$$

where v is a noise term with variance σ_v^2 , and ψ is a weight on current output.

There are two dimensions to the CB’s communication problem embedded in this signal structure. The first one, ψ , represents how strongly the CB weights today’s output relative to tomorrow’s. I refer to this dimension of dynamic communication as “targetedness.” When $\psi > 1$, the signal is targeted toward current output, with the limit of $\psi \rightarrow \infty$. By contrast, when $\psi < 1$, the signal is targeted toward future output, with a limiting case of $\psi \rightarrow 0$. In the intermediate case, when $\psi = 1$, the weight on both states is equal and thus the signal is not targeted.

The second dimension is σ_v , corresponding to the precision of communication. A low σ_v renders a signal with given targetedness more precise, while a high σ_v renders it more noisy. Most papers in the central bank communication literature focus on this dimension of communication, because asking how precisely the CB should communicate is a well-defined question also in a static communication setting. The targetedness dimension, instead, is a fundamentally new dimension to dynamic communication.²

Admittedly, the signal structure in Equation (5) is somewhat restrictive: it says that the CB makes statements about the economy that mix information about current and future conditions, and that such statements can be anywhere between very concrete and relatively vague. In fact, this signal structure is designed to mimic the way central bank statements about the economic outlook are set up. Consider for example the following excerpt from the Federal Open Market Committee (FOMC) Statement of the Federal Reserve from December 15, 2021:

With progress on vaccinations and strong policy support, indicators of economic activity and employment have continued to strengthen. The sectors most adversely affected by the pandemic have improved in recent months but continue to be affected by COVID-19. Job gains have been solid in recent months, and the unemployment rate has declined substantially. Supply and demand imbalances related to the pandemic and the reopening of the economy have continued to contribute to elevated levels of inflation. Overall financial conditions remain accommodative, in part reflecting policy measures to support the economy and the flow of credit to U.S. households and businesses.

The path of the economy continues to depend on the course of the virus. Progress on vaccinations and an easing of supply constraints are expected to support continued gains in economic activity and employment as well as a reduction in inflation. Risks to the economic outlook remain, including from new variants of the virus. (Federal Reserve, December 15, 2021)

The first paragraph is dedicated mainly to the current economic stance, while the second focuses first and foremost on the future outlook. Thus the statement contains information about both present and future. With four long sentences in the first paragraph compared to three shorter ones in the second, the statement also appears to place a higher weight on the current circumstances. In the language of Equation (5), the statement is more targeted toward the present.

The boundaries between present and future are blurred, however. For many of the sentences, it is not straightforward to make out what time horizon they pertain to. A good example of this is the second sentence of the second paragraph. Are “[progress] on vaccinations and easing of supply constraints” happening now, leading to “continued gains in economic activity” in the future? Or is the FOMC expecting “[progress] on vaccinations and easing of supply constraints” moving forward? Similarly, the last sentence of the paragraph is also somewhat ambiguous as to whether the current or the future economic outlook is risky (or both).

Thus Equation (5) offers a simple, reduced-form way of capturing the way central bank statements provide information about the current and future stance of the economy. Furthermore, it also aligns with the form of the signal that Main Street would choose if it were selecting the signal subject to rational inattention.³

In this setting, the communication policy of the CB consists of choosing at the

beginning of time the form of the signal to be sent to the FM in each period. In other words, in period 0, the CB chooses ψ and σ_v and sticks to this communication policy forever.

Noting that the FM's problem implies that the FM's investment choice corresponds to its expectation of tomorrow's output, the CB's problem can be stated as minimizing its loss function subject to optimal FM inference. Since the signal is linear with Gaussian noise, the optimal FM forecast is given by the Kalman filter formula

$$\theta_{t+1|t} = m_1\theta_{t|t-1} + m_2\theta_t + m_3\theta_{t+1} + m_4v_t, \quad (6)$$

where $\theta_{t+1|t} := \mathbb{E}_t[\theta_{t+1}|\mathcal{I}_t^{FM}]$, and $m_i, i = 1, \dots, 4$ are given by the Kalman filter, as derived in Appendix B. Then the CB's problem is

$$\min_{\psi, \sigma_v} \mathbb{E}_0^{CB} \sum_{t=0}^{\infty} \beta^t (\theta_{t+1|t} - b\theta_t)^2 \quad \text{s.t.} \quad \theta_{t+1|t} = m_1\theta_{t|t-1} + m_2\theta_t + m_3\theta_{t+1} + m_4v_t. \quad (7)$$

DEFINITION 1. Let $\mu_X(x)$ be the probability distribution of a variable X induced by the FM's beliefs. A Perfect Bayesian Equilibrium is an action rule I_t , belief system μ and a communication policy (ψ^*, σ_v^*) such that

- $I_t = \arg \min \mathcal{L}_t^{FM}(I_t, \theta_{t+1}) \quad \text{s.t.} \quad \mathbb{E}_t^{FM}(\theta_{t+1}|s_t),$
- $(\psi^*, \sigma_v^*) = \arg \min \mathcal{L}^{CB}(\{I_t, \theta_t\}_{t=0}^{\infty}) \quad \text{s.t.} \quad \mathbb{E}_t^{FM}(\theta_{t+1}|s_t) \quad \text{and}$
 $s_t = \theta_t + \frac{1}{\psi}\theta_{t+1} + v_t \quad \text{with} \quad v_t \sim \mathcal{N}(0, \sigma_v^2),$
- FM beliefs \mathbb{E}_t^{FM} come from $\mu \forall t$, and for the dynamic model, μ is consistent with Bayes' rule:

$$\mu_{\Theta|S=s}(\theta) = \frac{\mu_{S|\Theta=\theta}(s)\mu_{\Theta}(\theta)}{\mu_S(s)}.$$

3 THE STATIC BENCHMARK

I now provide a static model that is as close of an analogue to the dynamic model as possible. Let θ_1, θ_2 denote two fundamentals that are correlated in the cross-section with the following properties:

$$(\theta_1, \theta_2) \sim \mathcal{N}(0, \mathbf{V}), \quad \text{with} \quad \mathbf{V} = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}. \quad (8)$$

Like in the dynamic world, the two states have unit variance and their correlation is given by the parameter ρ . The only difference to the dynamic model is that here ρ is not a temporal correlation, but a cross-sectional one. Thus, stricly speaking, the interpretation of “current output” and “future output” is no longer valid. Therefore the temporal misalignment in preferences between the FM and the CB is also gone. In order to have the same payoff structure as in the dynamic model, then, let us imagine a situation in which for some reason, the CB targets some other aspect of the economic environment than the FM does. That gives rise to analogous payoffs as in the dynamic setting, where

$$\mathcal{L}^{FM,static} = \mathbb{E}^{FM}(I - \theta_2)^2, \quad (9)$$

$$\mathcal{L}^{CB,static} = \mathbb{E}^{CB}(I - b\theta_1)^2. \quad (10)$$

For ease of comparison with the dynamic model, I will abuse terminology a little and refer to θ_1 as “current output,” and to θ_2 as “future output.”

Again in analogy with the dynamic setting, I assume the CB knows more than the FM does

$$\mathcal{I}^{CB} = \{\theta_1, \theta_2\}, \quad \mathcal{I}^{FM} = \{s\}, \quad (11)$$

and, finally, I constrain the CB's signal to be identical to that in the dynamic model:

$$s = \theta_1 + \frac{1}{\psi}\theta_2 + v \quad v \sim \mathcal{N}(0, \sigma_v^2). \quad (12)$$

Setting up the CB's problem follows the same steps as in the dynamic model. The optimal FM forecast in the static model is given by

$$\theta_{2|s} = \phi s, \quad (13)$$

where $\theta_{2|s} := \mathbb{E}^{FM}[\theta_2|s]$, while the optimal projection is given by the ordinary-least-squares formula $\phi = \frac{Cov(\theta_2, s)}{Var(s)}$. Then the statement of the CB's problem is

$$\min_{\psi, \sigma_v} \mathbb{E}^{CB}(\theta_{2|s} - b\theta_1)^2 \quad \text{s.t.} \quad \theta_{2|s} = \phi s. \quad (14)$$

4 OPTIMAL TARGETEDNESS

Integrating out the states from the central bank's loss function, one obtains \mathcal{V} , the CB's expected loss function. In the static model, this takes the following form

$$\mathcal{V}^{static} = Var(\theta_{2|s}) - 2bCov(\theta_{2|s}, \theta_1) + b^2, \quad (15)$$

where I have utilized the fact that $Var(\theta_1) = 1$. Analogously, the expected loss function in the dynamic model is given by

$$\mathcal{V}^{dynamic} = \frac{1}{1-\beta} \left(Var(\theta_{t+1|t}) - 2bCov(\theta_{t+1|t}, \theta_t) + b^2 \right), \quad (16)$$

where I have similarly made use of the fact that $Var(\theta_t) = 1$. Since Equation (16) is quite involved, I now present the analytical solution for ψ in the static model, and contrast it with a numerical solution for the dynamic model. (See Appendix C for the full analytical expressions.)

PROPOSITION 1. *In the static model, optimal targetedness is given by*

$$\psi_{static}^* = \frac{\sqrt{\gamma} - \rho (2b\sigma_v^2 + \rho) + \sigma_v^2 + 1}{2(b\rho^2 + b - \rho)\sigma_v^2 - 2(\rho^2 - 1)(b - \rho)}, \quad (17)$$

$$\gamma := (-\rho^2 + \sigma_v^2 + 1) \left((1 - 2b\rho)^2 \sigma_v^2 - (\rho^2 - 1)(4b(b - \rho) + 1) \right). \quad (18)$$

Figure 1 illustrates the optimal ψ as a function of ρ for a benchmark calibration of $b = 1$ and $\beta = 0.99$. I also set $\sigma_v = 1$ throughout this section. The solid line pertains to the static, the dotted line to the dynamic model. Before turning to the differences in optimal targetedness between the static and dynamic models, let me first characterize the commonalities between the two. This is formalized in the following proposition.

PROPOSITION 2. *Shared features of the optimal targetedness policy*

Both in the static and dynamic model, $\psi^(\rho)$ has the following properties:*

- $\psi^* \geq 0 \forall \rho$.
- When $\rho \rightarrow -1$, $\psi^* \rightarrow 1$ from above.
- When $\rho \rightarrow 1$, $\psi^* \rightarrow 0$ from above.
- $\exists \rho^a$ such that for $\rho < \rho^a$, ψ^* is increasing in ρ , but for $\rho > \rho^a$, ψ^* is decreasing in ρ .
- $\exists \rho^s > \rho^a$ such that $\psi^*(\rho^s) = 1$, and for a small $\epsilon > 0$, $\psi^*(\rho^s - \epsilon) > 1 > \psi^*(\rho^s + \epsilon)$.

I refer to ρ^a and ρ^s as the alignment and switching threshold respectively.

Figure 2 explains Proposition 2 by replotting Figure 1 schematically, splitting the ρ -space into four different regions. In Region I, shown in gray shading, $\rho < 0$, rendering the preferences of the FM and the CB severely misaligned. This means that while the CB would like the signal to be targeted toward today's output, it cannot make ψ too

large, because a negative correlation between today's and tomorrow's output means that the FM would set investment exactly the opposite way. One can see this on Figure 3, which plots the FM's responsiveness to the CB's signal for various values of ρ . For example when $\rho = -0.99$, as in the leftmost panel of Figure 3, the response coefficient of the FM (κ in the dynamic, ϕ in the static model) has the opposite sign to ψ . In the limit, when $\rho = -1$, the best the CB can do is to send a perfectly confounding signal, $\psi = 1$, resulting in the FM ignoring the signal completely. As ρ approaches zero from below, however, the constraint on how strongly the CB can target current output loosens. Therefore, throughout Region I, ψ^* is increasing in ρ .

Once $\rho \geq 0$, we enter Region II, the region on Figure 2 shaded with diagonal lines. Here, since $\rho \geq 0$, there is no constraint on the CB in raising ψ and thus targeting current output more strongly. And it is desirable to raise ψ for the CB because ρ is sufficiently close to zero that the CB's and the FM's preferences are still misaligned. However, the CB cannot raise ψ indefinitely, since responsiveness falls the higher ψ is. For example, when $\rho = 0$, the middle panel of Figure 3 shows that $\partial\kappa/\partial\psi < 0$ for any $\psi > 1$.

The region shaded with dots, Region III, begins once $\rho > \rho^a$. At this point, ρ is high enough for the two players' preferences to become more aligned. Thus the CB becomes interested in raising the FM's responsiveness to tomorrow's output instead of today's. As shown in Appendix B, the FM's responsiveness to tomorrow's output is κ/ψ in the dynamic, while it is ϕ/ψ in the static model. Thus in both models the CB can raise this by lowering ψ . A lower ψ not only lowers the denominator, but also raises the numerator, as κ and ϕ are decreasing in ψ if $\psi > 1$ and ρ is positive and not too close to 1 (see third and fourth panels of Figure 3). Thus in this region, the CB

maximizes responsiveness to a weighted average of current and future output.

Once $\rho > \rho^s$, the two states are so strongly correlated that the CB's and the FM's preferences are strongly aligned. In this region, held in white and marked Region IV on Figure 2, the CB's primary concern is the noise in the signal. Since ρ is approaching 1, the CB is indifferent to whether the FM responds to today's or tomorrow's output, and the only thing that can get in the way of that is the noise. This way, the FM's responsiveness to future output increases, while its responsiveness to the signal in general, and thus to noise, decreases. Thus for a sufficiently high ρ , the CB switches the targetedness of the signal from today's to tomorrow's output.

In this manner, ρ , the correlation between the two states, determines the optimal choice of targetedness in both the static and the dynamic model. In the limit, severe misalignment results in a signal that does not target any of the two states, while the opposite extreme of perfect alignment leads the CB to target the FM's preferred state infinitely. Between the limiting cases, the CB balances the need for a high level of targeting current output if alignment is low against the shrinking responsiveness of the FM to the signal if it is too targeted toward today's output.

4.1 More targeted toward the present

As Figure 1 indicates, ψ^* is quantitatively different across the two models. Moreover, the threshold ρ^s after which the CB only aims to drive down the responsiveness to noise arrives later in the dynamic model than in the static one. The next proposition states these results.

PROPOSITION 3. *Higher weight on the central bank's target in the dynamic model*

$$\psi_{dynamic}^* \geq \psi_{static}^* \quad \forall \rho. \quad (19)$$

In the dynamic model, the central bank finds it optimal to target today's state more than it would in the static model for any level of alignment.

COROLLARY 3.1.

$$\rho_{dynamic}^s > \rho_{static}^s. \quad (20)$$

The switching threshold is higher in the dynamic model than in the static model.

Figure 4 illustrates Proposition 3. It plots expected payoffs of the CB as a function of targetedness across the static (solid line) and dynamic models (dotted line). One notices that for any value of ρ , that is on all panels of the figure, ψ^* is weakly larger in the dynamic model than in the static one.

Why is it optimal for the central bank to target current output more in the dynamic model? To understand the reason, it is helpful to consider some measure of how informative a given signal is about current and future output. To do so, let $\pi(\theta_T)$ and $p(\theta_T, s_t)$ denote prior and posterior variances of θ_T , where $T = t, t + 1$,

$$\pi(\theta_T) := \mathbb{E}[(\theta_T - \theta_{T|t-1})^2], \quad (21)$$

$$p(\theta_T, s_t) := \mathbb{E}[(\theta_T - \theta_{T|t})^2]. \quad (22)$$

Let us also define

$$I(\theta_T, s_t) := \pi(\theta_T) - p(\theta_T, s_t), \quad (23)$$

as the decrease in uncertainty in θ_T given s_t . I will refer to $I(\theta_T, s_t)$ as the “informativeness” of signal s_t about θ_T , as it captures how much the signal contributed to reducing the FM’s uncertainty about output at time T .⁴

Figure 5 plots the informativeness of today's signal about today's output (solid line) and tomorrow's (dotted line), both as a function of targetedness, ψ .⁵ The first thing one observes is that in general, the solid and dotted lines go in opposite ways. This is an intuitive feature which comes from the fact that higher targetedness toward current output (higher ψ) renders a signal more informative about today's output, and less informative about tomorrow's.

Notice, however, that for the static model, this pattern disappears when the two states are highly correlated, so that ρ is close to -1 or 1 (the far left and right panels). This makes sense, given that a high targetedness in either direction is informative about both current and future output when the two are very closely related. For example, when $\rho = 1$, this means that in the static model, the two states are exactly equal. Because the output process is an AR(1) in the dynamic model, however, a perfect correlation of $\rho = 1$ does not render today's and tomorrow's output exactly equal. Therefore, in the dynamic setting, this introduces the need for the CB to always target current output more in the signal, because even in a $\rho = 1$ situation, there is something to lose for the CB if the FM focuses on the future. This explains why for low ρ -values, the central bank's loss is higher in the dynamic than in the static model (see the left panels of Figure 4).

But the fact that the AR(1) process for output always introduces a meaningful distinction between today and tomorrow also shows up in another way. For example, look at the solid lines on the second from the right panel of Figure 5, when $\rho = 0.5$. What becomes apparent is that the informativeness of the signal about today's output is nonmonotonic in the static model, but stays monotonic in the dynamic model. In other words, for any correlation, in the static model there comes a point at which

increasing the targetedness of the signal further toward current output *decreases* the informativeness of the signal about current output.

In the dynamic model, this feature is absent because the additional distinction between present and future gives the CB the opportunity to provide information about today's output that is not relevant for tomorrow's. A cross-sectional correlation, as in the static model, precludes this possibility because any information about today is by construction relevant for tomorrow. Therefore, in a dynamic world, the CB exploits this possibility to provide more information about current conditions to push the FM's investment choice in the right direction. This also explains why the dynamic expected loss in Figure 4 is lower for sufficiently high (ψ, ρ) pairs.

4.2 *More information in the dynamic world*

Is the autoregressive process for output the only explanation for a stronger targeting of current output in the dynamic model? The answer is no: one also needs the financial market to understand the dynamic environment. To see this, let us examine how the financial market's beliefs affect the central bank's loss function across the two models. As stated in Equations (15) and (16), the expected loss function of the CB consists of the variance of the FM's beliefs about tomorrow's output and the covariance of the FM's beliefs of tomorrow's output with today's output:

$$\mathcal{V}^{static} = Var(\theta_{2|s}) - 2bCov(\theta_{2|s}, \theta_1) + b^2, \quad (24)$$

$$\mathcal{V}^{dynamic} = \frac{1}{1-\beta} \left(Var(\theta_{t+1|t}) - 2bCov(\theta_{t+1|t}, \theta_t) + b^2 \right). \quad (25)$$

For the static and dynamic loss functions to coincide, one thus needs that the covariances of beliefs are identical across the two models. One can compute the covari-

ances from manipulating the expressions for the evolution of beliefs. Those, in turn, are computed by substituting in the signal in the optimal forecast in the static model from Equation (14), and writing out the Kalman filter equation in the dynamic model, Equation (7), with the help of Appendix B:

$$\text{Static: } \theta_{2|s} = \phi\theta_1 + \frac{\phi}{\psi}\theta_2 + \phi v, \quad (26)$$

$$\text{Dynamic: } \theta_{t+1|t} = m_1\theta_{t|t-1} + \kappa\theta_t + \frac{\kappa}{\psi}\theta_{t+1} + \kappa v_t, \quad (27)$$

where

$$\phi = \frac{\rho + \frac{1}{\psi}}{1 + \frac{1}{\psi^2} + 2\frac{\rho}{\psi} + \sigma_v^2}, \quad \kappa = \frac{\rho p_4 + \frac{1}{\psi} p_1}{p_4 + \frac{1}{\psi^2} p_1 + 2\frac{\rho}{\psi} p_4 + \sigma_v^2}, \quad (28)$$

and $p_1 = \pi(\theta_{t+1})$ and $p_4 = \pi(\theta_t)$ are the two diagonal elements of the 2×2 forecast-error-variance matrix P , corresponding to the prior variances of tomorrow's output and today's, respectively.

One immediately notices that there is more information in the dynamic model. This is captured by two sources: the presence of the prior mean $\theta_{t|t-1}$ in Equation (27) and the presence of the prior variances p_1 and p_4 in κ , which render κ distinct from ϕ . Both allow the FM to carry over information from the past to the present.

One can shut off the flow of information from the past by constraining the FM to hold prior beliefs with mean 0 and variance 1, equal to the priors in the static setting. In this case, the response coefficients to incoming information become equal ($\kappa = \phi$), and the prior mean drops out of Equation (27), so that beliefs in the dynamic model evolve *exactly* as in the static model. Now, in each period of the dynamic model, the FM no longer uses any past information and thus has exactly as much information as in the static setting.

The crucial point is that even in this case of equal information, the CB's loss functions do not coincide as long as $m_1 \neq 0$. The reason is that the covariances that enter the loss functions are a function of m_1 (see Appendix B). What this is saying is that as long as the FM understands the dynamic nature of the problem, even if it were to mistakenly hold the wrong prior beliefs in every period, and therefore make erroneous point forecasts, the unconditional variance of the FM's beliefs would still reflect the knowledge that output is correlated over time. In order for dynamics not to matter at all, one thus needs to set $m_1 = 0$ as well, implying that the FM thinks that $\rho = 0$. In other words, one also needs to assume that the FM thinks that output is uncorrelated over time.

5 OPTIMAL PRECISION

Now I turn to the question of how precise the central bank's signal should be for a given targetedness. Noting that the precision of the signal is the inverse of the variance of the noise, I will use the terms "precision" and "noise" interchangeably when talking about σ_v . Thus, holding ψ fixed, the central bank now chooses σ_v to minimize the loss function in Equation (16). The next proposition gives an analytical characterization of the solution in the static and dynamic models.

PROPOSITION 4. *Corner versus interior solutions*

In the static model, optimal precision is a corner solution: σ_v^ is either 0 or ∞ . In the dynamic*

model, the roots of the first order condition are

$$\begin{pmatrix} 0 \\ \pm \frac{\sqrt{1-\rho^2}}{\sqrt{\rho}\sqrt{\psi}} \\ \pm \sqrt{\frac{\varphi_1+\varphi_2}{\varphi_3}} \\ \pm \sqrt{\frac{\varphi_1-\varphi_2}{\varphi_3}} \end{pmatrix}, \quad (29)$$

where φ_i , $i = 1, 2, 3$ are provided in Appendix E, and both corner and interior solutions exist.

This result suggests a qualitative difference between optimal precision in the static and dynamic models: whereas in the static model, the solution for σ_v is bang-bang, this is not necessarily the case in the dynamic model. Figure 6 provides a visual illustration of this point by showing σ_v^* as a function of ψ and ρ in the top panel, while the bottom panel depicts the cross-section with ψ set to 0.2, 1 and 1.8. The ψ -values are selected to encompass a case where the CB's signal is targeted toward tomorrow's output ($\psi < 1$), when the signal is not targeted ($\psi = 1$), and lastly when it is targeted toward today's output ($\psi > 1$). Also, I continue to use the calibration with $b = 1$ and $\beta = 0.99$ throughout the section.

5.1 Corner solutions in the static model

Focusing first on the solid line on the bottom panel of Figure 6, which corresponds to σ_v^* in the static model, one observes the bang-bang nature of optimal static precision. When $\psi = 0.2 < 1$, so that the CB targets future output in its signal (bottom left column of Figure 6), the CB sets $\sigma_v^* = 0$ only if $\rho \geq \rho^b(\psi)$, that is, if ρ is higher than a particular threshold which I will refer to as the bang-bang threshold, and which is a function of ψ . Note that if $\psi < 1$ and $\sigma_v = 0$, the CB's signal in Equation (5)

is an invertible moving average process, so in this case sending an infinitely precise signal ($\sigma_v = 0$) would mean that the CB reveals the output process fully. Therefore, in this case, when preference alignment is high enough ($\rho \geq \rho^b(\psi)$), the CB in the static model prefers to fully reveal the current and future output, and otherwise, it sends an infinitely imprecise signal.

A similar situation arises in the static model when $\psi > 1$, so that the signal is targeting current output. Again, there exists a bang-bang threshold $\rho^b(\psi)$ such that for $\rho > \rho^b(\psi)$, the CB's signal becomes infinitely precise, and is infinitely noisy otherwise. Since in this case the signal weights the CB's preferred state (current output) more strongly, the bang-bang threshold is lower than it is when the signal is targeted toward future output. This reflects that the CB reverts to communicating infinitely precisely already for a lower level of alignment if it gets to skew its signal towards the present.

The only ψ for which the CB chooses $\sigma_v^* = 0$ for any ρ is $\psi = 1$. The reason is that in this case, the signal is not targeted at all. Therefore the information provided by the CB improves the FM's expectations without pushing them in the wrong direction, even if $\rho \rightarrow -1$.

5.2 *Tightness of prior beliefs in the dynamic model*

Let us now examine the optimal precision σ_v^* in the dynamic model, shown by the dotted line in the bottom panel of Figure 6. The main qualitative difference is that, in contrast to the static model, the solution is not generally bang-bang. While there exist (ρ, ψ) pairs for which σ_v^* is corner (0 or ∞), for many (ρ, ψ) pairs the optimal precision is interior.

When $\psi = 0.2$ (bottom left panel of Figure 6), so that the CB's signal is targeted

toward future output, the FM's preferred state, σ_v^* shifts from ∞ to an interior value when ρ exceeds $\rho^i(\psi)$, a threshold value I will call interiority threshold. Optimal noise only shrinks to zero when $\rho = 1$. When $\psi = 1$ (bottom middle panel of Figure 6), so that the weight on the two states in the CB's signal is equal, the same pattern obtains. For $\rho < \rho^i(\psi)$, optimal noise is infinity, and as $\rho > \rho^i(\psi)$, σ_v^* drops below infinity, converging to zero for $\rho \rightarrow 1$. The only difference between the future-targeted and the untargeted signal is that for the untargeted signal, the interiority threshold arrives earlier. This resembles the intuition in the static model that for a signal that is more targeted towards today's output (has a higher value of ψ), the CB pivots away from infinite imprecision already for a lower level of ρ .

Something entirely different happens when the CB's signal is strongly targeted towards current output, the CB's preferred state. This situation, depicted on the bottom right panel of Figure 6, involves an initial bang-bang as σ_v^* transitions from ∞ to 0 when ρ hits a particular bang-bang threshold $\rho^b(\psi) < 1$. Afterwards, σ_v^* remains 0 as ρ increases, until ρ exceeds an interiority threshold. Here, σ_v^* becomes interior *from below*, surprisingly increasing as alignment between the two states grows. Finally, once alignment becomes very strong, σ_v^* reverses course and decreases back down to asymptote to 0 as $\rho \rightarrow 1$.

What is going on? Why is optimal precision often interior in the dynamic model, while it never is in the static model? And if σ_v^* passes from ∞ to 0 through an interiority region when $\psi \leq 1$, why does this interiority region disappear when $\psi > 1$, only to reappear for much higher ρ -values? To understand why all of this happens, it is helpful to investigate the tightness of prior and posterior beliefs in the two models. Recall the notation of $\pi(\theta_T)$ for prior and $p(\theta_T, s_t)$ for posterior variances of θ_T , intro-

duced in Equation (22), and also recall the measure of the informativeness of a signal s_t from Equation (23)

$$I(\theta_T, s_t) := \pi(\theta_T) - p(\theta_T, s_t), \quad (30)$$

which captures the decrease in uncertainty in θ_T given s_t . As demonstrated in Appendix D, the FM's prior on current or future output in the static model is equal to the unconditional variance of output, 1. In the dynamic model, instead, in any period, the posterior variance $p(\theta_{t+1}, s_t)$ is carried over to the next period, becoming the endogenous prior $\pi(\theta_t)$ in the following period. On the one hand, this has the important consequence that in the static model

$$I(\theta_T, s) = \text{Var}(\theta_{T|s}), \quad T = 1, 2, \quad (31)$$

implying that informativeness is a monotonically decreasing function in σ_v . It is easy to verify that this is true for informativeness about θ_{t+1} in the dynamic model as well, as

$$I(\theta_{t+1}, s_t) = (1 - \rho^2) \text{Var}(\theta_{t+1|t}). \quad (32)$$

On the other hand, it clearly does not always hold for θ_t :

$$I(\theta_t, s_t) = \text{Var}(\theta_{t|t}) - \text{Var}(\theta_{t|t-1}). \quad (33)$$

In other words, the dynamic model has the property that informativeness about θ_t is nonmonotonic in σ_v . The following proposition and corollary summarize this insight, together with its implications for the optimal precision choice.

PROPOSITION 5. *Tight priors dampen informativeness in the dynamic model*

Informativeness about θ_t is nonmonotonic in σ_v because the endogenous prior carries over information from period $t - 1$ to period t . Let s^h denote a high-precision signal, while s^l a low-precision one. If a precise signal in the previous period, s_{t-1}^h , tightens the prior in t sufficiently,

then the reduction in uncertainty at time t due to the signal pair $\{s_t, s_{t-1}^h\}$ is lower than for the pair $\{s_t, s_{t-1}^l\}$.

COROLLARY 5.1. *Information smoothing in the dynamic model*

In the dynamic model, a lower degree of precision serves to keep priors from getting too tight; the optimal precision choice involves the smoothing of information provision.

The nonmonotonicity of $I(\theta_t, s_t)$ comes from the signal structure in Equation (5), which implies that in any given period t , the current signal s_t is not the first to provide the FM with information about θ_t (see also the discussion in Section 4.2). If all signals have sufficiently high precision, then by the time period t arrives, the FM has already learned enough about θ_t from s_{t-1} , and will therefore find the current signal less informative than if it had acquired less information in the past.

Figure 7 illustrates Proposition 5, plotting the reduction in uncertainty in θ_{t+1} (dotted line) against that in θ_t (solid line), both as a function of σ_v for various (ρ, ψ) pairs. The dash-dotted line corresponds to the optimal precision σ_v^* . Just like Equation (32) suggests, $I(\theta_{t+1}, s_t)$ indeed decreases monotonically in σ_v . $I(\theta_t, s_t)$, instead, is hump-shaped in the majority of cases, assuming a maximum at an “informativeness-maximizing” precision level I will denote by σ_v^I :

$$\sigma_v^I := \arg \max I(\theta_t, s_t). \quad (34)$$

The optimal precision, σ_v^* , shown in the dash-dotted line on Figure 7, lends support to Corollary 5.1. The interior solutions for σ_v are visible as situations where the dash-dotted line is not stuck at the numerical bounds I impose on the figure. Whenever the optimal precision is interior, σ_v^* is in the neighborhood of σ_v^I . In other words, the CB finds it optimal not to shrink the signal noise to zero because this allows it

to push $I(\theta_t, s_t)$ up, that is, increase the informativeness of the signal about current output. This lets the CB ensure that in any period t , the current signal contributes to the FM's information set and thus exerts persuasion. Thus the CB wants the FM to receive approximately equal amounts of information in each period, a phenomenon I refer to as the smoothing of information provision.

Thus it is the nonmonotonicity of informativeness about θ_t that is behind the interior solutions in σ_v in the dynamic model. Intuitively, this nonmonotonicity captures the effect of the tightness of the FM's priors on the CB's communication problem. Clearly, there are parameter configurations for which it is optimal to communicate infinitely precisely. But for many (ρ, ψ) pairs, precise communication in period $t - 1$ would tighten the FM's prior in period t too much, making it hard for the CB to influence the FM's beliefs. Therefore the key driver of the difference between static and dynamic optimal precision is the varying tightness of the prior.

But it is not quite the only driver. If it were, then the CB would find it optimal to set precision exactly so as to always maximize informativeness about current output ($\sigma_v^* = \sigma_v^I$). This is however clearly not the case, as one can see by comparing the dash-dotted line (σ_v^*) with the informativeness-maximizing precision level (σ_v^I) on Figure 7. More often than not, the two are not equal. So to find out what other forces are behind the optimal precision choice, let us now investigate how σ_v^* depends on first ρ and then ψ .

Fixing ψ and increasing ρ means that for a fixed column, one moves down through the rows of Figure 7. When $\rho \ll 0$, $\sigma_v^* = \infty$. This tells us that when preference misalignment is strong, the CB wishes to minimize both $I(\theta_{t+1}, s_t)$ and $I(\theta_t, s_t)$, because in this case, any information in the signal will be misused by the FM in the sense

that its investment choice will be of opposite sign to θ_t . When alignment rises above the interiority threshold, $\rho^i(\psi)$, for example when ρ rises from -0.9 to 0 in the middle column of Figure 7, the CB seeks to maximize $I(\theta_t, s_t)$, i.e. to provide as much information about current output as possible. This requires both $\sigma_v < \infty$ and $\sigma_v > 0$ because of the nonmonotonicity of $I(\theta_t, s_t)$.

At the same time, depending on the extent of alignment, the CB sometimes sees it necessary to also minimize the information it reveals about future output. This is the case in the $(\rho, \psi) = (0, 1)$ panel of the figure, pushing σ_v^* above σ_v^I . As alignment becomes near perfect, for instance in the last row of the figure, where $\rho = 0.9$, the CB no longer simply maximizes $I(\theta_t, s_t)$, because providing information about tomorrow's output pushes the FM in the same direction as information about today's. Now instead it is optimal to maximize a weighted sum of information provision about both states. This pushes σ_v^* below σ_v^I toward zero.

Let us now hold alignment fixed and investigate how targetedness affects the optimal precision choice. For a fixed ρ , increase ψ by moving right along any row of Figure 7. Except for the first row, where $\rho \ll 0$ and thus $\sigma_v^* = \infty \forall \psi$, a higher ψ leads to a strictly lower σ_v^* . For example, in the middle row of the figure, where $\rho = 0$, optimal noise is infinity when $\psi = 0.2$. As ψ increases to 1, σ_v^* takes on an interior value around 1.4. Finally, when ψ rises to 1.8, σ_v^* drops to zero. Intuitively, what is happening here is that the CB renders those signals precise that target its favored state strongly.

There is something special about $\psi > 1$, however, i.e. when the CB's signal is targeted towards current output. As the bottom right panel of Figure 6 recalls, not only does this involve the only bang-bang σ_v^* -behavior in the dynamic model, so that

σ_v^* switches from ∞ to 0 once ρ exceeds the bang-bang threshold $\rho^b(\psi)$, but σ_v^* rises above 0 again for ρ high enough. We can read off why this happens from the right column of Figure 7. The cases of $\rho = -0.9$ or 0 are the only instances on the figure where $I(\theta_t, s_t) > I(\theta_{t+1}, s_t)$. This is because for a combination of sufficiently low alignment and sufficiently high targetedness towards the present, the CB's signal is very informative about current output, but this informativeness does not carry over to future output. In this case, the CB can max out informativeness about today's output by setting $\sigma_v^* = 0$.

If alignment is too low, however, as in the top right panel of Figure 7, where $\rho = -0.9$, the informativeness of the signal about current output is misused by the FM to deduce clues about the future, and thus the CB can do no better than to shut off precision and set $\sigma_v^* = \infty$. If alignment is instead too high, then it is no longer the case that the signal is more informative about the present than the future, and therefore adding a little noise is optimal to push $I(\theta_t, s_t)$ up and $I(\theta_{t+1}, s_t)$ down a bit.

The optimal precision choice of the CB is thus a much more complicated object in the dynamic model than in the static one. It is mainly driven by the tightness of the financial market's prior, which, in the dynamic world, is endogenous to central bank communication. A too tight prior is a constraint to the CB as it renders current communication ineffective. Therefore adding noise to the signal serves the function of loosening up the prior and distributing the information provided to the FM evenly across time.⁶

6 CONCLUSION

Central bank communication is a dynamic problem: in a world where economic fundamentals evolve over time and the monetary authority communicates repeatedly, central bankers have to provide information about both the present and the future. This introduces a novel tradeoff for central bankers, as their statements reach only the subset of the public that is concerned with the future business cycle. At the same time, the central bank has a mandate to maximize welfare of the entire public, rendering preferences misaligned between sender and receiver of central bank communication.

This paper is the first to investigate the implications of dynamics for central bank communication. Analyzing two models identical up to the interpretation of the correlation between two states allows me to isolate the effect of temporal dependence between the states on the communication problem. It turns out that the fact that the financial market's prior beliefs become endogenous in the dynamic model drives a wedge between the static and dynamic communication policies, dampening the persuasiveness of the central bank's signal.

In terms of optimal targetedness, the dynamic model involves a higher weight on today's state in the signal than the static model for any correlation. This is because, as opposed to cross-sectional correlation, temporal dependence renders the two states sufficiently distinct that the central bank needs to exert additional persuasion to push the financial market's beliefs in the right direction.

Optimal precision is always weakly lower in the dynamic model than in the static one. This is driven by the central bank's desire to loosen the financial market's priors, which, if too tight, render the central bank's signal less persuasive. Adding noise to the signal flattens the information profile of the financial market in the sense that it

does not receive too much information too early. The optimal precision choice thus involves the smoothing of information provision.

A A STYLIZED MODEL OF INVESTMENT

Consider a firm problem in partial equilibrium (gross interest rate R is fixed). The firm chooses the inputs to production: hours worked, h_t , and capital, k_{t+1} . Since the firm inherits yesterday's capital level today, it chooses current investment, i_t , to select tomorrow's capital level. The problem can be stated as

$$\max_{i_t, k_{t+1}, h_t} V = \mathbb{E}_0 \sum_{t=0}^{\infty} \left(\frac{1}{R} \right)^t \left(f(k_t, h_t) - w_t h_t - i_t - \Phi \left(\frac{i_t}{k_t} \right) k_t \right), \quad (\text{A.1})$$

$$\text{s.t. } k_{t+1} = (1 - \delta)k_t + i_t, \quad (\text{A.2})$$

where $f(k_t, h_t)$ is the production function, and $\Phi \left(\frac{i_t}{k_t} \right)$ are convex capital adjustment costs. Setting up a Lagrangian with a multiplier q_t on the law of motion of capital, taking the first order condition for capital, and iterating it forward gives

$$q_t = \frac{1}{1 - \delta} \mathbb{E}_t \sum_{s=1}^{\infty} \left(\frac{1 - \delta}{R} \right)^s \left(f_k(k_{t+s}, h_{t+s}) + \Phi' \left(\frac{i_{t+s}}{k_{t+s}} \right) \frac{i_{t+s}}{k_{t+s}} - \Phi \left(\frac{i_{t+s}}{k_{t+s}} \right) \right). \quad (\text{A.3})$$

This equation is the key Q-theory relation of optimal investment choice. What it is saying is that "little-q," the marginal value of capital, is the determinant of investment. This, in turn, is determined by the discounted future stream of profits. Thus, when selecting today's investment, firms first and foremost need to know what future economic conditions will be like, so that they can accurately forecast future profits.

Contrast this forward-looking optimization choice with the demand side of the GDP accounting identity

$$Y = C + I, \quad (\text{A.4})$$

where I assume a closed economy and abstract from government expenditures. This relation instead implies that high current investment raises demand for goods and services, leading firms to produce more today and thus hire more workers, raising

current employment. For those who reside on Main Street, it is thus clearly the present production and thus employment conditions which are relevant. The central bank's mandate of stabilizing current inflation and employment thus pitches it against Wall Street, who seeks to set investment with the future in mind.

B THE KALMAN FILTER FOR THE DYNAMIC MODEL

Denoting $\theta'_t := \theta_{t+1}$, set up the state-space system as:

$$x_{t+1} = hx_t + \eta\epsilon_{t+1} \quad \text{state equation,} \quad (\text{B.1})$$

$$y_t = gx_t + v_t \quad \text{observation equation,} \quad (\text{B.2})$$

with the observation vector given by $y_t = s_t$, and the state vector x_t and transition matrices are

$$x_t = \begin{bmatrix} \theta'_t \\ \theta_t \end{bmatrix}, \quad h = \begin{bmatrix} \rho & 0 \\ 1 & 0 \end{bmatrix}, \quad g = \begin{bmatrix} \frac{1}{\psi} & 1 \end{bmatrix}, \quad \eta = \begin{bmatrix} \sigma_\epsilon & 0 \\ 0 & 0 \end{bmatrix}, \quad \epsilon_t = \begin{bmatrix} \epsilon_t \\ 0 \end{bmatrix}, \quad Q = \eta\eta' = \begin{bmatrix} \sigma_\epsilon^2 & 0 \\ 0 & 0 \end{bmatrix},$$

$$R = \sigma_v^2.$$

The optimal forecast of the state vector at $t + 1$ is:

$$x_{t+1|t} = hx_{t|t-1} + h\kappa(y_t - gx_{t|t-1}), \quad (\text{B.3})$$

where $\kappa = Pg'\Omega^{-1}$ and the steady-state forecast-error-variance matrix P solves

$$hPh' - hPg'\Omega^{-1}gPh' + Q = P, \quad (\text{B.4})$$

$$\Omega = gPg' + R. \quad (\text{B.5})$$

Given the analytical solution to the 2×2 matrix $P = \begin{bmatrix} p_1 & p_2 \\ p_2 & p_4 \end{bmatrix}$, suppressed here, one obtains the law of motion of beliefs

$$\theta_{t+1|t} = m_1\theta_{t|t-1} + m_2\theta_t + m_3\theta_{t+1} + m_4v_t, \quad (\text{B.6})$$

$$\theta_{t|t} = n_1\theta_{t|t-1} + n_2\theta_t + n_3\theta_{t+1} + n_4v_t. \quad (\text{B.7})$$

The m_i, n_i $i = 1, \dots, 4$ coefficients are

$$m_1 = \rho - \kappa_1 \left(\frac{\rho}{\psi} + 1 \right), \quad (\text{B.8})$$

$$m_2 = m_4 = \kappa_1, \quad (\text{B.9})$$

$$m_3 = \frac{\kappa_1}{\psi}, \quad (\text{B.10})$$

$$n_1 = 1 - \kappa_2 \left(\frac{\rho}{\psi} + 1 \right), \quad (\text{B.11})$$

$$n_2 = n_4 = \kappa_2, \quad (\text{B.12})$$

$$n_3 = \frac{\kappa_2}{\psi}, \quad (\text{B.13})$$

where κ_1 is the first element of the 2×1 Kalman gain and is given by

$$\kappa_1 = \frac{\psi(p_1 + \rho p_4 \psi)}{p_1 + \psi(2\rho p_4 + \psi(p_4 + \sigma_v^2))}, \quad (\text{B.14})$$

and κ_2 is the second element, given by

$$\kappa_2 = \frac{\psi(\rho p_4 + p_4 \psi)}{p_1 + \psi(2\rho p_4 + \psi(p_4 + \sigma_v^2))}. \quad (\text{B.15})$$

For simplicity, I refer to κ_1 as κ in the main text.

C EXPECTED LOSS FUNCTIONS

The expected loss function in the static model is given by

$$\gamma^{static} = \frac{b^2 (\psi (2\rho + \psi \sigma_v^2 + \psi) + 1) - 2b(\rho + \psi)(\rho\psi + 1) + (\rho\psi + 1)^2}{\psi (2\rho + \psi \sigma_v^2 + \psi) + 1}, \quad (\text{C.1})$$

with an associated first order condition with respect to ψ of

$$\frac{2(b(\rho^2 + \psi^2(\rho^2(\sigma_v^2 - 1) + \sigma_v^2 + 1) + 2\rho\psi\sigma_v^2 - 1) + \psi(\rho\psi + 1)(\rho^2 - \sigma_v^2 - 1))}{(\psi(2\rho + \psi\sigma_v^2 + \psi) + 1)^2} = 0. \quad (\text{C.2})$$

As Equation (16) in the main text recalls, the expected loss function in the dynamic model is

$$\mathcal{V}^{dynamic} = \frac{1}{1 - \beta} \left(\text{Var}(\theta_{t+1|t}) - 2b\text{Cov}(\theta_{t+1|t}, \theta_t) + b^2 \right).$$

The expressions for the covariances in the above are

$$\text{Var}(\theta_{t+1|t}) = \frac{\rho m_2 + m_3}{1 - \rho m_1}, \quad (\text{C.3})$$

$$\text{Cov}(\theta_{t+1|t}, \theta_t) = m_1 \text{Var}(\theta_{t+1|t}) + m_2 + \rho m_3, \quad (\text{C.4})$$

where the m_i , $i = 1, 2, 3$, coefficients are given in Appendix B, and the first order condition is suppressed.

D PRIOR AND POSTERIOR VARIANCES IN THE TWO MODELS

Using the definitions of prior and posterior variances from the main text, in the static model these are given by

$$\pi(\theta_2) = 1, \quad (\text{D.1})$$

$$p(\theta_2) = 1 - \text{Var}(\theta_{2|s}), \quad (\text{D.2})$$

$$\pi(\theta_1) = 1, \quad (\text{D.3})$$

$$p(\theta_1) = 1 - \text{Var}(\theta_{1|s}). \quad (\text{D.4})$$

In the dynamic model, instead, they are

$$\pi(\theta_{t+1}) = 1 - \rho^2 \text{Var}(\theta_{t+1|t}), \quad (\text{D.5})$$

$$p(\theta_{t+1}) = 1 - \text{Var}(\theta_{t+1|t}), \quad (\text{D.6})$$

$$\pi(\theta_t) = 1 - \text{Var}(\theta_{t|t-1}) = p(\theta_{t+1}), \quad (\text{D.7})$$

$$p(\theta_t) = 1 - \text{Var}(\theta_{t|t}), \quad (\text{D.8})$$

where the equivalence between (D.6) and (D.7) uses covariance stationarity of steady-state beliefs.

E INFORMATIVENESS IN THE TWO MODELS

Using the expressions for prior and posterior variances from Appendix D, the notions of informativeness of today's signal in the static model are given by

$$I(\theta_2, s) = \text{Var}(\theta_{2|s}) = \frac{\rho^2 + \frac{2\rho}{\psi} + \frac{1}{\psi^2}}{\frac{2\rho}{\psi} + \sigma_v^2 + \frac{1}{\psi^2} + 1}, \quad (\text{E.1})$$

$$I(\theta_1, s) = \text{Var}(\theta_{1|s}) = \frac{\frac{\rho^2}{\psi^2} + \frac{2\rho}{\psi} + 1}{\frac{2\rho}{\psi} + \sigma_v^2 + \frac{1}{\psi^2} + 1}. \quad (\text{E.2})$$

In the dynamic model, instead, they are

$$I(\theta_{t+1}, s_t) = (1 - \rho^2) \text{Var}(\theta_{t+1|t}), \quad (\text{E.3})$$

$$I(\theta_t, s_t) = \text{Var}(\theta_{t|t}) - \text{Var}(\theta_{t+1|t}). \quad (\text{E.4})$$

The relevant variances are

$$\text{Var}(\theta_{t+1|t}) = \frac{\rho m_2 + m_3}{1 - \rho m_1}, \quad (\text{E.5})$$

$$\text{Var}(\theta_{t|t}) = n_1 \text{Var}(\theta_{t+1|t}) + n_2 + \rho n_3, \quad (\text{E.6})$$

where the m_i, n_i , $i = 1, 2, 3$, coefficients are given in Appendix B.

F OPTIMAL PRECISION

The first order condition in the static model is

$$-\frac{2\psi^2(\rho\psi + 1)\sigma_v(-2b(\rho + \psi) + \rho\psi + 1)}{(\psi(2\rho + \psi\sigma_v^2 + \psi) + 1)^2} = 0, \quad (\text{F.1})$$

which is solved by $\sigma_v = 0$. That this is either a minimum or a maximum can be verified by considering the second order condition, which is given by

$$\frac{2\psi^2(\rho\psi + 1)(2b(\rho + \psi) - \rho\psi - 1)(\psi(2\rho - 3\psi\sigma_v^2 + \psi) + 1)}{(\psi(2\rho + \psi\sigma_v^2 + \psi) + 1)^3}. \quad (\text{F.2})$$

The first and second order conditions in the dynamic model are too large and are thus suppressed. The expressions for φ_i , $i = 1, 2, 3$ from the main text are

$$\varphi_1 := (\rho - 1)(\rho + 1)\psi^2(\rho - 2b)(2b(\rho^2(\psi^2 + 1) + 3\rho\psi + \psi^2) - \psi(\rho\psi + 1)), \quad (\text{F.3})$$

$$\varphi_2 := \sqrt{2}\sqrt{b(\rho^2 - 1)^2\psi^4(2b - \rho)(\rho + \psi)^2(\rho\psi + 1)(-2b(\rho + \psi) + \rho\psi + 1)^2}, \quad (\text{F.4})$$

$$\varphi_3 := \rho\psi^4(\rho - 2b)(2b(\rho^3 - 2\rho - \psi) + \rho\psi + 1). \quad (\text{F.5})$$

REFERENCES

- Angeletos, George-Marios, and Alessandro Pavan. (2007) "Efficient Use of Information and Social Value of Information." *Econometrica*. 75 (4): 1103–1142.
- Barro, Robert J, and David B Gordon. (1983) "Rules, discretion and reputation in a model of monetary policy." *Journal of Monetary Economics*. 12 (1): 101–121.
- Binder, Carola. (2017) "Fed speak on main street: Central bank communication and household expectations." *Journal of Macroeconomics*. 52 238–251.
- Binder, Carola, and Gwangmin Kim. (Forthcoming) "Learning-through-Survey in Inflation Expectations." *American Economic Journal: Macroeconomics*.
- Campbell, Jeffrey R, Charles L Evans, Jonas DM Fisher, Alejandro Justiniano, Charles W Calomiris, and Michael Woodford. (2012) "Macroeconomic effects of Federal Reserve forward guidance [with comments and discussion]." *Brookings Papers on Economic Activity*. 1–80.
- Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko. (2021) "The Inflation Expectations of US Firms: Evidence from a New Survey." Manuscript.
- Chahrour, Ryan. (2014) "Public Communication and Information Acquisition." *American Economic Journal: Macroeconomics*. 6 (3): 73–101.
- Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar. (2018) "How do firms form their expectations? New survey evidence." *American Economic Review*. 108 (9): 2671–2713.

- Coibion, Olivier, Yuriy Gorodnichenko, Saten Kumar, and Mathieu Pedemonte. (2020) "Inflation expectations as a policy tool?" *Journal of International Economics*. 124, 103297.
- Frankel, Alex, and Navin Kartik. (2018) "What Kind of Central Bank Competence?" *Theoretical Economics*. 13 (2): 697–727.
- Gaballo, Gaetano. (2016) "Rational Inattention to News: The Perils of Forward Guidance." *American Economic Journal: Macroeconomics*. 8 (1): 42–97.
- Goldstein, Itay, and Yaron Leitner. (2018) "Stress tests and information disclosure." *Journal of Economic Theory*.
- Hansen, Stephen, and Michael McMahon. (2016) "First impressions matter: Signalling as a source of policy dynamics." *Review of Economic Studies*. 83 (4): 1645–1672.
- Hellwig, Christian, and Laura Veldkamp. (2009) "Knowing What Others Know: Coordination Motives in Information Acquisition." *The Review of Economic Studies*. 76 (1): 223–251.
- Herbert, Sylvérie. (2021) "State-Dependent Central Bank Communication with Heterogeneous Beliefs." Manuscript.
- Inostroza, Nicolas, and Alessandro Pavan. (2022) "Persuasion in global games with application to stress testing." Manuscript.
- Kamenica, Emir, and Matthew Gentzkow. (2011) "Bayesian persuasion." *The American Economic Review*. 101 (6): 2590–2615.
- Ko, Eunmi. (2019) "Optimal Central Bank Forward Guidance." Manuscript.

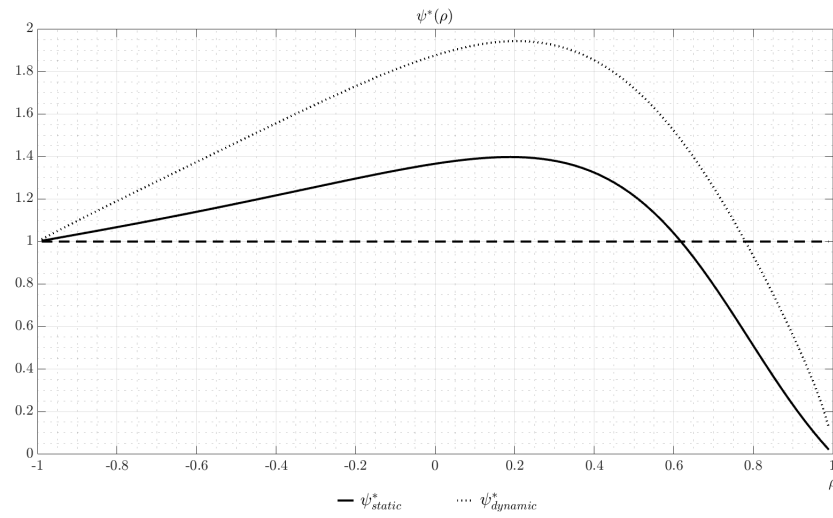
- Kumar, Saten, Hassan Afrouzi, Olivier Coibion, and Yuriy Gorodnichenko. (2015) "Inflation targeting does not anchor inflation expectations: Evidence from firms in New Zealand." *Brookings Papers on Economic Activity*. Fall 151–225.
- Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt. (2018) "Dynamic rational inattention: Analytical results." *Journal of Economic Theory*. 176 650–692.
- Morris, Stephen, and Hyun Song Shin. (2002) "Social Value of Public Information." *American Economic Review*. 92 (5): 1521–1534.
- Moscarini, Giuseppe. (2007) "Competence Implies Credibility." *The American Economic Review*. 97 (1): 37–63.
- Pfäuti, Oliver. (2021) "Inflation—who cares? Monetary Policy in Times of Low Attention." Manuscript.
- Reis, Ricardo. (2011) "When Should Policymakers Make Announcements?" 2011 Meeting Papers 122, Society for Economic Dynamics.
- Sims, Christopher A. (2003) "Implications of rational inattention." *Journal of Monetary Economics*. 50 (3): 665–690.
- Svensson, Lars E. O. (2006) "Social Value of Public Information: Comment: Morris and Shin (2002) Is Actually Pro-Transparency, Not Con." *American Economic Review*. 96 (1): 448–452.
- Tobin, James. (1969) "A general equilibrium approach to monetary theory." *Journal of Money, Credit and Banking*. 1 (1): 15–29.

Van Rooij, Maarten, Annamaria Lusardi, and Rob Alessie. (2011) “Financial literacy and stock market participation.” *Journal of Financial Economics*. 101 (2): 449–472.

NOTES

1. This is not a binding constraint as lying turns out to never be optimal in equilibrium.
2. Online Appendix [A](#) considers an alternative formalization of the targetedness dimension via a signal structure comprising two independent signals, where the choice variables of the CB are the variances of the noise terms in each signal.
3. This follows from a direct application of [Maćkowiak et al. \(2018\)](#)’s Propositions 1 and 2, noting that Main Street internalizes that what is tomorrow’s output today will become today’s output tomorrow.
4. This concept is related to the notion of mutual information, the expected difference between the log posterior and the log prior distribution. My “informativeness” concept simply compares the level difference between the second moments of the prior and posterior.
5. For analytical expressions for prior and posterior variances and informativeness, see Appendices [D](#) and [E](#).
6. I provide an algorithm for choosing the optimal precision through an indirect utility formulation in Online Appendix [B](#).

Figure 1: Optimal targetedness and FM responsiveness as a function of ρ



The figure shows the optimal targetedness, ψ^* , as a function of ρ , for σ_v normalized to 1. The solid line corresponds to the static, and the dotted line to the dynamic model, while the dashed line indicates the no-targetedness case of $\psi = 1$.

Figure 2: Optimal targetedness policy across the two models

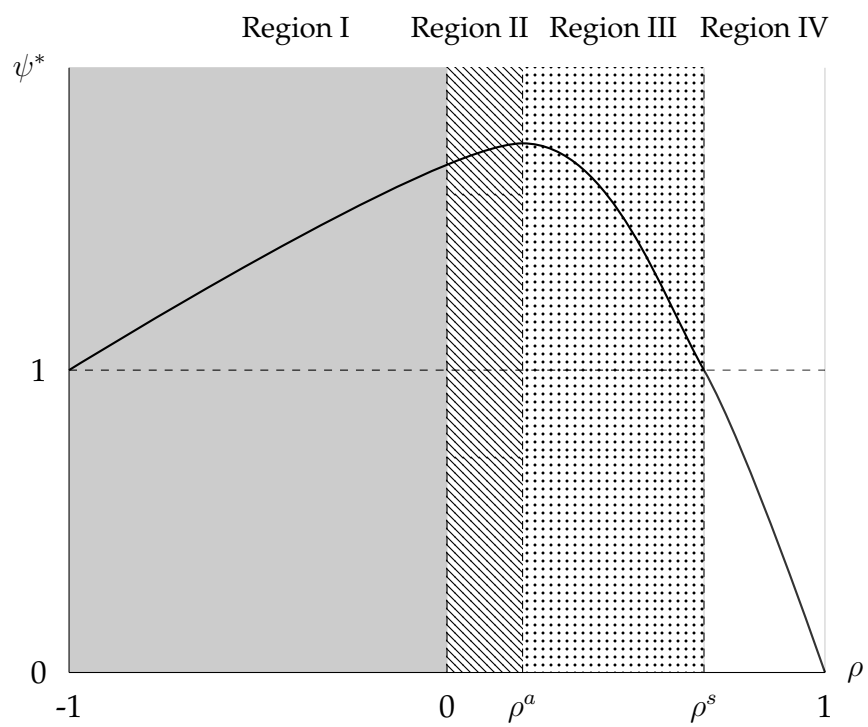
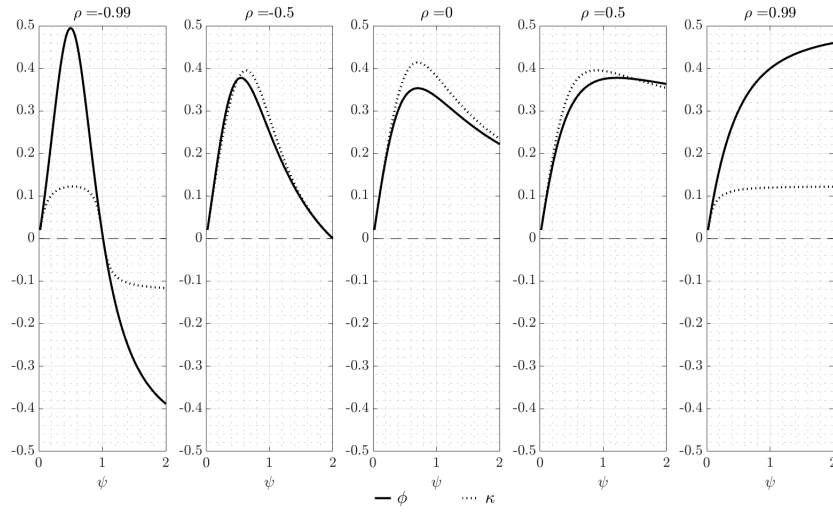
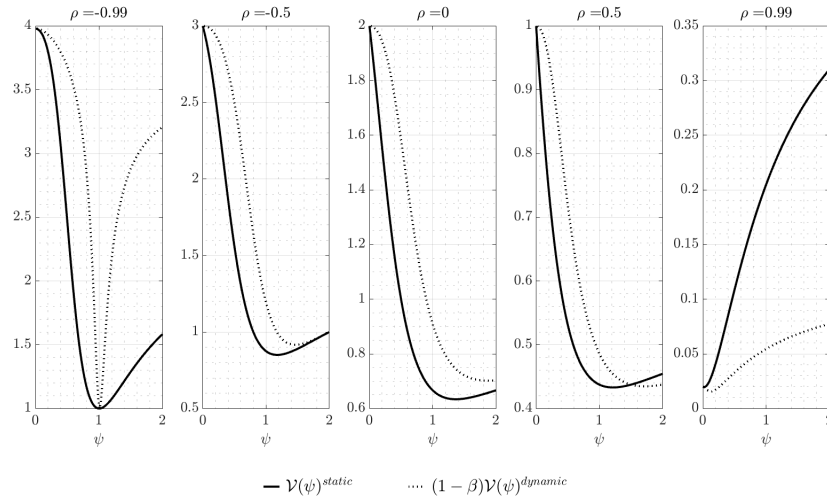


Figure 3: FM responsiveness to the signal as a function of ψ for various values of ρ



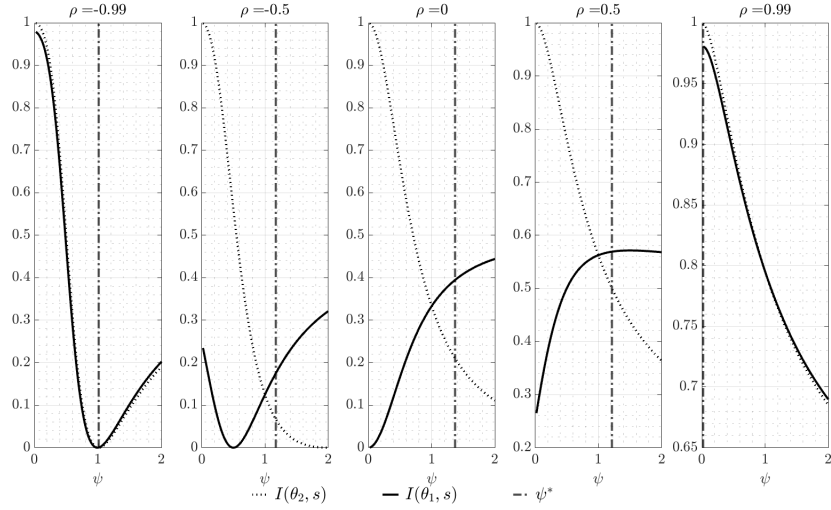
The figure shows the FM's responsiveness to the signal, ϕ in the static model (solid line) and κ in the dynamic model (dotted line), as a function of targetedness, ψ , for various levels of correlation between the states (ρ).

Figure 4: Expected losses as a function of ψ for various ρ

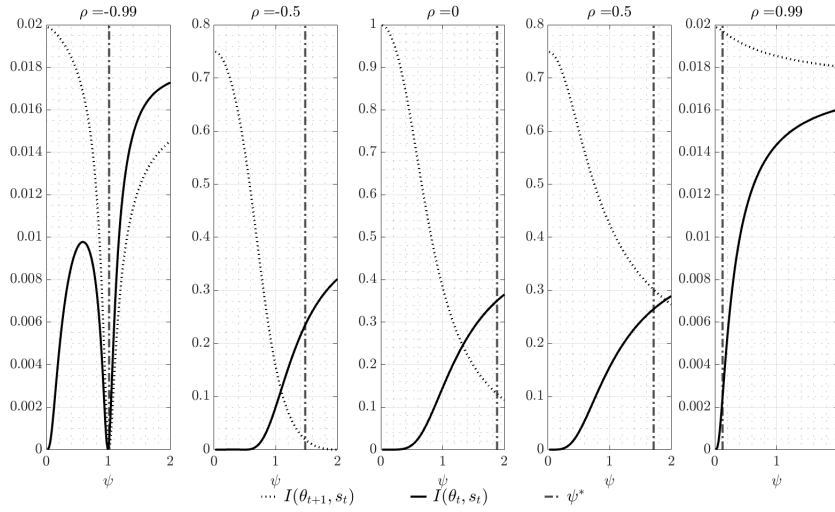


The figure plots expected losses in the static model (solid line), and in the dynamic model (dotted line), both as a function of targetedness, and for various values of ρ . Since the expected loss in the dynamic model is a discounted sum of period losses, I normalize it by $1 - \beta$ (see Equations (15) and (16)).

Figure 5: Informativeness of the current signal as a function of ψ for various ρ



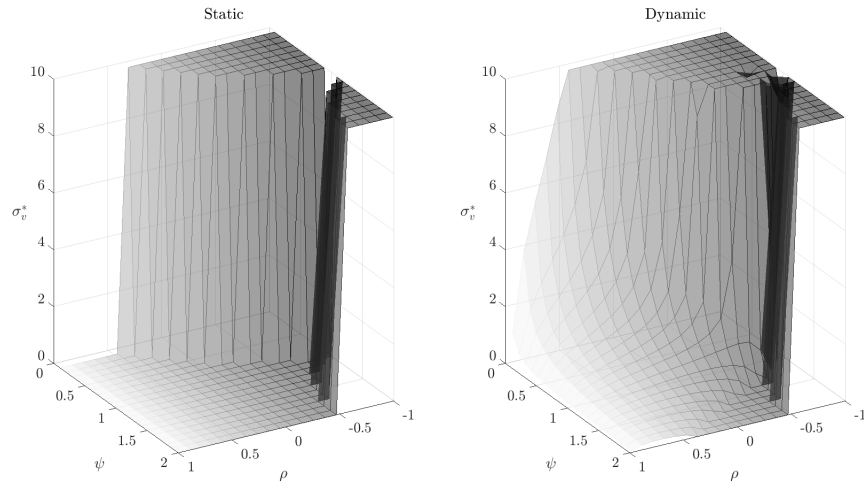
(a) Static model



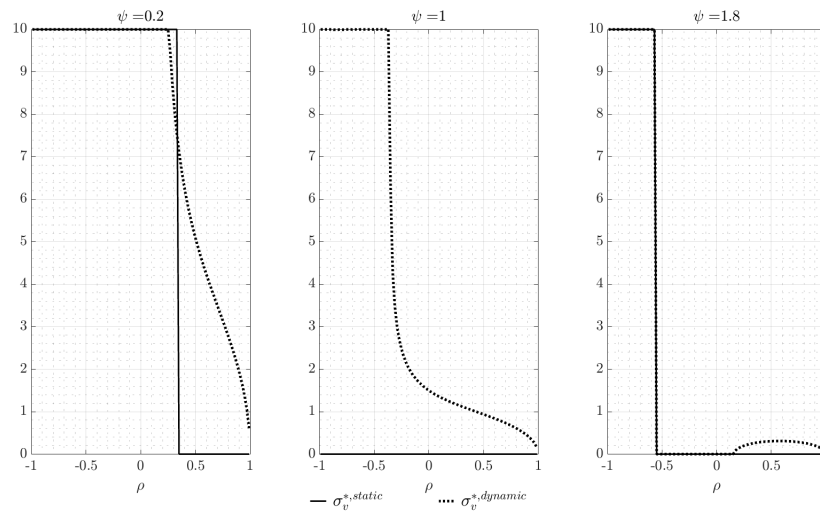
(b) Dynamic model

The figure plots the informativeness of today's signal about tomorrow's state (dotted line), and about today's state (solid line). The top panel refers to the static, and the bottom to the dynamic model. For both models, the dash-dotted line indicates the optimal targetedness ψ^* . Recall that I define informativeness as $I(\theta_T, s_t) := \pi(\theta_T) - p(\theta_T, s_t)$, so that it is the level difference between prior and posterior variances.

Figure 6: Optimal precision σ_v^*



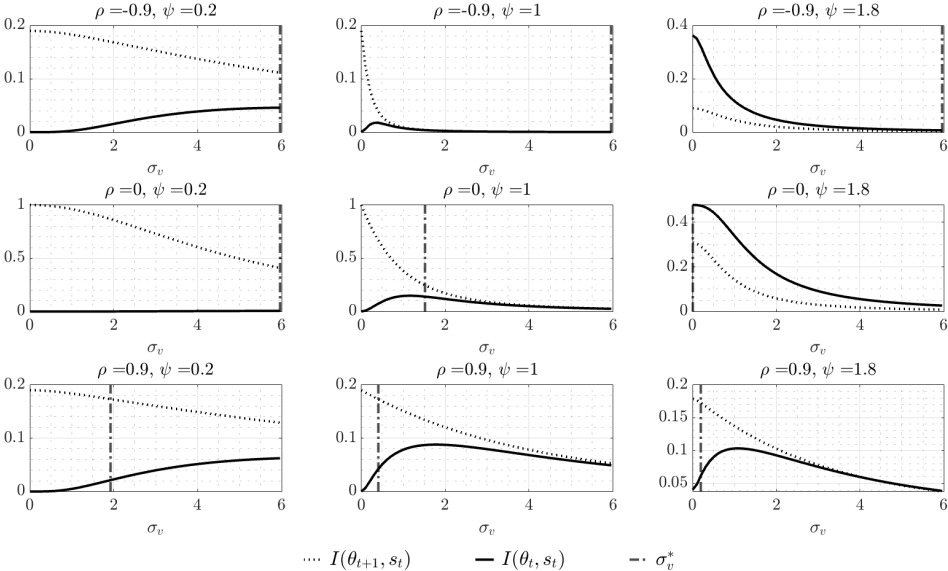
(a) 3D plot



(b) Cross-section for various values of ψ

The top panel shows optimal precision σ_v^* as a function of ψ and ρ . The bottom panel shows cross-sections for various values of ψ , with the solid line pertaining to the static, and the dotted line to the dynamic model.

Figure 7: Informativeness of today's signal as a function of σ_v for various (ρ, ψ) pairs



The figure plots the informativeness of signal s_t about θ_{t+1} (dotted line) and about θ_t (solid line). The dash-dotted line indicates the optimal precision σ_v^* . Recall that I define informativeness as $I(\theta_T, s_t) := \pi(\theta_T) - p(\theta_T, s_t)$, so that it is the level difference between prior and posterior variances.