Managing Beliefs about Monetary Policy under Discretion *

Elmar Mertens†
Federal Reserve Board

October 30, 2014

Abstract

In models of monetary policy, discretionary policymaking is typically constrained in its ability to manage public beliefs. However, when a policymaker possesses private information, policy actions serve as signals to the public about unobserved economic conditions and belief management becomes an integral part of optimal discretion policies. My paper derives the optimal time-consistent policy for a general linear-quadratic setting.

The optimal policy is illustrated in a simple New Keynesian model, where analytical solutions can be derived as well. In this model, imperfect information about the policymaker's output target leads to lower policy losses.

JEL Classification: E31, E37, E47, E52, E58

Keywords: Optimal Monetary Policy, Discretion, Time-Consistent Policy, Markov-Perfect Equilibrium, Incomplete Information, Kalman Filter

---

*I would like to thank the members of my dissertation committee, Jean-Pierre Danthine, Philippe Bacchetta, Robert G. King, and Peter Kugler as well as the editor (Kenneth West) and two anonymous referees for their advice and stimulating inputs. Furthermore I would like to thank seminar participants at the University of Basel, Boston University, the Federal Reserve Bank of Richmond, the Federal Reserve Board as well as the University of Lausanne for their comments. All remaining errors are of course mine.

†For correspondence: Elmar Mertens, Board of Governors of the Federal Reserve System, Washington D.C. 20551. email `elmar.mertens@frb.gov`. Tel.: +(202) 452 2916. The views in this paper do not necessarily represent the views of the Federal Reserve Board, or any other person in the Federal Reserve System or the Federal Open Market Committee. Any errors or omissions should be regarded as solely those of the author.

1. INTRODUCTION

Starting with Kydland and Prescott (1977) and Barro and Gordon (1983), the theoretical literature on rules versus discretion has documented clear benefits from commitment in monetary policy. Many economic decisions in the private sector are forward-looking and depend on policy expectations. In such an environment, policy rules and commitment strategies benefit from their ability to manage public beliefs about future policies. Most of the rules-versus-discretion literature is based on models of perfect information, symmetrically shared between the central bank and the public. However, when a policymaker possesses private information, belief management becomes an integral part of optimal discretion policies, too, as demonstrated by Walsh (2000), Faust and Svensson (2001, 2002), Gaspar et al. (2006, 2010) and Molnár and Santoro (2010), for example.

In reality, monetary policy is often conducted under imperfect information of various sorts. On the one hand, policymakers may face uncertainty about the state of the business cycle, the nature of structural relationships in the economy or a lack of access to timely data. But on the other hand, policymakers are at times also privy to confidential information, for example arising from staff efforts in gathering and analyzing economic data or supervisory activities.[1] My paper considers the design of optimal monetary policy for the latter case, when the policymaker has private information.

In particular, I derive a general linear-quadratic (LQ) framework for the optimal discretion policy, when the policymaker is better informed. My methods are illustrated with a small New Keynesian example, that is similar to the models of Cukierman and Meltzer (1986), Faust and Svensson (2001, 2002), Jensen (2002) and Westelius (2009). To gain further intuition about the role of imperfect information, a variant of the example economy is also considered, where the true values of the economy's state variables are exogenously revealed at the end of each period. In this case, public beliefs do not emerge as dynamic state variables on their own and there is no *dynamic* belief management, whereby a current policymaker can influence future policy decisions through

[1]Evidence in favor of such private information is for example provided by Romer and Romer (2000) and Sims (2002), who show how the confidential forecasts made by the staff of the U.S. Federal Reserve System tend to outperform forecasts based on publicly available data.

the dependence of future beliefs on policies observed in previous periods. However, even in this simplified setting, where the imperfect-information problem is purely static, discretionary policy is affected by what Walsh (2000) called the "disciplinary channel" of imperfect information in an economy where current outcomes depend on forward-looking behavior of the private sector.[2]

The problem of "public learns about central bank" studied here is distinct from settings of "policymaker learns about economy" analyzed by Sargent (1999), Aoki (2003) or Svensson and Woodford (2004). In the latter settings, atomistic individuals take policy as given without regard for inference problems faced by the policymaker, and policy constraints like the Phillips curve are largely preserved. Here however, when the "public learns about the central bank", optimal policy must take into account the public's inference problem, since the central bank is a strategic, not an atomistic player. [3] This changes the policy constraints in non-trivial ways and belief management becomes an integral part of discretionary policy.[4] While the framework adopted here exclusively assigns the policymaker, and not the public, with superior information, reality is probably best described by dispersed information, endowing different bits and pieces of hidden knowledge to the private sector and policymakers. By studying the polar case of a strictly better informed policymaker, I focus on the strategic effects arising from the public's information problem on the design of optimal policy, whereby policy actions gain an additional role as information signals. By assuming a homogeneously informed private sector, I abstract from the potential role of monetary policy in coordinating dispersed private sector beliefs, which has been highlighted for example by Morris and Shin (2002, 2005).[5]

Only Markov-perfect equilibria are considered here. In the spirit of "bygones are bygones",

[2]A similar information structure has also recently been used by Tang (2014).

[3]Such issues do not arise in the studies of Collard and Dellas (2010, 2007), Collard et al. (2009), and Orphanides and Williams (2005) who take the design of policy as given when studying the economic consequences of "public learns about central bank".

[4]Of course, in equilibrium inflation always depends on monetary policy. But in perfect information models of discretionary policy, this dependence occurs mostly indirectly and in ways beyond the control of a current period policymaker.

[5]Morris and Shin (2002, 2005) focus on static problems of dispersed information. The analysis of optimal policy in dynamic models of dispersed information is complicated by the infinity of lagged higher-order expectations that is typically relevant for the characterization of equilibrium outcomes in such settings. Nimark (2008), Berkelmans (2011) and Kohlhas (2013) study the effects of monetary policy in the presence of dynamic higher-order expectations when policy is given by an interest rate rule, without explicit consideration of the policymaker's problem.

Markov-perfect state variables must be relevant for current payoffs. When the public is imperfectly informed, its prior beliefs matter for public payoffs and they become a state variable of the policy problem, which is influenced by policy actions. By managing this state of beliefs, the policymaker responds (indirectly) to past policies, even when reputational mechanisms in the form of history-dependent strategies, known from Barro and Gordon (1983), Chari and Kehoe (1990), or Atkeson et al. (2010) are excluded from the analysis.

My specific contribution is to model these issues in a fairly general, but also tractable way. With a few exceptions, problems of this kind have mostly been analyzed in highly stylized and often static settings.[6] But the models used for policy analysis are typically dynamic and of larger scale. A key complication for models with imperfect information is to track the distribution of public beliefs. The tractability of the solution method presented here stems from the unobservable states following smooth, Gaussian processes in a linear-quadratic setting.[7] In this case, the public's uncertainty about unobserved state variables is characterized by normal distributions, arising from the solution of a Kalman filtering problem. Since shocks are assumed to be homoscedastic, changes in beliefs are then completely summarized by the linear evolution of posterior means in the Kalman filter.[8] Except for the opposite nature of the informational asymmetry, my linear-quadratic setting is similar to Svensson and Woodford (2003; 2004), Aoki (2006) who study optimal policy when the central bank is imperfectly informed.

Imperfect information gives also rise to information shocks as a source of business cycle fluctuations. As in Lorenzoni (2009), Angeletos and La'O (2009) such shocks shift public percep-

[6]Examples of static models — including also models of repeated play in one-period games — are the classic contributions by Backus and Driffill (1985a), Canzoneri (1985) and Cukierman and Liviatan (1991). More recent work includes the papers by Ball (1995) and Walsh (2000), Geraats (2007) and Sibert (2009). Dynamic models, but relatively small in size, are studied by Cukierman and Meltzer (1986), Jensen (2002), Faust and Svensson (2001, 2002) as well as Gaspar et al. (2006, 2010) and Molnár and Santoro (2010).

[7]In contrast to the Gaussian signal extraction problems studied here, discrete regime switches are an attractive alternative for modeling central bank "types" like weak/soft or commitment/discretion as in Backus and Driffill (1985b), Cukierman and Liviatan (1991), Ball (1995), Walsh (2000) and King et al. (2008). Unobserved regime switches lead to however to important non-linearities in the public's inference problem, which complicate the constraints in an optimal policy problem considerably.

[8]In general, normal distributions like the posteriors generated by the Kalman filter are characterized by their first two moments (means and variances). In the homoscedastic case, posterior variances are however constant and changes in posterior beliefs are captured by changes in the posterior means alone.

tions about fundamentals, whilst their actual values remain unchanged. Most recently, Angeletos et al. (2013) have documented evidence suggesting a persuasive empirical role of informational shocks in accounting for business cycle fluctuation in U.S. data. Non-fundamental shocks that affect economic outcomes are, of course, relevant for optimal policy design as illustrated within my framework further below.

The remainder of this paper is structured as follows. Section 2 describes the optimal discretion policy, when the policymaker is better informed than the public in a general linear-quadratic framework. The solution is illustrated with a simple New Keynesian model with stochastic output targets in Section 3. The role of information shocks is discussed in Section 4. Section 5 concludes the paper. Analytical derivations for the results of Section 3 are given in Appendix A at the end of this paper. A separate technical appendix describes the concepts and methods for the general case in more detail.

2. THE LINEAR-QUADRATIC FRAMEWORK

2.1. Dynamic information problem

This section presents a generic, linear-quadratic framework for studying optimal policy choices (under discretion) when the policymaker has private information. Attention is limited to a Markov perfect, discretionary policy problem.[9] As will be seen below, a general result is how the private sector's prior expectations — which will be referred to as "beliefs" throughout this paper — emerge as endogenous state variables in the policy problem. The time-consistent choices of the policymaker react to these public beliefs, imparting history dependence on the optimal policy under discretion.

The general setting is similar to the class of linear-quadratic economies studied by Svensson and Woodford (2003; 2004); except for the reverse asymmetry in the partial problem. There are four types of variables:

[9] If not otherwise indicated, the term "optimal policy" shall always refer to the optimal discretion policy — that is the optimal, time-consistent, Markov-perfect policy — which is typically not identical to the optimal commitment policy, since the latter is typically time-inconsistent in the kind of forward-looking economies considered here.

1. Backward-looking variables, X_t
2. Policy controls, U_t
3. Publicly observable variables, Z_t
4. Forward-looking decision variables of the private sector, Y_t

These variables will be treated as vectors of dimensions N_x, N_u, N_y, and N_z respectively. The backward-looking variables can capture exogenous forcing variables but also endogenous states like capital, habits, or lagged variables. The backward-looking variables evolve as

$$X_{t+1} = A_{xx}X_t + A_{xy}Y_t + B_xU_t + Dw_{t+1} \tag{1}$$

where w_{t+1} is an exogenous N_w-dimensional white noise process with variance $Ew_tw_t' = I$.

The policymaker observes the entire history of w_t, denoted w^t and will thus have complete information about the realization of all variables until time t. In contrast, the private sector observes only a linear combination of policy controls and backward-looking variables:

$$Z_t = C_xX_t + C_uU_t \tag{2}$$

Typically, the policy instrument will be directly observable, which can be represented by setting a sub-vector of Z_t equal to U_t. For example, in the simple model presented in the next section, the policy control will constitute the only observable and we will have $Z_t = U_t = U_{t|t}$. The history $Z^t = \{Z_t, Z_{t-1}, Z_{t-2}, \ldots\}$ spans the public information set.[10] For any variable v_t, $v_{t|t} \equiv E(v_t|Z^t)$ denotes the expectation of v_t conditional on the private sector's information set. Synonymously these expectations will be called "public beliefs" or just "beliefs".[11] In particular, $X_{t|t-1}$ are the prior beliefs about X_t before observing Z_t.

[10] In addition, there is no uncertainty about the structure of the economy — including the equilibrium behavior of the policymaker as explained below — and the public will know all parameters of the model, for example the matrices A_{xx}, A_{xy}, B_x and D of equation (1).

[11] Also, private agents will synonymously be referred to as "the public" or "the private sector". Throughout the paper, the public is assumed to be homogeneously informed.

A sufficient condition to ensure superior information of the policymaker is that $N_z < N_w$ which prevents Z^t from spanning w^t. In principle, the public may not even be able to completely recover lagged values of the structural shocks (w_{t-k}) or lagged backward-looking variables (X_{t-k}) from Z^t, which makes the information problem "dynamic" as current signals (Z_t) remain informative about past conditions. Further below, I will also study the consequences of a simplified, purely static information problem, where the true values of lagged policy actions and lagged backward-looking variables will always be revealed, thus reducing the scope of the signaling channel for monetary policy to signals about the realization of current structural shocks (w_t) alone. This simplified setup will mimic the role of the signaling channel studied for example by Walsh (2000) and Tang (2014).

Public decisions are based on public information and thus $Y_t = Y_{t|t}$ always holds by construction. Since Y_t is restricted to lie in the span of Z^t, adding Y_t to the measurement vector would not provide any new signal to the public's information set. The optimality conditions of private sector behavior are represented by an expectational linear difference equation involving only publicly observable variables and public sector expectations:

$$A_{yy}^1 Y_{t+1|t} = A_{yy} Y_{t|t} + A_{yx} X_{t|t} + B_y U_{t|t} \tag{3}$$

The policymaker seeks to minimize the expected present value of current and future losses

$$E_t \sum_{k=0}^{\infty} \beta^k L_{t+k} \qquad \text{with} \quad L_t = \begin{bmatrix} X_t \\ Y_t \\ U_t \end{bmatrix}' Q \begin{bmatrix} X_t \\ Y_t \\ U_t \end{bmatrix} \tag{4}$$

where Q is a positive definite matrix, and the expectation operator E_t conditions on w^t.[12]

Attention is limited here to Markov-perfect equilibria, which exclude reputational mechanisms via the kind of history-dependent strategies considered by Barro and Gordon (1983) or Chari and

[12] In principle, one could also allow for public beliefs $X_{t|t}$ and $U_{t|t}$ to enter the loss function. Except for adding algebraic complexity, this would not raise any further methodological issues. Likewise, linear terms in $X_{t|t}$ and $U_{t|t}$ could be added to the transition equation for the backward-looking variables. In its current form, the loss function (4) depends on public beliefs via $Y_t = Y_{t|t}$.

Kehoe (1990). In the spirit of "bygones are bygones", state variables in a Markov-perfect equilibrium must be relevant for current payoffs.[13] Under perfect information, only the backward-looking variables would qualify here as Markov-perfect state variables. In this case, a current decision-maker can influence a future decision-maker only if some of the backward-looking variables were endogenous, like capital or government debt. Under hidden information, the public's prior beliefs become part of the Markov states since they matter for the public's current decisions and payoffs. The prior beliefs, and not the posterior, beliefs enter the state vector of the policy problem, since the latter will be formed after observing current data which is influenced by current policy.[14]

In principle, the entire distribution of public beliefs needs to be tracked by the policy problem. The framework presented here affords a considerable simplification, which makes the problem very tractable: The model is cast in a Gaussian framework with constant variances, and the equilibrium is required to be linear and time-invariant. In this case, tracking entire distributions collapses to following changes in their means (and knowing their constant variances), which is easily provided by the Kalman filter. As a further simplification, it is also assumed that the private sector can be represented by a set of aggregate equations, see (3), based on a homogeneous information set (Z^t). This abstracts from the potential role for policy to coordinate between heterogeneously informed private-sector agents studied by Morris and Shin (2002, 2005), and Lorenzoni (2010) — though only in the context of static information problems — in order to emphasize the effects on the design of optimal policy resulting from the signaling role of policy actions in the formation of public beliefs.

Discretionary policy is time-consistent. At each point in time the policymaker can reoptimize while having to take future policy decisions as given. As shown in the technical appendix, the optimization problem can be cast as a dynamic program, whose state variables are the backward-looking variables (X_t) and the prior beliefs of the public ($X_{t|t-1}$). Crucially, these public beliefs reflect past policy actions through their effect on the history of observables Z^t, and these beliefs are

[13] Persson and Tabellini (2000, Chapter 11) review applications of Markov-perfect equilibria to macroeconomic policy problems.

[14] From a purely econometric perspective, the resulting state space system could equivalently by described in terms of the posteriors $X_{t|t}$ instead of $X_{t|t-1}$.

absent from the policy problem under perfect information, where Markov-perfect policies respond solely to the publicly observed, true values of the backward-looking variables; as in Oudiz and Sachs (1985) and Söderlind (1999).

In a rational expectations equilibrium the public forms its posterior beliefs consistently with the optimal policy function. The policymaker is free to choose policies which are inconsistent with the public's belief system, but equilibrium requires deviations from the equilibrium policy not to be optimal. The optimal, time-consistent policy and private sector decisions are then linear functions of the Markov states:

$$U_t = F_1 X_t + F_2 X_{t|t-1} \qquad\qquad Y_t = G_1 X_t + G_2 X_{t|t-1} \tag{5}$$

In contrast, under perfect information the optimal discretion policy and associated outcomes would depend only on the actual values of the backward-looking variables (X_t).

A detailed derivation of this optimal policy and the coefficient matrices F_1 and F_2 has been relegated to the technical appendix. But two aspects of the optimal policy are particularly worth mentioning: First, despite the linear-quadratic structure of the problem, the optimal policy coefficients (F_1 and F_2) are not certainty equivalent. They depend on the entire information structure, including the relative variances and covariances of the underlying driving variables.[15] Second, since the public information set depends on observed policies, via (2) above, public beliefs $X_{t|t-1}$ are an *endogenous* state variable, through which past policies affect current and future policies.[16] For example, even when the backward-looking variables are purely exogenous, the optimal discretion policy will exhibit history dependence via its responses to the public's prior beliefs, which is not the case under perfect information.

[15]This is a common theme in models of dispersed information, as in Lorenzoni (2010) or Berkelmans (2011).

[16]As described in a separate technical appendix, the evolution of beliefs is given by the updating equation of a Kalman filter that has the form $X_{t+1|t} = \boldsymbol{A_1} X_{t|t-1} + \boldsymbol{A_2} K(Z_t - Z_{t|t-1})$ for some $\boldsymbol{A_1}$, $\boldsymbol{A_2}$. and Kalman gain K. This updating equation enters the policymaker's policy problem as transition equation for the endogenous state of beliefs.

2.2. Static information problem

A crucial aspect of the general class of models outlined above is the emergence of endogenous beliefs $X_{t|t-1}$ and their relevance for characterizing equilibrium policies and outcomes. However, it is also instructive to consider a simplified information where lagged backward-looking variables and policy actions are directly revealed to the public. In particular, consider the case where public beliefs are not only conditioned on Z^t as defined in (2) but also on $E_{t-1}X_t \equiv E(X_t|w^{t-1})$.

The availability of $E_{t-1}X_t$ in the public information set at time t reduces the public's information problem to a static signal extraction effort, since the only unobserved component of the backward-looking variables are then the serially uncorrelated shocks (w_t). Computationally, this obviates the need for nesting a dynamic Kalman filter inside the policymaker's regulator problem. This even allows the derivation of analytical results for a small-scale model in Section 3.

More fundamentally, the Markov-perfect state vector will be different under this simplified information structure as $E_{t-1}X_t$ replaces $X_{t|t-1}$. Optimal policy and outcomes will thus respond only to X_t and $E_{t-1}X_t$

$$U_t = F_1 X_t + F_2 E_{t-1} X_t \qquad\qquad Y_t = G_1 X_t + G_2 E_{t-1} X_t \tag{6}$$

using different coefficient values than in (5). In particular, when the backward-looking variables are purely exogenous, partial information does not add any history dependence to the optimal policy.[17] Even when the backward-looking variables are partly endogenous — for example in the case of indexation in the Phillips curve or the presence of habit formation — this static partial information structure will not impart history dependence on policy choices above and beyond the full-information case.

[17]The backward-looking variables are exogenous when $A_{xy} = 0$ and $B_x = 0$ in (1).

3. A SIMPLE EXAMPLE

This section uses a very simple textbook version of the New Keynesian model to illustrate the optimal discretion policy under hidden information. This simple model is a natural benchmark, since different variants of it been studied already by Cukierman and Meltzer (1986), Faust and Svensson (2001, 2002), who found evidence that imperfect information might actually lead to better outcomes (under discretion) as measured by policy losses or social welfare. Moreover, the New Keynesian model has been extensively used for the study of optimal policy under perfect information — see for example Clarida et al. (1999) or the monograph by Woodford (2003).

The model is largely identical to what is used in textbook treatments of the perfect-information case of optimal policy in the simple New Keynesian model as in Clarida et al. (1999), Walsh (2003), Woodford (2003) or Galì (2008). The only difference is a stochastic preference shock to the policymaker's objective function, which is unobservable to the public.[18] Westelius (2009) has also studied the consequences of imperfect information in this setting, but in his case, policy is prescribed by the discretionary targeting rule that would be optimal under full information, without considering its (sub-)optimality and time-(in)consistency under imperfect information.[19]

A key feature of this simple New Keynesian model is that inflation is purely determined by public expectations of current and future policies, which puts centerstage the public's concerns about the policymaker's intentions. Moreover, the backward-looking variables are all exogenous, such that the introduction of belief management will be the only possible source of history dependence in the imperfect information case.

[18] This stochastic preference shock replaces the cost-push shocks commonly used for causing a policy-trade off in the simple New Keynesian model. By keeping the number of exogenous shocks to a minimum, this approach simplifies the analysis of the imperfect information case since the level of a cost-push shock would have to be contemporaneously observable by the private sector (due to its presence in the Phillips curve and the assumption of a homogeneously informed private sector). Together with an equally observable output gap, this would then require at least three shocks for the existence of an interesting signal extraction problem.

[19] A short note, comparing the two approaches and their results is available at the author's website.

3.1. The Setup

In this simple model, aggregate decisions of the private sector are represented by the New Keynesian Phillips curve. Furthermore, for a given output gap policy, the computation of a short-term interest rate via the New Keynesian IS curve is redundant, and the output gap can be used directly as the policy control.

The private sector is populated by a continuum of identical firms and households, which trade goods and labor services. There is no capital accumulation and output equals consumption. Firms are monopolistically competitive and use staggered price-setting as in Calvo (1983). Optimal pricing decisions lead to the New Keynesian Phillips curve as in Yun (1996) and King and Wolman (1996). The log-linearized Phillips curve is

$$\pi_t = \beta \pi_{t+1|t} + \kappa\, x_t \tag{7}$$

where π_t is inflation and x_t is the output gap.[20] The parameter β is the representative agent's discount factor — assumed to be positive but smaller than one — and κ is a reduced form parameter influenced amongst others by the frequency of price-setting (Galì, 2003, p. 159). As in Section 2, for any variable v_{t+1}, $v_{t+1|t}$ denotes the public's conditional expectation using its time t information set. (Below, I will consider three different information structures — including full information.) The policymaker seeks to minimize a present value of expected losses

$$E_t \sum_{k=0}^{\infty} \beta^k L_{t+k} \qquad\qquad L_t = \alpha_x (x_t - \bar{x}_t)^2 + \alpha_\pi\, \pi_t^2 \tag{8}$$

where the weights have been normalized to $\alpha_\pi = 1 - \alpha_x$ and $0 \leq \alpha_x \leq 1$. The expectations operator E_t reflects the policymaker's information set, conditioning on the full-information set spanned by w^t. The non-standard feature of the loss function is the time-varying target for the

[20] As it is standard in the New Keynesian model, the output gap measures the difference between actual output and the hypothetical level of output if there were no nominal frictions (Galì, 2003). Also, throughout the paper, all variables are in log-deviations from steady state, which implicitly assumes the existence and uniqueness of a steady state under discretionary policy.

output gap, $\bar{x}_t$, which will be specified as an exogenous stochastic process.

In principle, one could think of various ways to motivate the presence of $\bar{x}_t$ in the loss function. However, the information structure used below will require that $\bar{x}_t$ is not observed by the private sector.[21] To keep the model close to the New Keynesian benchmark, I maintain the assumption of a homogeneously informed private sector and follow Cukierman and Meltzer (1986) who interpret the output target as arising from time-varying preferences of the policymaker. Under this view, $\bar{x}_t$ represents the outcome of political influences on monetary policy to stimulate the economy. These preferences are assumed to vary exogenously with political representation in the government and the makeup of central banker's preferences.[22] Absent further specification of the underlying heterogeneity implied by variations in $\bar{x}_t$, the loss function (8) does however not necessarily represent a social welfare function and a micro-founded derivation — including consideration of the appropriate information set for welfare evaluation — is beyond the scope of this paper.[23]

In order to make the public's signal extraction interesting, the output gap target is assumed to be driven by two components, one persistent, one transitory:[24]

$$\bar{x}_t = \gamma_t + \varepsilon_t \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2) \tag{9}$$

$$\gamma_{t+1} = \rho\, \gamma_t + \eta_{t+1} \qquad \eta_t \sim N(0, \sigma_\eta^2) \quad \text{and} \quad |\rho| < 1 \tag{10}$$

[21] This assumption for example precludes the notion that $\bar{x}_t$ might measure an output gap defined relative to a first-best benchmark economy — and thus relevant for social welfare — whereas x_t might measure an output gap defined relative to a benchmark economy absent nominal frictions — and thus relevant for the Phillips curve — but subject to other real distortions, unless the policymaker could plausibly be better informed than a homogeneously informed private sector about the nature of such distortions.

[22] Such hidden pressures could arise even when the independence of the central bank is formally enshrined in law, since actual independence is a more fragile concept. For example Abrams (2006) gives a striking account of hidden but forceful policy influences. His study documents how U.S. President Nixon covertly pressured the then Chairman of the Federal Reserve, Arthur Burns to ease policy in the run-up to the Great Inflation. Under either interpretation, the output target captures a form of heterogeneity otherwise not present in the model.

[23] In a similar model, Faust and Svensson (2001) use for example a notion of representative welfare that corresponds to evaluating (8) at the average output target (here: zero), $L_t^R = \alpha_\pi \pi_t^2 + \alpha_x x_t^2$. As it turns out in this simple example, there will not be a conflict between ranking outcomes under the policymaker's objective as opposed to representative welfare in the sense of Faust and Svensson (2001).

[24] For the sake of brevity the term "transitory" will be used synonymously with "serially correlated" in this paper. Having said that, none of the two target shocks have permanent effects, and, in a broader sense, both shocks could thus be qualified as transitory as well.

The remainder of this section will consider the implications of three different information structures for the design of optimal discretionary policy:

Full information where the public directly observes the individual components of the output gap target (γ_t and ε_t).

Imperfect information, with known lagged states where the public observes the current policy x_t and the lagged values of the individual target components (γ_{t-1} and ε_{t-1}); thus corresponding to the framework laid out in Section 2.2 above.

Imperfect information in the general case where the public observes only the current policy x_t and remains uncertain about current and lagged values of the individual target components, corresponding to the framework of Section 2.1.

For either case, the model maps easily into the general LQ framework of Section 2 as shown in the technical appendix.

3.2. Full Information

Under discretion, a current policymaker takes the mapping between future policies and future states as given. In this simple model, the only backward-looking variables are the exogenous components γ_t and ε_t for the output gap target. When both components are directly observable, private-sector forecasts about future target values are exogenous, too, and policy actions do not convey any additional information. Under these circumstances, future policies and outcomes are beyond the control of a discretionary, current-period policymaker. The policymaker's first-order conditions then yield a myopic targeting rule that relates only current values of output gap and inflation with each other:

$$\alpha_x(x_t - \bar{x}_t) + \alpha_\pi \kappa \, \pi_t = 0 \tag{11}$$

As inflation is determined via the forward-looking Phillips curve (7), the effect of expected future policies are present in this first-order condition, but only indirectly. For a given equilibrium level of inflation, the full-information policy does not account for its effects on expected future inflation. Optimal policy and inflation outcomes are characterized by:[25]

$$x_t = \bar{f}_\gamma \gamma_t + \bar{f}_\varepsilon \varepsilon_t \qquad \bar{f}_\gamma = \frac{(1-\beta\rho)\alpha_x}{(1-\beta\rho)\alpha_x + \alpha_\pi \kappa^2} \qquad \bar{f}_\varepsilon = \frac{\alpha_x}{\alpha_x + \alpha_\pi \kappa^2} \tag{12}$$

$$\pi_t = \bar{g}_\gamma \gamma_t + \bar{g}_\varepsilon \varepsilon_t \qquad \bar{g}_\gamma = \frac{\kappa}{(1-\beta\rho)} \bar{f}_\gamma \qquad \bar{g}_\varepsilon = \kappa \bar{f}_\varepsilon. \tag{13}$$

In this Markov-perfect equilibrium, output gap policy exhibits the same persistence as the output gap target. The forward-looking Phillips curve amplifies the effect of persistent policy shocks on current inflation. Faced with a higher sensitivity of inflation to $\bar{f}_\gamma$, optimal policy chooses a more muted policy response to the persistent shock, $\bar{f}_\gamma < \bar{f}_\varepsilon$.

3.3. Imperfect information when lagged states are known

The optimal policy changes in significant ways, when the public cannot directly infer the target components, but must rather solve a signal extraction problem that uses the observed policy x_t as its input. For now, it will be assumed that the public also knows lagged values of the target components, but cannot directly observe their current values. In addition to direct observations of lagged states, the only current-period variable available to the public is assumed to be the actual output gap x_t set by the policymaker. This assumption simplifies the policy problem under imperfect information considerably, since there is no need for tracking the history of public beliefs, allowing the derivation of analytical results— in particular the result that expected policy losses are lower under hidden information. Similar results also hold in the more general case of imperfect information when lagged states are not known, but can then be established at best only under stronger assumptions or through the use of numerical methods.

[25] Appendix A.1 provides further detail about the policymaker's optimization problem.

$E_{t-1}\gamma$ is a sufficient statistic for the information content of lagged states in this model.[26] The public's information set can thus be described by the measurement vector

$$Z_t = \begin{bmatrix} x_t \\ E_{t-1}\gamma_t \end{bmatrix}. \tag{14}$$

As in the full information case, the discretionary policymaker does not commit to particular future policies and must thus take the private sector's mapping between future states and future inflation as given: $\pi_{t+1|t} = \hat{g}\gamma_{t+1|t}$, for some coefficient $\hat{g}$.[27] The Phillips curve constraint then becomes $\pi_t = \beta\rho\,\hat{g}\,\gamma_{t|t} + \kappa x_t$. Assuming a linear equilibrium, public beliefs about the persistent target component solve a static signal extraction problem,

$$\gamma_{t|t} = E_{t-1}\gamma_t + K(x_t - E_{t-1}x_t), \tag{15}$$

where the gain K and expectations $E_{t-1}x_t$ — to be derived shortly below — depend on the public's anticipation of the policymaker's behavior in equilibrium. This is a static signal extraction problem since beliefs are constructed from a linear combination of current-period signals alone. The public's signal extraction in (15) embodies the dependence of public beliefs on observed policy actions, which is a central theme of this paper.

Importantly, the discretionary policymaker is assumed to consider only the direct effect of a policy choice for x_t on beliefs in (15), not the indirect effects of the systematic behavior of policy on the gain K and expected policy $E_{t-1}x_t$. This extends the inability to commit to future actions — discussed above in the context of full information — to an inability to commit to a particular current-period reaction function between policy actions and state variables.[28]

[26]By construction, $E_{t-1}\gamma_t$ is a sufficient statistic for projections of γ_t on lagged state variables. Since ε_t is serially uncorrelated, knowledge of lagged states is irrelevant for inference about the current value of ε_t. Tracking $E_{t-1}\gamma_t = \rho\gamma_{t-1}$ instead of γ_{t-1} involves only a mere rescaling in this case. This normalization is however useful for comparison with the more general case of a dynamic signal extraction problem, leading to the emergence of public beliefs $\gamma_{t|t-1}$ as state variable.

[27]Beliefs about future values of the transitory target component can be ignored since $\varepsilon_{t+1|t} = 0$.

[28]Similar assumptions are also implicitly (or explicitly) made by Cukierman and Meltzer (1986), Faust and Svensson (2001, 2002), Walsh (2000), and Tang (2014) for example.

The private sector is assumed to (correctly) anticipate a time-invariant equilibrium, where policy is linear in a Markov-perfect state vector. In the present case, the Markov-perfect states are the true target components γ_t and ε_t as well as the information about lagged states captured by $E_{t-1}\gamma_t$, and policy has the form

$$x_t = \tilde{f}_\gamma \gamma_t + \tilde{f}_\varepsilon \varepsilon_t + f_b E_{t-1}\gamma_t \tag{16}$$

$$= \tilde{f}_\gamma \eta_t + \tilde{f}_\varepsilon \varepsilon_t + \hat{f}\ E_{t-1}\gamma_t \qquad (\hat{f} = \tilde{f}_\gamma + f_b). \tag{17}$$

The public assumes particular values for the above coefficients when solving its signal extraction problem; in equilibrium these coefficient values must coincide with the policymaker's optimal choices. Equation (17) reflects a convenient rotation of the state variables into orthogonal components whose sum equals the total target level for the output gap; $\bar{x}_t = \eta_t + \varepsilon_t + E_{t-1}\gamma_t$.[29] Analytical solutions for the policy coefficients in (16) and (17), respectively, will be described next. Before analyzing the properties of the policy coefficients, it will be convenient to make the following, innocuous assumption:

Assumption 1 (Policy responses to each target component are positive)**.** *The three policy coefficients in (17) are positive:* $\tilde{f}_\gamma > 0$, $\tilde{f}_\varepsilon > 0$, $\hat{f} > 0$.

Assumption 1 is a natural choice in this model, since the policymaker has to trade off attaining the output gap target — by setting each of the coefficients $\tilde{f}_\gamma$, $\tilde{f}_\varepsilon$, and $\hat{f}$ close to one — and keeping inflation close to zero. The latter objective is best achieved by setting the output gap, and thus all of the three coefficients in (17), to zero. Choosing a negative response to any of the three target components in (17), would thus always result in unwanted inflation variance while missing the output gap target by even more than in the case when the policymaker ignored the target and kept the output gap at zero.[30]

[29]Since $E(\varepsilon_t \eta_t) = E(\eta_t\, E_{t-1}\gamma_t) = 0$, the variance of the output gap equals $\mathrm{Var}\,(x_t) = \tilde{f}_\gamma^2 \sigma_\eta^2 + \tilde{f}_\varepsilon^2 \sigma_\varepsilon + \hat{f}^2\,\mathrm{Var}\,(E_{t-1}\gamma_t)$ with $\mathrm{Var}\,(E_{t-1}\gamma_t) = \sigma_\eta^2 \cdot \rho^2/(1-\rho^2)$.

[30]Notice that Assumption 1 does not preclude a negative coefficient $f_b = \hat{f} - \tilde{f}_\gamma$ in (16). As will be seen below, this is actually the case under optimal policy.

Given the linear policy in (17), beliefs $\gamma_{t|t}$ and thus also inflation can be expressed as a linear combination of the public's observables at time t:

$$\gamma_{t|t} = (1 - K\hat{f})E_{t-1}\gamma_t + Kx_t \qquad \text{with} \quad K = \frac{f_\gamma \sigma_\eta^2}{f_\gamma^2 \sigma_\eta^2 + f_\varepsilon^2 \sigma_\varepsilon^2} > 0, \tag{18}$$

$$\pi_t = \beta\rho\,\hat{g}(1 - K\hat{f})E_{t-1}\gamma_t + \tilde{\kappa}x_t, \qquad \tilde{\kappa} = \kappa\left(1 + \frac{\beta\rho}{1-\beta\rho}K\hat{f}\right) > \kappa. \tag{19}$$

where the inequalities are assured by Assumption 1. Under imperfect information, public expectations of future inflation, $\pi_{t+1|t}$, react directly to policy choices for the output gap, raising the Phillips curve slope from κ to $\tilde{\kappa}$. The hidden-information slope $\tilde{\kappa}$ reflects a sort of average persistence of the output gap, that is shaped by the relative occurrence of each shock and the strength of the policy response to each of them. The increase in the slope coefficient corresponds to the extent to which hidden information makes the discretionary policymaker internalize the effect of current policy on inflation while taking the remaining expectations component — captured by the term in $E_{t-1}\gamma_t$ in (19) — as given.

The optimal policy problem under imperfect information is given by the objective function (8), the Phillips curve constraint (19) and the definition of the target processes (9). Key properties of the optimal policy are summarized in the following proposition.

Proposition 1 (Optimal policy under imperfect information when lagged states are known)**.** *In the simple New Keynesian model with imperfect information, but known lagged states, optimal discretionary policy is characterized by a myopic targeting rule*

$$\alpha_x(x_t - \bar{x}_t) + \alpha_\pi \tilde{\kappa}\pi_t = 0, \tag{20}$$

which is identical to the full information rule (11) except for reflecting the steeper slope coefficient $\tilde{\kappa}$ of the Phillips curve (19) under imperfect information.

The optimal policy is linear as in (17). Notably, optimal policy responds identically to transi-

tory and persistent target shocks on impact, $\tilde{f}_\gamma = \tilde{f}_\varepsilon \equiv \tilde{f}$. The optimal policy coefficients are

$$\tilde{f} = \frac{\alpha_x}{\alpha_x + \alpha_\pi \tilde{\kappa}^2}, \qquad \hat{f} = \frac{(1-\beta\rho(1-\hat{R}^2))\alpha_x}{(1-\beta\rho(1-\hat{R}^2))\alpha_x + \alpha_\pi\tilde{\kappa}^2}, \tag{21}$$

$$\text{with} \quad \tilde{\kappa} = \kappa \cdot \left(1 + \frac{\beta\rho}{1-\beta\rho}\hat{R}^2\right), \qquad \hat{R}^2 = K \cdot \hat{f} = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_\varepsilon^2} \cdot \frac{\hat{f}}{\tilde{f}}. \tag{22}$$

For a given ratio $\phi \equiv \hat{f}/\tilde{f}$, the coefficients $\tilde{f}$ and $\hat{f}$ are uniquely given by (21) and (22). The ratio ϕ is determined by a second-order polynomial whose only positive root is smaller or equal to one. Since $\tilde{f} > 0$, the negative root of the polynomial would imply $\hat{f} < 0$, which is ruled out by Assumption 1; furthermore, the corresponding policy would be dominated by alternative policies leading to lower losses.

Proof. See Appendix A.2 □

There is thus a unique, linear time-invariant optimal policy with $\tilde{f} > \hat{f}$.[31] While $\hat{f}$ measures the total response of policy to lagged knowledge about the persistent target component, it is also instructive to consider the marginal response $f_b = \hat{f} - \tilde{f} < 0$ defined in (16), which measures the policy response to $E_{t-1}\gamma_t$, holding the optimal response to the actual level γ_t fixed.[32] The negative, marginal policy response to $E_{t-1}\gamma_t$ counteracts the otherwise positive effects of $E_{t-1}\gamma_t$ on inflation.[33] While a highly negative f_b would also generate a negative inflation response, the optimal trade-off prescribed by the targeting rule (20) balances a negative marginal response of the output gap to $E_{t-1}\gamma_t$ against a positive inflation response to $E_{t-1}\gamma_t$ (for a given level of γ_t).[34]

An important parameter for the hidden information problem is the relative share of persistent to transitory target shocks, R^2. Under full information, the optimal policy in this linear-quadratic model is certainty equivalent, and R^2 plays no role in determining optimal policy. Considering

[31] Blake and Kirsanova (2012) study the potential of multiple equilibria in linear-quadratic economies and find that a unique equilibrium is assured when all state variables are exogenous, as here the case.

[32] While the full-information benchmark for $\hat{f}$ is $\bar{f}_\gamma$, the corresponding reference value for f_b is zero, since lagged information about γ_t does not matter given perfect knowledge about its current value.

[33] Recall that $1 - K\hat{f} = 1 - \hat{R}^2 > 0$ in (19).

[34] Denoting the marginal response of inflation to $E_{t-1}\gamma_t$ by g_b, the targeting rule (20) implies $\alpha_x f_b + \alpha_\pi\tilde{\kappa}g_b = 0$ which requires f_b and g_b to have opposite signs.

the limiting behavior of the imperfect-information equilibrium, the slope of the Phillips curve approaches its full-information values κ and the policy response $\tilde{f}$ approaches the full-information response to a transitory shock $\bar{f}_\varepsilon$, when R^2 goes to zero and transitory shocks become the dominant source of target changes; see Proposition A.1 in the appendix.

However, when persistent target shocks become dominant as R^2 goes to one, the limiting behavior of the imperfect information equilibrium does not approach the outcomes induced by a persistent shock under full information. As persistent shocks become the sole source of variations in the output gap target, the optimal discretionary policy rather converges towards a policy that minimizes expected losses amongst linear rules that respond only to γ_t; see Proposition A.1 in the appendix. This discontinuity is revealing about the mechanism at play in the hidden-information economy: Through the dependence of expected inflation on beliefs, which again depend on observed policy via (18), optimal policy internalizes at least in part its effects on inflation expectations, and increasingly so the higher R^2.[35]

The optimal targeting rule under imperfect information differs from the full-information case only in using the Phillips-curve coefficient $\tilde{\kappa}$ rather than κ. The discretionary policymaker takes $\tilde{\kappa}$ as given, and the optimal policy response to either target shock is identical to the optimal full-information response to a serially uncorrelated shock when the Phillips-curve slope were $\tilde{\kappa}$. Everything else equal, the higher sensitivity of inflation to the current output gap raises the marginal policy loss from pursuing a given level of the output gap target under hidden information, which makes it less attractive for the policymaker to set the output gap close target.

According to the targeting rule (20), the optimal shortfall of the output gap from its target has to be balanced against the level of inflation in equilibrium. Since inflation depends not only on the current output gap, but also on the policy response to $E_{t-1}\gamma_t$, the worsening in the trade-off between inflation and the output gap due to $\tilde{\kappa} > \kappa$ need not necessarily translate into more muted policy responses to target shocks. The differences in the responses of output gap and inflation under perfect and imperfect information are summarized in the next proposition.

[35]Faust and Svensson (2001, 2002) find similar discontinuities in the limiting behavior of optimal policy as hidden information gets resolved when the distribution of noise shocks shrinks to zero.

Proposition 2 (Comparison of policy coefficients under imperfect and full information)**.** *Optimal discretionary policy in the case of imperfect information when lagged states are known is characterized by outcomes for output gap and inflation of the form*

$$x_t = \tilde{f}(\eta_t + \varepsilon_t) + \hat{f}E_{t-1}\gamma_t \qquad \pi_t = \tilde{g}(\eta_t + \varepsilon_t) + \hat{g}E_{t-1}\gamma_t. \tag{23}$$

The policy coefficients compare against their full-information counterparts given in (12) and (13) as follows: $\tilde{f} < \bar{f}_\varepsilon$, $\hat{f} < \bar{f}_\gamma$; and the inflation coefficients obey $\tilde{g} < \hat{g}_\gamma < \bar{g}_\gamma$. Depending on parameters $\tilde{f}$ can be larger or smaller than $\bar{f}_\gamma$, and $\tilde{g}$ can be larger or smaller than $\bar{g}_\varepsilon$.[36]

Proof. See Appendix A.2 □

Under hidden information, the public cannot disentangle transitory shocks — with otherwise benign inflation effects under full information — from persistent shocks. As a result, the policy response to the transitory target shock is strictly lower for all admissible parameter values of the model. But depending on whether the policy coefficient $\tilde{f}$ is sufficiently lower than $\bar{f}_\varepsilon$ to offset the increased steepness of the Phillips Curve, the corresponding inflation response to a transitory target shock can then be higher or lower than in the full-information case.

Since policy responses to persistent target shocks can only "benefit" from being confused with transitory changes in the output gap, the inflation response to a persistent target shock is strictly lower than under full information. Depending on parameter values, optimal policy may support this outcome by choosing a lower output gap response to a persistent target shock, but for other parameter values — notably low R^2 and low α_x as shown in a separate technical appendix— policy may also take advantage from the lower inflation response and rather strive for tracking the output gap target more closely than in the full-information case by choosing a value of $\tilde{f}$ that exceeds $\bar{f}_\gamma$.

For any set of parameter values, optimal policy under imperfect information responds less strongly to the part of the persistent target that is observed exogenously ($E_{t-1}\gamma_t$) than it would under full information, accompanied by an unambiguously lower inflation response. The lower

[36]The relevant model parameters are $0 < \beta < 1$, $|\rho| < 1$, $0 < \kappa$, $0 \leq \alpha_x \leq 1$, $\alpha_\pi = 1 - \alpha_x$ and $0 < R^2 < 1$. The limiting behavior of outcomes as R^2 tends to zero or one are discussed in the appendix.

value of $\hat{f}$ compared to $\bar{f}_\gamma$ is due to the appearance of $E_{t-1}\gamma_t$ in the Phillips curve (20), where this known component of γ_t acts as a shifting intercept that raises the inflation response to $E_{t-1}\gamma_t$.[37]

The imperfect-information setting studied here, considers optimal policy choices while taking it as given that the public cannot observe any of the target components directly. In a wider sense, the direct revelation of the target components — effectively instituting a full-information regime — constitutes a policy choice as well, and it is instructive to compare policy losses arising in either case. The comparison of such regimes is best conducted from an ex-ante perspective, made independently of a particular realization shocks and a handy measure is provided by the unconditional expectation of policy losses $E(L_t)$, or shortly "expected losses".[38] Comparing outcomes under perfect and hidden information, some of the hidden-information responses are closer to the policymaker's targets for output gap and inflation than in the full information case, while some others are not (and a few others can be higher or lower than in the full-information case depending on the parametrization of the problem).[39] Nevertheless, for all admissible parameter values, the unconditional expectation of policy losses under imperfect information turns out to be lower than in the full-information case.[40]

Proposition 3 (Lower policy losses under imperfect information)**.** *Expected losses and the variance of inflation generated by optimal policy are lower in the hidden-information case when compared against full-information outcomes:* $E(L_t^{FI}) > E(L_t^{HI})$ *and* $\mathrm{Var}\,(\pi_t^{FI}) > \mathrm{Var}\,(\pi_t^{HI})$.

Proof. See Appendix A.2 □

[37]As a side effect, the lower $\hat{f}$ also reduces the equilibrium value of the Phillips curve slope $\tilde{\kappa}$ through $\hat{R}^2$ by lowering the public's anticipation of future policies; however this is not internalized by the discretionary policymaker when taking first-order conditions.

[38]Expected losses $E(L_t)$ integrate over all possible realization of the appropriate state vector in each setting based in its ergodic distribution. Nevertheless, notice that the optimization problems studied adopt the usual convention of maximizing the state-dependent value function $V_t = \max L_t + \beta E_t V_{t+1}$, Damjanovic et al. (2008) derive optimal policies designed to minimize unconditionally expected losses $E(L_t) = E((1-\beta)V_t)$, albeit for the full-information case.

[39]For example, the unambiguously lower output gap responses to $E_{t-1}\gamma_t$ and ε_t induce larger shortfalls from the output gap target, while the lower inflation responses to $E_{t-1}\gamma_t$ and ε_t contribute to lower policy losses.

[40]Up to a scaling, the unconditional expectation of per-period losses L_t is also identical to the unconditional expectation of the policymaker's objective function, which consists of the discounted present value of current and future losses defined in (8), since $E(L_t) = E\left(\sum_{k=0}^{\infty}(1-\beta)\cdot\beta^k L_{t+k}\right)$.

Proposition 3 compares discretionary policies under shocks that do not implement a first-best outcome (Clarida et al., 1999). As an application of the theory of the second best, adding a further friction (hidden information) to this economy generally need not result in higher losses. As shown in Appendix A.2, a central element behind the derivation of Proposition 3 is an unambiguous reduction in the variance of inflation under hidden information (integrating over all shocks), which occurs for two reasons: First, there is less variability — in fact there is zero variability — between the hidden information responses to the two target shocks than there is under full information. By smoothing shock responses across states of nature, inflation variance is reduced overall even when the conditional variance of inflation due to transitory shocks is higher in the hidden information case. Second, the average response of inflation to both target shocks is lower when comparing the hidden information case against full information.

3.4. Imperfect information (general case)

To consider a more realistic case of imperfect information, it is now assumed that the public observes only current and past policies, (x^t), without being able to directly observe past states as assumed above.[41] In this case, it becomes necessary to track the endogenous evolution of the public's prior beliefs about the output gap target ($\gamma_{t|t-1}$) and beliefs are represented by a dynamic Kalman filter rather than a static signal extraction.[42] In many aspects, public beliefs $\gamma_{t|t-1}$ assume now the role previously played by $E_{t-1}\gamma_t$, however with the crucial difference that $\gamma_{t|t-1}$ is an endogenous state variable that reflects past policies.[43] While the introduction of hidden information with known lagged states substantially affected the scale of responses to target shocks, the transmission of shocks via beliefs $\gamma_{t|t-1}$ further alters the propagation mechanism of the model. However, the added generality comes at the expense of tractability, even in this small-scale model.

[41] As before, inflation is a function of the public's information set and will thus also be perfectly known to the public, but without providing any useful signal above and beyond what is conveyed by the vector of observables, which is here $Z_t = x_t$.

[42] For brevity the two model variants under imperfect information will henceforth also be referred to as "static" and "dynamic" cases, respectively, notwithstanding the overall dynamic nature of the New Keynesian model.

[43] As before, there is no need to track beliefs about the serially uncorrelated target component since they are always zero, $\varepsilon_{t|t-1} = 0$.

By strengthening a little the assumptions on the optimal policy coefficients made in Assumption 1, many of the previous results can be established here analytically as well.

As in the static case, a time-invariant, linear equilibrium is assumed, with endogenous beliefs $\gamma_{t|t-1}$ replacing $E_{t-1}\gamma_t$ as state variable. Optimal policy then has the form

$$x_t = \tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}\tilde{\gamma}_t + \tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}\varepsilon_t + \hat{\boldsymbol{f}}\gamma_{t|t-1} \tag{24}$$

where $\tilde{\gamma}_t \equiv \gamma_t - \gamma_{t|t-1}$ measures the innovation in γ_t relative to the public's prior beliefs; policy innovations are analogously defined as $\tilde{x}_t \equiv x_t - x_{t|t-1}$.[44] Similar to the the static case, the same policy as in (24) can also be written in terms of γ_t, ε_t and $\gamma_{t|t-1}$ and the coefficient on $\gamma_{t|t-1}$, denoted $\boldsymbol{f_b} \equiv \hat{\boldsymbol{f}} - \tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}$, then measures the marginal policy response to beliefs for a given level of γ_t.

In order to derive some analytical results, it will again be useful to make some assumptions about the policy coefficients. The first part of Assumption 2 below is identical to Assumption 1 from the static case. The second part corresponds to the result $\hat{f}/\tilde{f} < 1$ that held in the static case, ensuring $\hat{\boldsymbol{R}}^{\boldsymbol{2}} < 1$ and thus $\tilde{\boldsymbol{\kappa}} > \kappa$.[45]

Assumption 2 (Policy responses to target components in the dynamic case)**.** *The three policy coefficients in (24) are positive: $\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}} > 0$, $\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}} > 0$, $\hat{\boldsymbol{f}} > 0$. In addition, the optimal response to the public's prior beliefs about the persistent target component is more muted than the response to innovations in γ_t: $\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}} > \hat{\boldsymbol{f}}$.*

When the output gap is the only signal available to the public, and provided the public correctly anticipates an equilibrium policy as in (24), optimal beliefs are characterized by a Kalman filter:

$$\gamma_{t+1|t} = \rho\gamma_{t|t} \qquad\qquad \gamma_{t|t} = \gamma_{t|t-1} + K(x_t - x_{t|t-1}) \tag{25}$$

[44] While recycling some notation from the static case of imperfect information, endogenous coefficients that take on different values in the two cases — in particular $\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}$, $\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}$, $\hat{\boldsymbol{f}}$, $\tilde{\boldsymbol{\kappa}}$ and $\hat{\boldsymbol{R}}^{\boldsymbol{2}}$ but also $\boldsymbol{R^2}$— will be typeset in bold fonts when referring to the dynamic case of imperfect information.

[45] While the condition $\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}} > \hat{\boldsymbol{f}}$ was never violated in numerical simulations (that did not impose the condition) over a fairly exhaustive range of parameter values, it proved hard to establish its general validity analytically, and so the condition has rather been imposed here.

$$\text{with} \quad \boldsymbol{K} = \frac{\text{Cov}(\tilde{\gamma}_t, \tilde{x}_t)}{\text{Var}(\tilde{x}_t)} = \frac{\boldsymbol{R^2}}{\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}} \qquad \text{and} \quad 0 < \boldsymbol{R^2} \equiv \frac{\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}^{\,2} \text{Var}(\tilde{\gamma}_t)}{\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}^{\,2} \text{Var}(\tilde{\gamma}_t) + \tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}^{\,2} \sigma_\varepsilon^2} < 1 \tag{26}$$

In the static signal extraction problem of Section 3.3, the variance ratio R^2 was exogenously given whereas here $\boldsymbol{R^2}$ endogenously depends on the choice of optimal policy.[46] The endogeneity of $\boldsymbol{R^2}$ sufficiently complicates the fixed-point analysis behind the optimal policy problem to require numerical methods for its solution. As before, it is useful to define $\hat{\boldsymbol{R}}^{\boldsymbol{2}} \equiv \boldsymbol{R^2} \cdot \hat{\boldsymbol{f}}/\tilde{\boldsymbol{f}}$, and Assumption 2 assures that $0 \leq \hat{\boldsymbol{R}}^{\boldsymbol{2}} \leq 1$. Similar to the static imperfect-information case, the Phillips curve can then be written as in (19), again featuring a larger slope coefficient than in the full-information benchmark, now denoted $\tilde{\boldsymbol{\kappa}} = \kappa(1 + \hat{\boldsymbol{R}}^{\boldsymbol{2}} \cdot (\beta\rho)/(1-\beta\rho)) > \kappa$, that raises the marginal policy loss of setting the output gap closer to target. Given the policy in (24), inflation is again linear in the states with

$$\pi_t = \tilde{\boldsymbol{g}}_{\boldsymbol{\gamma}} \tilde{\gamma}_t + \tilde{\boldsymbol{g}}_{\boldsymbol{\varepsilon}} \varepsilon_t + \hat{\boldsymbol{g}} \gamma_{t|t-1}, \qquad \tilde{\boldsymbol{g}}_{\boldsymbol{\gamma}} = \tilde{\boldsymbol{\kappa}} \tilde{f}_\gamma, \qquad \tilde{\boldsymbol{g}}_{\boldsymbol{\varepsilon}} = \tilde{\boldsymbol{\kappa}} \tilde{f}_\varepsilon, \qquad \text{and} \quad \hat{\boldsymbol{g}} = \frac{\kappa}{1-\beta\rho} \cdot \hat{\boldsymbol{g}} \tag{27}$$

and $\boldsymbol{g_b} \equiv \hat{\boldsymbol{g}} - \tilde{\boldsymbol{g}}_{\boldsymbol{\gamma}}$ measures the marginal response of inflation to beliefs for a given level of γ_t.

Proposition 4 (Optimal policy in the general imperfect-information case.)**.** *When the public's only observable signal is x_t, optimal policy is characterized by a forward-looking targeting rule:*

$$\alpha_x(x_t - \bar{x}_t) + \alpha_\pi \hat{\boldsymbol{\kappa}} \pi_t + \alpha_\pi \kappa \frac{\beta\rho}{1-\beta\rho}(1 - \hat{\boldsymbol{R}}^{\boldsymbol{2}})\hat{\boldsymbol{R}}^{\boldsymbol{2}} \beta\rho \sum_{j=0}^{\infty} \left(\beta\rho(1 - \hat{\boldsymbol{R}}^{\boldsymbol{2}})\right)^j E_t \pi_{t+1+j} = 0. \tag{28}$$

As anticipated, optimal policy has the form given in (24), and generates inflation responses as in (27). In contrast to the static case of imperfect information when lagged states are known, the impact responses of optimal policy to innovations in the two target components generally differ,

[46] The endogeneity of $\boldsymbol{R^2}$ arises for two reasons: First, as will be seen shortly, optimal policy differentiates its responses to different target innovations in the general case ($\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}} \neq \tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}$ as opposed to $\tilde{f}_\gamma = \tilde{f}_\varepsilon$) such that the policy coefficients in numerator and denominator of $\boldsymbol{R^2}$ do not cancel as they did in the case of R^2. Second, while the dynamics of γ_t, including σ_η^2, are exogenous, the innovation dynamics of $\tilde{\gamma}_t$, including $\text{Var}(\tilde{\gamma}_t)$, depend on the policy coefficients in (24). As shown in Appendix A.3, $\text{Var}(\tilde{\gamma}_t)$ solves a Riccati equation involving the endogenous policy coefficients $\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}$ and $\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}$. Reflecting the additional uncertainty arising from not knowing $E_{t-1}\gamma_t$, the Riccati equation implies $\text{Var}(\tilde{\gamma}_t) > \sigma_\eta^2$.

and the impact response to an innovation in the persistent target is lower than the impact response to a transitory shock, $\tilde{\boldsymbol{f}}_\gamma < \tilde{\boldsymbol{f}}_\varepsilon$. The response to prior beliefs is smaller than any of the innovation responses, $\hat{\boldsymbol{f}} < \tilde{\boldsymbol{f}}_\gamma < \tilde{\boldsymbol{f}}_\varepsilon$. Furthermore, the comparison between the response coefficients of output gap and inflation under perfect and imperfect information corresponds to what has been found in the static case: $\bar{f}_\varepsilon > \tilde{\boldsymbol{f}}_\varepsilon$, $\bar{f}_\gamma > \hat{\boldsymbol{f}}$, $\tilde{\boldsymbol{g}}_\varepsilon > \tilde{\boldsymbol{g}}_\gamma$, $\bar{g}_\gamma > \hat{\boldsymbol{g}} > \tilde{\boldsymbol{g}}_\gamma$. Depending on parameter values, $\tilde{\boldsymbol{f}}_\gamma$ and $\tilde{\boldsymbol{g}}_\varepsilon$ may be larger or smaller than $\bar{f}_\gamma$ and $\bar{g}_\varepsilon$, respectively.

Proof. See Appendix A.3 □

Despite some commonalities, there are profound differences in the policy prescriptions and resulting equilibria arising from the two cases of imperfect information.[47] In the general case, imperfect information does not only operate through a disciplinary channel by raising the Phillips curve slope — affecting directly the second term in each of the targeting rules (20) and (28) — but also through an explicitly forward-looking component — the third term in (28). As shown in Appendix A.3, the forward-looking term in (28) reflects the effect of current policy x_t on the continuation value of the discretionary policymaker's dynamic program arising from the endogeneity of public beliefs $\gamma_{t|t}$. Notably, this forward-looking term captures expectations of future inflation rates that are conditioned on full knowledge of the current target components.[48] The explicit consideration of inflation expectations in the targeting rule induces the policymaker to manage the stock of beliefs $\gamma_{t+1|t}$ that will be inherited by next period's discretionary policymaker.[49]

The policymaker cares about managing this belief stock, since the public's posterior beliefs $\gamma_{t|t}$ matter for future prior beliefs via the transition equation of the Kalman filter (25); this channel is absent in the cases of full information or imperfect information with known lagged states. In

[47] As shown by Blake and Kirsanova (2012), multiple equilibria can arise in linear-quadratic discretion problems when there are endogenous state variables as is here the case ($\gamma_{t|t-1}$) — and none is here assumed, except for the existence of a unique steady state. In numerical simulations reported in the technical appendix I have however never encountered multiple solutions within the class of time-invariant linear equilibria described here.

[48] Since the policymaker's information is strictly larger than the public's, prediction errors of the public are forecastable from the vantage point of the policymaker as will be illustrated below.

[49] However, in contrast to a commitment problem, the policymaker considers only the marginal effects of $\gamma_{t|t}$ on the problem's continuation value while taking the reaction of future policy on (expected) future output gap targets as given instead of being able to commit to a specific set of state-contingent choices that would most appropriately balance current and future losses.

the dynamic imperfect-information problem, shocks are propagated through the evolution of beliefs $\gamma_{t|t-1}$, and transitory target shocks have persistent effects on output gap and inflation that are entirely absent in the previous cases.

As in the static case of imperfect information, policy losses generated by optimal policy in the general case of imperfect information are lower than under full information. This result is established by numerical simulations, conducted over a comprehensive grid of admissible parameter values, which are described in the technical appendix.

3.5. Comparison of impulse responses

The differences in outcomes arising from optimal policy under full information and the two cases of imperfect information described above can be summarized by a comparison of the impulse responses generated in each case. For illustrative purpose, a calibration has been chosen that adopts the private-sector parameters used in Galì (2003)'s calibration of the New Keynesian model to quarterly U.S. data: $\beta = 0.99$, $\kappa = 0.17$ (corresponding to a Calvo-probability of not repricing equal to 0.75 and unit elasticities of intertemporal substitution and inverse Frisch labor supply). Furthermore, the policymaker is assumed place equal weight on minimizing variations in inflation and any shortfall between the output gap and the exogenous target $\bar{x}_t$. Shock variances are each normalized to unity and not intended to match the scale of any second moments. The persistence of γ_t, measured by ρ, has been set to 0.9.

[Figure 1 about here.]

As shown on Figure 1, full-information outcomes for output gap and inflation inherit the persistence properties of the underlying shocks. Reflecting the persistent nature of a shock to γ_t, the inflation response to the persistent target shock is considerably larger than to the transitory target shock. The introduction of imperfect information with known lagged states — also referred to as "static" case — changes mostly the impact responses of output gap and inflation. As an instance of the disciplinary channel discussed by Walsh (2000), output gap responses to either shock are lower under hidden information. As has been found in Proposition 2 above, this is a general feature across all admissible model calibrations, except for the impact response to γ_t, which can also exceed its full-information value; as demonstrated in a separate technical appendix, this occurs for low values of R^2 or α_x that lead the public to anticipate persistent output gap changes either with little probability or to be small in size, thus causing only little effect in inflation. While the inflation response to ε_t in Figure 1 is higher under hidden information, there is a much larger reduction in the inflation response to γ_t, reflecting the results of Proposition 3, that hidden information (in the static case) lowers the variance of inflation as well as expected policy losses compared to the full-information case.

Figure 1 also shows that the qualitative differences discussed so far between outcomes under perfect information and those obtained under hidden information with known lagged states, also carry over to the dynamic case of hidden information. Importantly however, when lagged states are known, transitory target shocks have no persistent effects on the economy and their propagation is identical to the full-information case. Things are different in the more general case of hidden information when transitory target shocks lead to persistent effects on beliefs, which in turn affect future outcomes. In particular, a positive shock to ε_t causes the output gap, and thus also $\gamma_{t+1|t}$ to rise, prompting a contractionary policy one period after the shock ($f_b < 0$) that persists for another couple of periods. Given that the simple New Keynesian model generally lacks endogenous persistence, the extended propagation of shocks due to this learning process mirrors the results of Erceg and Levin (2003), Milani (2007) and Collard et al. (2009). In addition, this pattern is similar (though not fully identical) to commitment policies under perfect information (Galì, 2008, Chapter 5). In both cases, an expectation that expansionary shocks will be undone in the future lowers current inflation. In the current setting, it is remarkable that this expectation is fully time-consistent and thus credible.

[Figure 2 about here.]

In the dynamic hidden-information case, the better-informed policymaker is able to anticipate future innovations to the public's limited information set. Public beliefs are rational however, and while the public's forecasts suffer from persistent prediction errors conditional on knowing the true shocks, these errors cancel out under the public information set. These predictable forecast errors are captured by differences in the impulse responses obtained from the different information sets available to public and policymaker, which are illustrated in Figure 2.

For a given variable, like the output gap x_t, the policymaker's impulse responses constructed at time t reflect a change in the policymaker's expectations due to a shock occurring at time t and are given by the sequence $E_t x_{t+k} - E_{t-1} x_{t+k}\ \forall k \geq 0$, For brevity the policymaker's impulse responses will also be referred to as true impulse response since they reflect knowledge of the true shocks. Likewise, the public's impulse responses at t are given by $x_{t+k|t} - x_{t+k|t-1}$ and reflect the update in beliefs caused by an innovation in the public information set at t ($\tilde{x}_t$). Figure 2 considers two different scenarios: Impulse responses of the output gap generated after a shock to the persistent target component (γ_t) are shown in the left-hand panels and those generated by a transitory shock are displayed in panels on the right of the figure.[50] Under both scenarios, a single shock is assumed to occur at time $t = 0$ and has been scaled to generate a unit innovation to the public information set; no further shocks are assumed to occur in either case.[51] Since both scenarios generate an identical innovation to the public's information set at $t = 0$, the public's impulse responses constructed at $t = 0$ are identical; they correspond to the probability-weighted average of the true impulse responses. The two bottom panels of the figure then compare the impulse responses constructed at $t = 0$ with public impulse responses that have been updated at $t = 1$ and $t = 2$ to reflect the actual evolution of shocks under each scenario.

As seen in the left-hand panels of Figure 2, the public persistently underpredicts the true trajectory of the output gap in the first scenario that is generated by a persistent target shock with an offsetting pattern of overpredictions following a transitory target shock shown in the right-hand panels. The remaining differences between true and public impulse responses at time $t = 2$ (the set of solid lines in the two bottom panels) foreshadow the extent of prediction errors in following

[50] A qualitatively similar figure for inflation can be generated as well.

[51] Assuming prior beliefs equal to zero ($\gamma_{0|-1} = 0$), a unit innovation, $\tilde{x}_0 = \tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}} \tilde{\gamma}_0 + \tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}} \tilde{\varepsilon}_0 = 1$, corresponds to ($\gamma_0 = 1/\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}, \varepsilon_0 = 0$) and ($\gamma_0 = 0, \varepsilon_0 = 1/\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}$), respectively. in the two scenarios.

periods; as discussed before, it takes about four periods for the effects of hidden information to dissipate after a shock. When particular periods are supposed to have been dominated by one set of shocks rather than another, such patterns of persistent forecast errors in public beliefs should be reflected in survey data. For example, Erceg and Levin (2003) use survey data to characterize the Volcker disinflation as a period of persistently excessive inflation forecasts, matching predictions from a linear-quadratic model with a Gaussian information structure similar to mine, but for a fixed policy rule. The methods presented here can be used to derive the parameters of such a rule within an explicitly optimizing framework of monetary policy under hidden information.

4. INFORMATION SHOCKS

The simple New Keynesian model analyzed so far has only one communication channel between policymaker and public: Policy actions themselves. In reality, there are however other communication channels than the policy instrument itself. If these channels are informative, they will alleviate the public's inference problems and affect the scope of belief management for policy. This section extends the information structure of the simple model to a multivariate setting, combining exogenous and endogenous signals. The particular example studied in this section nests the cases of perfect and (dynamic) imperfect information considered in Section 3.

In addition to observing policy, the public is now assumed to receive a noisy signal about the current value of the persistent component of the output target. These are the only two observables and the public cannot directly observe lagged states, extending the general hidden-information case described above. The target signal is contaminated by noise shocks n_t, which will be called "information shocks". and the public's measurement vector is

$$Z_t = \begin{bmatrix} x_t \\ s_t \end{bmatrix} \qquad \text{where} \quad s_t = \gamma_t + n_t \qquad \text{and} \quad n_t \sim N(0, \sigma_n^2)$$

The state vector of the model thus needs to be augmented by n_t.[52] The model is straightforward

[52]Since the noise in s_t is assumed to be serially uncorrelated it follows that $n_{t|t-1} = 0$ and there is no need to track

to solve numerically within the framework presented in Section 2. Recycling previously used notation, optimal policy has then the form

$$x_t = f_\gamma \gamma_t + f_\varepsilon \varepsilon_t + f_n n_t + f_b \gamma_{t|t-1}.$$

The information shocks are uncorrelated with fundamentals (γ_t and ε_t) and play no role under symmetric information. But under asymmetric information they matter since they are correlated with an informative signal about fundamentals, giving rise to fluctuations driven by "non fundamental" shocks.[53] Inflation is affected by information shocks via the forward-looking Phillips curve (7), making it suboptimal for policy to ignore information shocks. To the extent that people put positive weight on the exogenous signal s_t when forming beliefs about future inflation, a positive information shock will raise inflation and optimal policy should want to counteract the resulting effects on inflation by contracting output.

In the limit, as the volatility of noise shocks goes to either to zero (infinity), the extended model replicates the cases of full (hidden) information analyzed in the previous section. For $\sigma_n = 0$, the model is identical to the full information model, since γ_t is perfectly observable, and the transitory target component will be perfectly revealed by any linear policy with non-zero weight on ε_t. When σ_n is very large, the signal s_t becomes useless and the model converges to the imperfect information setting from the previous section where policy is the only observable. The exogenous signal s_t thus puts a ceiling on the public's uncertainty about realizations of the persistent target component, which will be illustrated shortly.

[Figure 3 about here.]

Impulse responses to a noise shock are shown in Figures 3, again using the baseline calibration described above, and setting the volatility of noise shocks equal to one. Under this configuration, each of the three shocks in this model occurs with the same probability. When the target signal

these constant beliefs in state vector of the policy problem.

[53]The meaning of "fundamentals" is intended here in the sense of the full-information economy.

$\gamma_t + n_t$ goes up because of a noise shock, this leads to ample confusion for the public. Current and expected inflation rise, since the public attributes part of the signal to the persistent target γ_t, To counteract these erroneous beliefs, policy contracts output. This is sensible in two ways: First it directly lowers inflation via the output term in (7). Second, it signals that the target γ_t may in fact not have gone up and thus reduces expected inflation. In the baseline calibration, it takes about four periods (one year) to fight these erroneous beliefs.

The resulting pattern of contracting output and elevated inflation is similar to the dynamics known from cost-push shocks, a pattern that Angeletos and La'O (2009) also found in response to information shocks in a model with dispersed beliefs in the private sector. Figure 3 documents that public beliefs of future output and inflation are *both* elevated during the entire episode, which distinguishes information-shock-induced dynamics from cost-push behavior, since the latter would typically be accompanied by an *expected* recession as well. Direct measures of people's beliefs, like surveys, might thus be useful to distinguish between these stories,

[Figure 4 about here.]

How does policy change with the volatility of information shocks? To answer this question, Figure 4 shows how policy coefficients and expected losses change when σ_n is varied between zero and infinity. The limit points in this experiment are the symmetric information model, respectively the previously studied model with no target signal except for policy. The policy coefficients f_γ, f_ε and f_b vary smoothly and monotonically between the comparative statics of hidden vs full information studied before. As the extent of hidden information increases with σ_n, optimal policy accepts larger shortfalls of the output gap from target — f_γ and f_ε become smaller — while reaping the benefits of reduced variability in inflation and, on net, lower policy losses. The policy response to noise shocks (f_n) is always negative. The reasons are similar to what has been seen before in terms of the negative marginal response to *prior* beliefs for a given level of γ_t. A contraction lowers inflation directly via the Phillips curve and indirectly via beliefs. For better comparison with the other coefficients, the middle left panel of Figure 4 the policy reaction to noise shocks has been normalized to measure the response to a one-standard deviation shock in n_t ($f_n \cdot \sigma_n$). In absolute value, the normalized noise response peaks at an intermediary level of σ_n, where the public places roughly equal weight (not shown) on policy and the target signal in its updating of beliefs. At such intermediary levels of σ_n policy react strongest to noise shocks, since the signal s_t is sufficiently informative to affect public beliefs, while also leaving ample uncertainty about γ_t for policy signal to matter as well.

The interplay of exogenous and endogenous information is further illustrated in the lower-left panel of Figure 4, which compares the information content about γ_t provided by s^t alone and (s^t, x^t) jointly; both measured by the share of variations in γ_t that is explained by projections of γ_t on either information set. By construction, the information content provided by s^t varies between zero and one as σ_n is varied from zero to infinity. Adding observed policy to the information set increases, of course, the precision of the public's information about γ_t. Strikingly however, optimal policy maintains a fairly constant share $\mathrm{Var}\left(\gamma_{t|t}\right)/\,\mathrm{Var}\left(\gamma_t\right)$ over almost all values of σ_n, except for very low values of σ_n where almost the entire variability of γ_t is explained by s_t alone. In part, this result reflect the fact that even when the public could only observe x^t optimal policy would convey enough information to account for more than eight-tenth of the variability in γ_t. With decreasing σ_n, optimal policy then lets f_ε increase more swiftly to its full-information value than f_γ, effectively rendering the policy signal about γ_t more noisier, and thus offsetting, at least in part, the higher precision of s_t.

Variations in noise variance affect the (ex-ante) transparency about the central bank's output target. The above results then document clear disadvantages from transparency. In a somewhat related model, Faust and Svensson (2001, Proposition 6.3) appear to establish the opposite: Namely that central bank losses were increasing, not decreasing, in transparency. The difference lies here in the definition of "transparency", and it is instructive to see how apparently innocuous differences in a model's setting can lead to different conclusions. In the experiments of Faust and Svensson, transparency means that targets can be perfectly inferred *once policy is observed*. In the experiments above, transparency ($\sigma_n = 0$) makes the target component γ_t directly observable, *regardless of policy*. Both imply the same information sets in equilibrium. But the constraints faced by the discretionary policymaker differ in profound ways. Under discretion, the policymaker takes the public beliefs system and its Kalman gains as given. When the exogenous signal s_t perfectly reveals the persistent target component — as is the case here when $\sigma_n = 0$ — public beliefs about the target are independent from the policy decision, and the optimal policy reduces to the full information case.[54] In contrast, when targets are perfectly inferable from observed policies as in Faust and Svensson (2001, 2002), the link from policy choices to beliefs is retained, which explains the difference in outcomes.

5. CONCLUSIONS

This paper has derived the optimal discretion policy for dynamic, linear-quadratic economies, when the policymaker is better informed than the public. In this case public beliefs about the policymaker's hidden information become a distinct, endogenous state variable of the policy problem. These public beliefs are shaped by observed policy actions, giving a scope for managing beliefs about future policies that is otherwise absent in a discretionary policy problem.

For specific models, variations of this result have already been reported for example by Faust and Svensson (2001, 2002), Gaspar et al. (2006, 2010) and Jensen (2006).[55] Using a similar example, these results are confirmed by my framework in a small New Keynesian economy with stochastic output targets. In this case, belief management disciplines the policymaker's pursuit of his output targets, which unambiguously lowers policy losses. For a version of the model with a simplified information structure, this result is derived analytically here, extending results from the earlier literature. Against a persuasive (and well-founded) trend towards increased openness of central banks, starting with Goodfriend (1986) and recently surveyed by Reis (2013), this is an interesting result. The underlying mechanism is second-best theory: The New Keynesian economy

[54] More precisely, any policy that places non-zero weight on ε_t will then perfectly reveal both target components, independently of the choice of policy coefficients. This setup is similar to what Faust and Svensson, p. 374 call the regime "OG: observable goal and intention", for which they find results similar to what is described here.

[55] Walsh (2000) and Tang (2014) also study the effects of hidden information on optimal discretionary policy, but within models that preclude the emergence of public beliefs as an endogenous state variable with effects for the propagation of shocks (instead of their amplification) .

cannot attain first best — which corresponds to zero policy losses in this environment — when subjected to output gap target shocks; furthermore only discretion policies are considered in this paper. As a general principle, adding a further friction (hidden information) may lead to better outcomes, when the economy is already in a second-best equilibrium. Of course, adding information frictions may but need not lower policy losses.[56]

Imperfect information gives also rise to information shocks as a source of business cycle fluctuations. Similar to Lorenzoni (2009), such shocks shift public perceptions about fundamentals, whilst their actual values remain unchanged. The optimal policy seeks to quell the erroneous beliefs arising from these shocks. In the New Keynesian model studied here, the dynamics induced by information shocks resemble those known from cost push shocks, similar to what Angeletos and La'O (2009) found based on higher-order dynamics in a setting of dispersed information.

APPENDIX

A. THE POLICYMAKER'S PROBLEM IN THE SIMPLE MODEL

A.1. Full information

Under discretion, the policymaker seeks to minimize the discounted sum of current and expected future losses (8), subject to the Phillips curve (7) by choosing a time t policy for the output gap without committing to future policy choices. Price setters are forward looking and form expectations about the reaction of future policies to future states, which in turn implies a mapping from expected future states to expected future inflation. Anticipating a linear equilibrium, this takes the form $\pi_{t+1|t} = \hat{g}\gamma_{t+1|t}$ (recall $\varepsilon_{t+1|t} = 0$).[57] The policymaker takes $\hat{g}$ as given; in equilibrium this coefficient must, of course, be consistent with optimal policy choices and private sector behavior. In the full-information case, the only Markov-perfect state variables are the exogenous compo-

[56]For example, as illustrated in the technical appendix, a lack of transparency about permanent changes in policymakers' inflation target can cause losses to be larger than in the case of a fully transparent inflation target.

[57]Notice that under full information $\pi_{t+1|t} = E(\pi_{t+1}|w^t) = E_t\pi_{t+1}$. For ease of comparison with the imperfect information cases, the notation $\pi_{t+1|t}$ is kept to explicitly distinguish terms involving public beliefs.

nents of the output gap target, γ_t and ε_t, and public beliefs $\gamma_{t+1|t} = \rho\gamma_t$ are exogenous as well. The policy problem can thus be written as

$$V(\gamma_t, \varepsilon_t) = \min_{x_t, \pi_t} \alpha_x (x_t - \bar{x}_t)^2 + \alpha_\pi \pi_t^2 + \beta E_t V(\gamma_{t+1}, \varepsilon_{t+1}) \qquad \text{s.t.} \quad \pi_t = \beta \hat{g} \rho \gamma_t + \kappa x_t$$

As the states are entirely exogenous, the expected continuation value $E_t V(\gamma_{t+1}, \varepsilon_{t+1})$ is independent of current policy and the first-order condition is the myopic targeting rule (11), resulting in the response coefficients given in (12) and (13).

A.2. Incomplete information when lagged states are known

As in the full information case, the policymaker takes as given the public's mapping of anticipated future states into future inflation $\pi_{t+1|t} = \hat{g}\gamma_{t+1|t}$ for some $\hat{g}$. In equilibrium, $\hat{g}$ must be consistent with an optimal policy of the form (17) and the Phillips curve; it follows that $\hat{g} = \hat{f} \cdot \kappa/(1-\beta\rho)$. Furthermore it will be useful to define $\hat{R}^2 \equiv K\hat{f}$. Assumption 1 implies $\hat{R}^2 > 0$; furthermore, $\hat{R}^2 < 1$ in equilibrium, as will be seen below.

Proof of Proposition 1. The optimal policy problem is

$$V(\gamma_t, \varepsilon_t, E_{t-1}\gamma_t) = \min_{x_t, \pi_t} \alpha_x (x_t - \bar{x}_t)^2 + \alpha_\pi \pi_t^2 + \beta E_t V(\gamma_{t+1}, \varepsilon_{t+1}, E_t\gamma_{t+1}) \tag{29}$$

$$\text{s.t.} \quad \pi_t = \kappa \frac{\beta\rho}{1-\beta\rho}(1-\hat{R}^2)E_{t-1}\gamma_t + \underbrace{\kappa\left(1 + \frac{\beta\rho}{1-\beta\rho}\hat{R}^2\right)}_{\equiv\tilde{\kappa}} x_t. \tag{30}$$

The Phillips curve constraint (30) follows from the signal extraction formula (15) and the expressions above for $\hat{g}$ and $\hat{R}^2$; $\tilde{\kappa} > \kappa$ is a consequence of $\hat{R}^2 > 0$. Taking first-order conditions yields the myopic targeting rule (20). Combining the targeting rule with the Phillips curve (30) results in the policy coefficients shown in (21). Since $\tilde{f}_\gamma = \tilde{f}_\varepsilon = \tilde{f}$ the gain (18) simplifies to $K = R^2/\tilde{f}$.

Inspection of (21) shows that a given ratio $\phi \equiv \hat{f}/\tilde{f}$, uniquely determines the individual policy coefficients $\tilde{f}$ and $\hat{f}$. The ratio ϕ itself solves a second-order polynomial:

$$\phi = \frac{\hat{f}}{\tilde{f}} = \frac{(1-\beta\rho)^2\alpha_x + (1-\beta\rho(1-R^2\phi))^2\alpha_\pi\kappa^2)}{(1-\beta\rho)^2\alpha_x + (1-\beta\rho(1-R^2\phi))\alpha_\pi\kappa^2)} \qquad \Leftrightarrow \qquad q\phi^2 = c - l\phi \tag{31}$$

for some $q, c, l > 0$.[58] The roots of (31) are given by the intersection of an upwardly open parabola and a downward sloping straight line. The parabola has its minimum at the origin of the coordinate system whereas the line has a positive intercept; there are thus two real roots, one positive and one negative. Since $q > c - l$ the positive root must lie between zero and one. (There is no such constraint on the negative root.) A graphical illustration is provided in the technical appendix.

The negative root would imply $\hat{f} < 0$ which is excluded by Assumption 1. Furthermore, the per-period loss of such a solution would be dominated by an alternative policy that chooses the same $\tilde{f}$ but uses $-\hat{f}$ instead of $\hat{f}$. This policy would generate the same inflation variance — notice that the gain does not depend on $\hat{f}$ — while yielding smaller shortfalls from the output gap target since $(-\hat{f}-1)^2 < (\hat{f}-1)^2$ for $\hat{f} < 0$. Consequently, the policymaker would prefer to deviate from the policy prescribed by a negative ϕ. □

Proof of Proposition 2. Substitution of the assumed policy (17) into the Phillips curve (19) yields the expression for inflation stated in (23) of the proposition with $\tilde{g} = \tilde{\kappa}\tilde{f}$ and the above expression for $\hat{g}$. Since $\tilde{\kappa} > \kappa$ it follows directly that $\hat{f} < \bar{f}_\gamma$ and thus $\hat{g} < \bar{g}_\gamma$. The technical appendix provides numerical examples of different parametrizations where $\tilde{f}_\gamma$ and $\tilde{g}_\varepsilon$ are either larger or smaller than $\bar{f}_\gamma$ and $\bar{g}_\varepsilon$, respectively. While $\tilde{f}$ can be larger or smaller than $\bar{f}_\gamma$ (depending on parameter values), it is straightforward to show that $(1-\beta\rho(1-\hat{R}^2))\tilde{f} < \bar{f}_\gamma$ and thus

$$\tilde{g} = \frac{\kappa}{1-\beta\rho}(1-\beta\rho(1-\hat{R}^2))\tilde{f} < \frac{\kappa}{1-\beta\rho}\bar{f}_\gamma = \bar{g}_\gamma.$$

Alternatively, note that the targeting rule (20) implies $\alpha_x f_b + \alpha_\pi\tilde{\kappa}g_b = 0$, such that f_b and g_b must have opposite signs; $g_b > 0$ then follows directly from $f_b < 0$. □

[58] Specifically, $q = \alpha_\pi\kappa^2\beta\rho R^2(1-\beta\rho R^2)$, $l = (1-\beta\rho)^2\alpha_x + \alpha_\pi\kappa^2(1-\beta\rho)(1-2\beta\rho R^2)$, and $c = (1-\beta\rho)^2\alpha_x + \alpha_\pi\kappa^2(1-\beta\rho)$.

Proposition A.1. *As a function of R^2 the limiting behavior of various coefficients is as follows:*

$$\lim_{R^2\to 0}\phi = \frac{(1-\beta\rho)^2\alpha_x + (1-\beta\rho)^2\alpha_\pi\kappa^2}{(1-\beta\rho)^2\alpha_x + (1-\beta\rho)\alpha_\pi\kappa^2}, \qquad \lim_{R^2\to 0}\hat{R}^2 = 0, \qquad \lim_{R^2\to 0}\tilde{f} = \bar{f}_\varepsilon. \tag{32}$$

Thus, as $R^2 \to 0$, the economy approaches the full-information equilibrium of an economy absent fluctuations in γ_t and $E_{t-1}\gamma_t$; the limiting behavior of $\hat{f}$ becomes uninteresting in this case. Furthermore,

$$\lim_{R^2\to 1}\phi = 1, \quad \lim_{R^2\to 1}\hat{R}^2 = 1, \quad \lim_{R^2\to 1}\tilde{f} = \lim_{R^2\to 1}\hat{f} = \frac{(1-\beta\rho)^2\alpha_x}{(1-\beta\rho)^2\alpha_x + \alpha_\pi\kappa^2} \equiv f^*_\gamma < \bar{f}_\gamma, \tag{33}$$

and f^*_γ is the optimal policy coefficient amongst the class of policy rules $x_t = f\cdot\gamma_t$ that minimizes expected losses $E(L_t|\gamma_t)$.

Proof. The limiting behavior is straightforward to verify given the policy coefficients derived in Proposition 1. To see the optimality of f^*_γ, notice that for a policy of the form $x_t = f\cdot\gamma_t$ expected losses — conditional on γ_t being the only source of target shocks — are given by

$$E(L_t|\gamma_t) = \left(\alpha_x(f-1)^2 + \alpha_\pi\left(\frac{\kappa}{(1-\beta\rho)}\cdot f\right)^2\right)\cdot \mathrm{Var}\left(\gamma_t\right)$$

such that $f^*_\gamma = \mathrm{argmin}_f\, E(L_t|\gamma_t)$. □

Proof of Proposition 3. Expected losses are defined as $E(L_t) = \alpha_x \mathrm{Var}\left(x_t - \bar{x}_t\right) + \alpha_\pi \mathrm{Var}\left(\pi_t\right)$. In order to establish that $E(L_t^{FI}) > E(L_t^{HI})$ it will be useful to decompose expected losses into $E(L_t) = E(\tilde{L}_t) + E(\hat{L}_t)$ with $E(\hat{L}_t) = \alpha_x \mathrm{Var}\left(E_{t-1}(x_t - \bar{x}_t)\right) + \alpha_\pi \mathrm{Var}\left(E_{t-1}(\pi_t)\right)$ and $E(\tilde{L}_t) = \alpha_x \mathrm{Var}_{t-1}\left(x_t - \bar{x}_t\right) + \alpha_\pi \mathrm{Var}_{t-1}\left(\pi_t\right)$.[59] The same decomposition can be applied to losses under full and hidden information, denoted by superscripts FI and HI, respectively.

[59] This decomposition is an application of the law of total variance when conditional variances are constant.

Since the hidden-information coefficient $\hat{f}$ lies between its full-information counterpart $\bar{f}_\gamma$ and the optimal rule coefficient f^*_γ, defined in Proposition A.1, it is straightforward to show that $E(\hat{L}^{FI}_t) > E(\hat{L}^{HI}_t)$.[60] Given $E(\hat{L}^{FI}_t) > E(\hat{L}^{HI}_t)$ it is then sufficient to show that $E(\tilde{L}^{FI}_t) > E(\tilde{L}^{HI}_t)$ in order to establish $E(L^{FI}_t) > E(L^{HI}_t)$.

The targeting rule (20) implies that $E(L^{HI}_t) = (\alpha_x/\alpha_\pi) \cdot (\alpha_x + \alpha_\pi \tilde{\kappa}^2) \cdot \mathrm{Var}\,(\pi^{HI}_t)$ and except for differences in parameter value (κ vs. $\tilde{\kappa}$), the same holds in the full-information case. For $E(\tilde{L}^{FI}_t) > E(\tilde{L}^{HI}_t)$ to hold, it is then required that

$$\frac{\mathrm{Var}_{t-1}\,(\pi^{FI}_t)}{\bar{f}_\varepsilon} > \frac{\mathrm{Var}_{t-1}\,(\pi^{HI}_t)}{\tilde{f}}. \tag{34}$$

Without loss of generality, the sum of shock variances can be normalized to $\sigma^2_\eta + \sigma^2_\varepsilon = 1$; conditional variances of inflation are then given by $\mathrm{Var}_{t-1}\,(\pi^{FI}_t) = R^2 \bar{g}^2_\varepsilon + (1 - R^2)\bar{g}^2_\gamma$ and $\mathrm{Var}_{t-1}\,(\pi^{HI}_t) = \tilde{g}^2$.[61]

A necessary condition for (34) is $\mathrm{Var}_{t-1}\,(\pi^{FI}_t) > \mathrm{Var}_{t-1}\,(\pi^{HI}_t)$ which will be verified first. To establish this, it turns out to be important to properly account for the endogenous wedge between R^2 and $\hat{R}^2 = R^2 \cdot \hat{f}/\tilde{f}$. As stated in Proposition 1, this wedge is determined by a second-order polynomial, which risks obtaining highly unwieldy expressions when substituting out $\hat{R}^2$ for the derivation of the proof here. Instead, it is more convenient to substitute out R^2 against $\hat{R}^2 \cdot \tilde{f}/\hat{f}$; this is admissible since for given parameters α_x, κ, β, ρ there is a one-to-one mapping from R^2 to $\hat{R}^2$, see Propositions 1 and A.1. The range of relevant parameters can thus also be described by varying $\hat{R}^2$ between zero and one and substituting out R^2.

To show that $\mathrm{Var}_{t-1}\,(\pi^{FI}_t) > \mathrm{Var}_{t-1}\,(\pi^{HI}_t)$ notice that, because of Jensen's inequality, it is sufficient to compare the average inflation response under full information, denoted $\bar{g}$, to the uniform

[60] For an arbitrary coefficient f expected losses due to $E_{t-1}\gamma_t$ — either under hidden or perfect information — are given by $E(\hat{L}_t) = \left(\alpha_x(f-1)^2 + \alpha_\pi\,(f \cdot \kappa/(1-\beta\rho))^2\right) \cdot \mathrm{Var}\,(E_{t-1}\gamma_t)$ such that $f^*_\gamma = \mathrm{argmin}_f\, E(L_t|\gamma_t)$ as in Proposition A.1. As a function of f, $E(\hat{L})$ can be represented as an upwardly open parabola whose minimum lies at f^*_γ. Losses generated by alternative choices for f can then be ranked by the distance between f and f^*_γ.

[61] For any variable v_t, conditional variances $\mathrm{Var}_t\,(v_t) = E\left((v_t - E_{t-1}v_t)^2\right)$ are defined with respect to the information set spanned by w^t using $E_{t-1}v_t = E(v_t|w^{t-1})$ as elsewhere in this paper,

inflation response $\tilde{g}$ under hidden information:

$$\operatorname{Var}_{t-1}(\pi_t^{FI}) = R^2\bar{g}_\varepsilon^2 + (1-R^2)\bar{g}_\gamma^2 > \left(R^2\bar{g}_\varepsilon + (1-R^2)\bar{g}_\gamma\right)^2 \equiv \bar{g}^2$$

$$\bar{g} = \bar{f}_\varepsilon\left(1 + \frac{\beta\rho}{1-\beta\rho}\hat{R}^2\frac{\tilde{f}}{\hat{f}}\bar{f}_\gamma\right) > \bar{f}_\varepsilon\left(1 + \frac{\beta\rho}{1-\beta\rho}\hat{R}^2\tilde{f}\right) > \tilde{f}\left(1 + \frac{\beta\rho}{1-\beta\rho}\hat{R}^2\right) = \tilde{g}$$

As a result, the variance of inflation under hidden information is lower than in the full-information case.[62] To establish (34) it can then be verified that:[63]

$$\frac{\operatorname{Var}_{t-1}(\pi_t^{FI})}{\bar{f}_\varepsilon} > \frac{\bar{g}^2}{\bar{f}_\varepsilon} = \bar{f}_\varepsilon\left(1 + \frac{\beta\rho}{1-\beta\rho}\hat{R}^2\frac{\tilde{f}}{\hat{f}}\bar{f}_\gamma\right)^2 > \tilde{f}\left(1 + \frac{\beta\rho}{1-\beta\rho}\hat{R}^2\right)^2 = \frac{\operatorname{Var}_{t-1}(\pi_t^{HI})}{\tilde{f}}.$$

□

A.3. Incomplete information (general case)

The evolution of public beliefs $\gamma_{t|t}$ is described by a Kalman filter that uses the equilibrium policy in (24) as measurement equation and the exogenous AR(1) dynamics of γ_t, see (10), as state transition. The Kalman filter and its gain are given in (25) and (26). The variance of state innovations is described by a Riccati equation, that depends on the policy responses to target innovations, $\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}$ and $\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}$, but not on the response to prior beliefs, $\hat{\boldsymbol{f}}$; the variance ratio $\boldsymbol{R^2} = \operatorname{Var}(\tilde{\gamma}_{t|t}) / \operatorname{Var}(\tilde{\gamma}_t)$ is also described by a Riccati equation:

$$\operatorname{Var}(\tilde{\gamma}_t) = \sigma_\eta^2 + \rho^2\frac{\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}^2\sigma_\varepsilon^2}{\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}^2\operatorname{Var}(\tilde{\gamma}_t) + \tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}^2\sigma_\varepsilon^2}\operatorname{Var}(\tilde{\gamma}_t) = \frac{\sigma_\eta^2}{1-\rho^2(1-\boldsymbol{R^2})} > \sigma_\eta^2 \tag{35}$$

$$\boldsymbol{R^2} = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \left(\frac{\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}}{\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}}\right)^2\sigma_\varepsilon^2\cdot(1-\rho^2(1-\boldsymbol{R^2}))} > \frac{\sigma_\eta^2}{\sigma_\eta^2 + \left(\frac{\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}}{\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}}\right)^2\sigma_\varepsilon^2} < \frac{\sigma_\eta^2}{\sigma_\eta^2+\sigma_\varepsilon^2} = R^2 \tag{36}$$

[62] Because of $\hat{f} < \bar{f}_\gamma$, $\bar{g} > \tilde{g}$ does not only establish $\operatorname{Var}_{t-1}(\pi_t^{FI}) > \operatorname{Var}_{t-1}(\pi_t^{HI})$ but also that the unconditional variance of inflation is lower under hidden information.

[63] The verification involves substitution and further manipulation of the above-defined policy coefficients which is straightforward, though tedious.

As suggested by the conflicting inequalities in (36), $\boldsymbol{R^2}$ and R^2 — the exogenous counterpart to $\boldsymbol{R^2}$ from the static imperfect-information case — cannot generally be ranked. The first inequality is caused by the additional uncertainty arising from not observing $E_{t-1}\gamma_t$ whereas the second inequality reflects $\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}} > \tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}$ under the optimal policy (Proposition 4). As shown in the technical appendix, for different parameter values the latter may outweigh the former, causing $\boldsymbol{R^2}$ to be larger or smaller than R^2. (For the calculation of expected losses, the latter is the appropriate probability weight in the static case but the former needs to be used in the dynamic case, complicating the analysis of expected losses in the general case considerably.)

Optimal policy under discretion solves a dynamic program consisting of the objective function (8), a Phillips curve constraint analogous to (19), a transition for beliefs given by the Kalman filter (25) plus the exogenous laws of motion for the output gap target and its components given in (9) and (10).[64] The state variables of the problem are γ_t, ε_t and $\gamma_{t|t-1}$ (or alternatively $\tilde{\gamma}_t$, ε_t and $\gamma_{t|t-1}$) and the Lagrangian is

$$\begin{aligned} V(\gamma_t, \varepsilon_t, \gamma_{t|t-1}) = \max_{x_t, \pi_t, \gamma_{t+1|t}} \; & \frac{1}{2}\left(\alpha_x (x_t - \bar{x}_t)^2 + \alpha_\pi \pi_t^2\right) + \beta E_t V(\gamma_{t+1}, \varepsilon_{t+1}, \gamma_{t+1|t}) \\ & + \lambda_t \left(\pi_t - \beta \hat{\boldsymbol{g}}_{\boldsymbol{\gamma}} \gamma_{t+1|t} - \kappa x_t\right) + \xi_t \left(\gamma_{t+1|t} - \rho(1 - \hat{\boldsymbol{R}}^{\boldsymbol{2}})\gamma_{t|t-1} - \rho \boldsymbol{K} x_t\right). \end{aligned} \tag{37}$$

The targeting rule (28) then follows from the first-order conditions and the envelope theorem:

$$\alpha_x (x_t - \bar{x}_t) - \kappa \lambda_t - \rho \boldsymbol{K} \xi_t = 0, \qquad \alpha_\pi \pi_t + \lambda_t = 0, \tag{38}$$

$$\beta \frac{\partial V_{t+1}}{\partial \gamma_{t+1|t}} - \beta \hat{\boldsymbol{g}}_{\boldsymbol{\gamma}} \lambda_t + \xi_t = 0, \qquad \frac{\partial V_t}{\partial \gamma_{t|t-1}} = -\rho(1 - \hat{\boldsymbol{R}}^{\boldsymbol{2}})\xi_t. \tag{39}$$

The forward-looking component of the targeting rule is obtained from combining the two expressions in (39) and iterating forward, which translates the marginal effects of beliefs on the continuation value ($\partial V_{t+1}/\partial \gamma_{t+1|t}$) into a discounted sum of the policymaker's expectations of future inflation rate.

[64] Similar to the static case of hidden information, the Kalman gain $\boldsymbol{K}$ and anticipated policy $x_{t|t-1} = \hat{\boldsymbol{f}} \gamma_{t|t-1}$ are taken as given in the optimization.

To derive the optimal policy coefficients it is necessary to evaluate the policymaker's expectations $E_t\pi_{t+1+j}$ in (28). Conditional on the equilibrium policy in (24) inflation evolves as in (27) and beliefs follow $\gamma_{t+1|t} = \rho(1-\boldsymbol{R^2})\gamma_{t|t-1} + \rho\boldsymbol{R^2}\gamma_t + \rho\boldsymbol{K}\cdot\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}\varepsilon_t$. Let $\boldsymbol{\Theta} \equiv \beta\rho(1-\hat{\boldsymbol{R}}^{\boldsymbol{2}})$ and inflation expectations can be computed from

$$E_t\pi_{t+1+j} = \rho^j\left[\hat{\boldsymbol{g}}_{\boldsymbol{\gamma}} - \boldsymbol{g_b}(1-\boldsymbol{R^2})^j\right]\rho\gamma_t + \boldsymbol{g_b}\left(\rho(1-\boldsymbol{R^2})\right)^j\rho\gamma_{t|t} \tag{40}$$

$$\sum_{j=0}^{\infty}\boldsymbol{\Theta}^j E_t\pi_{t+1+j} = \left[\frac{\hat{\boldsymbol{g}}_{\boldsymbol{\gamma}}}{1-\boldsymbol{\Theta}\rho} - \frac{\boldsymbol{g_b}}{1-\boldsymbol{\Theta}\rho(1-\boldsymbol{R^2})}\right]\rho\gamma_t + \frac{\boldsymbol{g_b}}{1-\boldsymbol{\Theta}\rho(1-\boldsymbol{R^2})}\rho\gamma_{t|t} \tag{41}$$

The targeting rule (28) can then be written as

$$\alpha_x(x_t - \bar{x}_t) + \alpha_\pi\hat{\boldsymbol{\kappa}}\pi_t + \hat{\boldsymbol{\Psi}}\hat{\boldsymbol{f}}_{\boldsymbol{\gamma}}\gamma_t - \boldsymbol{\Psi_b}\left(\hat{\boldsymbol{f}}_{\boldsymbol{\gamma}} - (1-\boldsymbol{\Theta}\boldsymbol{f}_{\boldsymbol{\gamma}})\right)(\gamma_t - \gamma_{t|t}) = 0. \tag{42}$$

where $\hat{\boldsymbol{\Psi}} \equiv \alpha_\pi\dfrac{\kappa^2}{1-\beta\rho}\dfrac{\beta\rho}{1-\beta\rho}\dfrac{\boldsymbol{\Theta}\rho}{1-\boldsymbol{\Theta}\rho}\hat{\boldsymbol{R}}^{\boldsymbol{2}}$, and $\boldsymbol{\Psi_b} \equiv \alpha_\pi\dfrac{\kappa^2}{1-\beta\rho}\dfrac{\beta\rho}{1-\beta\rho}\dfrac{\boldsymbol{\Theta}\rho}{1-\boldsymbol{\Theta}\rho(1-\boldsymbol{R^2})}\hat{\boldsymbol{R}}^{\boldsymbol{2}}$,

Substituting the responses of output gap and inflation from (24) and (27) into (42) and matching coefficients, the optimal policy coefficients are then given by:

$$\begin{aligned}
\hat{\boldsymbol{f}}_{\boldsymbol{\gamma}} &= \frac{\alpha_x}{\alpha_x + \alpha_\pi\frac{\kappa^2}{1-\beta\rho}\frac{1-\boldsymbol{\Theta}}{1-\beta\rho} + \hat{\boldsymbol{\Psi}}}, \\
\hat{\boldsymbol{f}}_{\boldsymbol{\varepsilon}} &= \frac{\alpha_x}{\alpha_x + \alpha_\pi\hat{\boldsymbol{\kappa}}^2 + \boldsymbol{\Psi_b}\left(\hat{\boldsymbol{R}}^{\boldsymbol{2}} - (1-\boldsymbol{\Theta})\boldsymbol{R^2}\right)}, \\
\hat{\boldsymbol{f}}_{\boldsymbol{\gamma}} &= \frac{\alpha_x - \left(\hat{\boldsymbol{\Psi}} - \boldsymbol{\Psi_b}(1-\boldsymbol{R^2})\right)\hat{\boldsymbol{f}}_{\boldsymbol{\gamma}}}{\alpha_x + \alpha_\pi\hat{\boldsymbol{\kappa}}^2 + \boldsymbol{\Psi_b}(1-\boldsymbol{R^2})(1-\boldsymbol{\Theta})}.
\end{aligned}$$

Proof of Proposition 4. Given the policy coefficients derived above it is straightforward to verify that $\bar{f}_\varepsilon > \tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}} > \tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}} > \hat{\boldsymbol{f}}$ and $\hat{\boldsymbol{f}} < \bar{f}_\gamma \Leftrightarrow \hat{\boldsymbol{g}} < \bar{g}_\gamma$. Assumption 2 implies $\boldsymbol{f_b} = \hat{\boldsymbol{f}} - \tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}} < 0$ and the targeting rule (42) requires $\alpha_x\boldsymbol{f_b} + \alpha_\pi\tilde{\boldsymbol{\kappa}}\left(1 + \frac{\beta\rho}{1-\boldsymbol{\Theta}}\frac{\boldsymbol{\Theta}\rho}{1-\boldsymbol{\Theta}\rho(1-\boldsymbol{R^2})}\hat{\boldsymbol{R}}^{\boldsymbol{2}}\right)\boldsymbol{g_b} = 0$ from which follows $\boldsymbol{g_b} > 0$ and thus $\hat{\boldsymbol{g}} > \tilde{\boldsymbol{g}}_{\boldsymbol{\gamma}} \Rightarrow \bar{g}_\gamma > \tilde{\boldsymbol{g}}_{\boldsymbol{\gamma}}$. The technical appendix provides numerical examples of different parameterizations where $\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}$ and $\tilde{\boldsymbol{g}}_{\boldsymbol{\varepsilon}}$ are either larger or smaller than $\bar{f}_\gamma$ and $\bar{g}_\varepsilon$, respectively. $\square$

REFERENCES

Abrams, B. A. (2006, Fall). How Richard Nixon pressured Arthur Burns: Evidence from the Nixon tapes. *Journal of Economic Perspectives 20*(4), 177–188.

Angeletos, G.-M., H. Dellas, and F. Collard (2013). Beliefs and fluctuations: A quantitative exploration. *mimeo* Massachusetts Institute of Technology.

Angeletos, G.-M. and J. La'O (2009). Incomplete information, higher-order beliefs and price inertia. *Journal of Monetary Economics 56*(Supplement 1), S19 – S37. Supplement issue: December 12-13, 2008 Research Conference on 'Monetary Policy under Imperfect Information' Sponsored by the Swiss National Bank (www.snb.ch) and Study Center Gerzensee (www.szgerzensee.ch).

Aoki, K. (2003, April). On the optimal monetary policy response to noisy indicators. *Journal of Monetary Economics 50*(3), 501–523.

Aoki, K. (2006, January). Optimal commitment policy under noisy information. *Journal of Economic Dynamics and Control 30*(1), 81–109.

Atkeson, A., V. V. Chari, and P. J. Kehoe (2010, February). Sophisticated Monetary Policies. *The Quarterly Journal of Economics 125*(1), 47–89.

Backus, D. and J. Driffill (1985a, June). Inflation and reputation. *The American Economic Review 75*(3), 530–538.

Backus, D. and J. Driffill (1985b, April). Rational expectations and policy credibility following a change in regime. *The Review of Economic Studies 52*(2), 211–221.

Ball, L. (1995, November). Time-consistent policy and persistent changes in inflation. *Journal of Monetary Economics 36*(2), 329–350.

Barro, R. J. and D. B. Gordon (1983, July). Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics 12*, 101–121.

Berkelmans, L. (2011). Imperfect information, multiple shocks, and policy's signaling role. *Journal of Monetary Economics 58*(4), 373–386.

Blake, A. P. and T. Kirsanova (2012). Discretionary policy and multiple equilibria in LQ RE models. *Review of Economic Studies 79*(4), 1309–1339.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics 12*(3), 383–398.

Canzoneri, M. B. (1985, December). Monetary policy games and the role of private information. *The American Economic Review 75*(5), 1056–1070.

Chari, V. V. and P. J. Kehoe (1990, August). Sustainable plans. *The Journal of Political Economy 98*(4), 783–802.

Clarida, R., J. Galì, and M. Gertler (1999, December). The science of monetary policy: A New Keynesian perspective. *Journal of Economic Literature 37*, 1661–1707.

Collard, F. and H. Dellas (2007, March-April). The great inflation of the 1970s. *Journal of Money, Credit and Banking 39*(2-3), 713–731.

Collard, F. and H. Dellas (2010, March-April). Monetary misperceptions, output, and inflation dynamics. *Journal of Money, Credit and Banking 42*(2-3), 483–502.

Collard, F., H. Dellas, and F. Smets (2009). Imperfect information and the business cycle. *Journal of Monetary Economics 56, Supplement*(0), S38 – S56. Supplement issue: December 12-13, 2008 Research Conference on 'Monetary Policy under Imperfect Information' Sponsored by the Swiss National Bank (www.snb.ch) and Study Center Gerzensee (www.szgerzensee.ch).

Cukierman, A. and N. Liviatan (1991, February). Optimal accommodation by strong policymakers under incomplete information. *Journal of Monetary Economics 27*(1), 99–127.

Cukierman, A. and A. H. Meltzer (1986, September). A theory of ambiguity, credibility, and inflation under discretion and asymmetric information. *Econometrica 54*(5), 1099–1128.

Damjanovic, T., V. Damjanovic, and C. Nolan (2008). Unconditionally optimal monetary policy. *Journal of Monetary Economics 55*(3), 491 – 500.

Erceg, C. J. and A. T. Levin (2003, May). Imperfect credibility and inflation persistence. *Journal of Monetary Economics 50*(4), 915–944.

Faust, J. and L. E. O. Svensson (2001, May). Transparency and credibility: Monetary policy with unobservable goals. *International Economic Review 42*(2), 369–97.

Faust, J. and L. E. O. Svensson (2002, May). The equilibrium degree of transparency and control in monetary policy. *Journal of Money, Credit and Banking 34*(2), 520–39.

Galì, J. (2003). New perspectives on monetary policy, inflation, and the business cycle. In M. Dewatripont, L. P. Hansen, and S. J. Turnovsky (Eds.), *Advances in Economics and Econometrics: Theory and Applications*, Eighth World Congress of the Econometric Society, Chapter 6, pp. 151–197. Cambridge, UK: Cambridge University Press.

Galì, J. (2008). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton, New Jersey: Princeton University Press.

Gaspar, V., F. Smets, and D. Vestin (2006, 04-05). Adaptive learning, persistence, and optimal monetary policy. *Journal of the European Economic Association 4*(2-3), 376–385.

Gaspar, V., F. Smets, and D. Vestin (2010). Chapter 19 - Inflation Expectations, Adaptive Learning and Optimal Monetary Policy. In B. M. Friedman and M. Woodford (Eds.), *Handbook of Monetary Economics*, Volume 3, pp. 1055 – 1095. Elsevier.

Geraats, P. M. (2007, March). The mystique of central bank speak. *International Journal of Central Banking 3*(1), 37–80.

Goodfriend, M. (1986, January). Monetary mystique: Secrecy and central banking. *Journal of Monetary Economics 17*(1), 63–92.

Jensen, C. (2006, December). Expectations, learning, and discretionary policymaking. *International Journal of Central Banking 2*(4), 137 – 156.

Jensen, H. (2002, September). Optimal degrees of transparency in monetary policymaking. *Scandinavian Journal of Economics 104*(3), 399–422.

King, R. G., Y. K. Lu, and E. S. Pasten (2008). Managing expectations. *Journal of Money, Credit and Banking 40*(8), 1625–1666.

King, R. G. and A. L. Wolman (1996, May/June). Inflation targeting in a St. Louis model of the 21st century. *Review 78*(3), 83–107.

Kohlhas, A. N. (2013, December). Learning-by-sharing: Monetary policy and hte nformation content of aggregate variables. *mimeo* University of Cambridge.

Kydland, F. E. and E. C. Prescott (1977, June). Rules rather than discretion: The inconsistency of optimal plans. *The Journal of Political Economy 85*(3), 473–492.

Lorenzoni, G. (2009, December). A theory of demand shocks. *The American Economic Review 99*(5), 2050–2084.

Lorenzoni, G. (2010, January). Optimal monetary policy with uncertain fundamentals and dispersed information-super-. *Review of Economic Studies 77*(1), 305–338.

Milani, F. (2007, October). Expectations, learning and macroeconomic persistence. *Journal of Monetary Economics 54*(7), 2065–2082.

Molnár, K. and S. Santoro (2010, May). Optimal monetary policy when agents are learning. Working Paper 2010/08, Norges Bank.

Morris, S. and H. S. Shin (2002, December). Social value of public information. *American Economic Review 92*(5), 1521–1534.

Morris, S. and H. S. Shin (2005). Central Bank Transparency and the Signal Value of Prices. *Brookings Papers on Economic Activity 36*(2), 1–66.

Nimark, K. (2008, March). Dynamic pricing and imperfect common knowledge. *Journal of Monetary Economics 55*(2), 365–382.

Orphanides, A. and J. C. Williams (2005, April). Inflation scares and forecast-based monetary policy. *Review of Economic Dynamics 8*(2), 498–527.

Oudiz, G. and J. Sachs (1985). International policy coordination in dynamic macroeconomic models. In W. H. Buiter and R. C. Marston (Eds.), *International Economic Policy Coordination*, NBER Chapters, Chapter 7, pp. 274–330. Cambridge University Press.

Persson, T. and G. Tabellini (2000). *Political Economics*. Zeuthen lecture book series. Cambridge, MA: The MIT Press.

Reis, R. (2013). Central bank design. *Journal of Economic Perspectives 27*(4), 17–44.

Romer, C. D. and D. H. Romer (2000, June). Federal Reserve information and the behavior of interest rates. *American Economic Review 90*(3), 429–457.

Sargent, T. J. (1999). *The Conquest of American Inflation*. Princeton (NJ): Princeton University Press.

Sibert, A. (2009). Is transparency about central bank plans desirable? *Journal of the European Economic Association 7*(4), 831–857.

Sims, C. A. (2002). The role of models and probabilities in the monetary policy process. *Brookings Papers on Economic Activity 2002*(2), 1–40.

Söderlind, P. (1999, April). Solution and estimation of RE macromodels with optimal policy. *European Economic Review 43*(4-6), 813–823.

Svensson, L. E. O. and M. Woodford (2003, April). Indicator variables for optimal policy. *Journal of Monetary Economics 50*(3), 691–720.

Svensson, L. E. O. and M. Woodford (2004, January). Indicator variables for optimal policy under asymmetric information. *Journal of Economic Dynamics and Control 28*(4), 661–690.

Tang, J. (2014, August). Uncertainty and the signaling channel of monetary policy. *mimeo* Federal Reserve Bank of Boston.

Walsh, C. E. (2000, April). Market discipline and monetary policy. *Oxford Economic Papers 52*(2), 249–271.

Walsh, C. E. (2003). *Monetary Theory and Policy* (2nd edition ed.). Cambridge, MA: The MIT Press.

Westelius, N. J. (2009, April). Imperfect transparency and shifts in the central bank's output gap target. *Journal of Economic Dynamics and Control 33*(4), 985–996.

Woodford, M. (2003). *Interest and Prices*. Princeton, NJ: Princeton University Press.

Yun, T. (1996, April). Nominal price rigidity, money supply endogeneity, and business cycles. *Journal of Monetary Economics 37*(2), 345–370.

Figure 1: Impulse Responses in the Simple Model

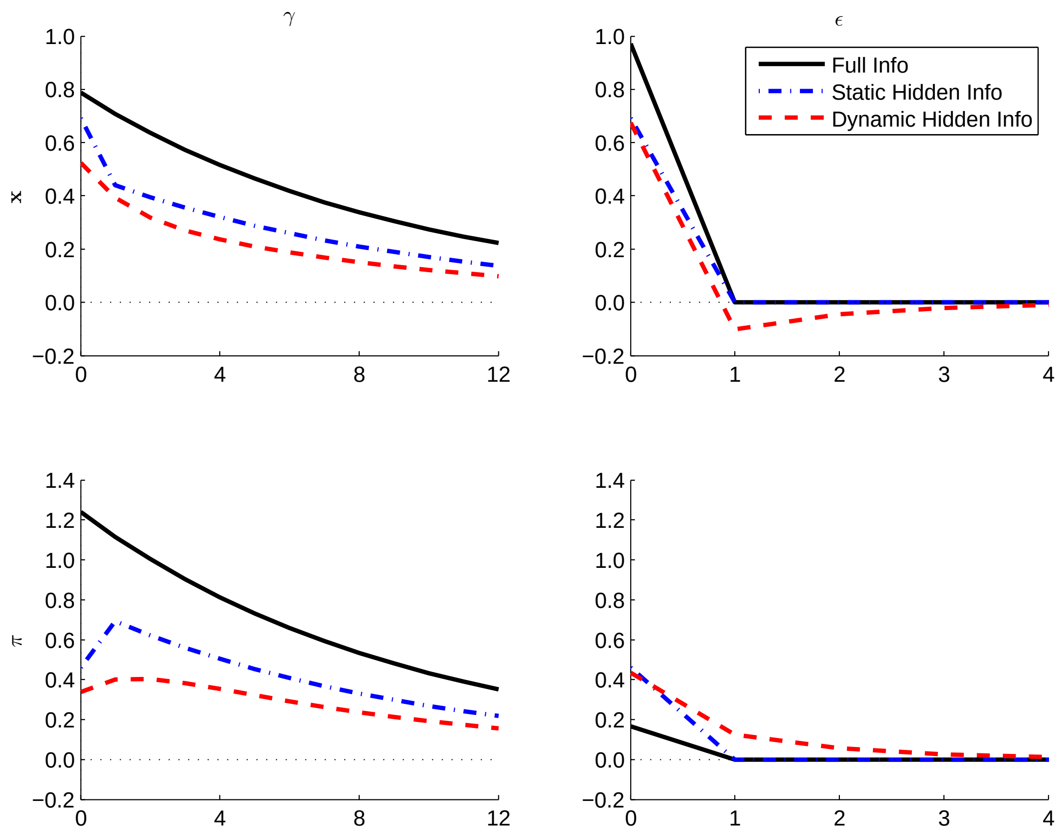

Note: Impulse responses of output gap (x_t) and inflation (π_t) under full information (straight black lines), hidden information with known lagged states (dash-dotted blue lines) and the general case of hidden information (dashed red lines) to unit standard deviation shocks to γ_t and ε_t.

Figure 2: True Impulse Responses and Public Beliefs about the Output Gap

Time $t = 0$

(a) True impulse is $\gamma_0 = 1/\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}$

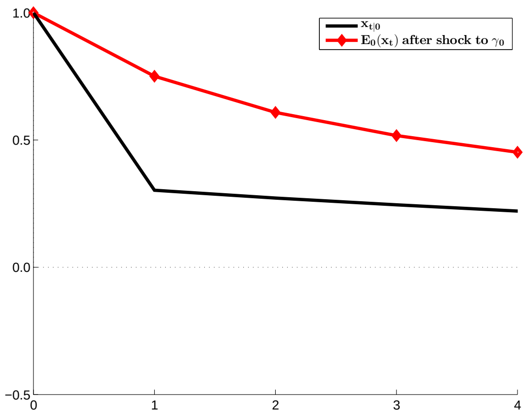

(b) True impulse is $\varepsilon_0 = 1/\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}$

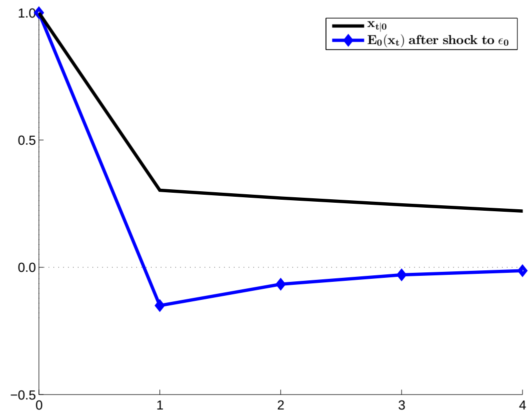

Time $t = 2$

(c) True impulse was $\gamma_0 = 1/\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}$

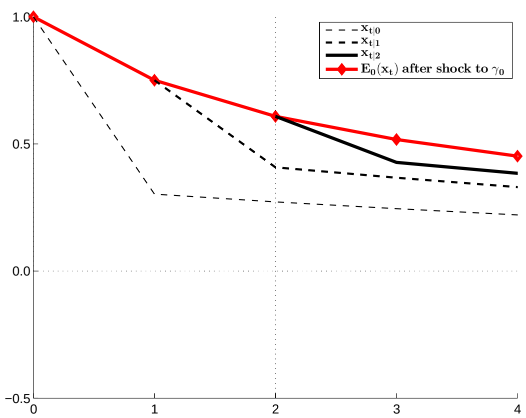

(d) True impulse was $\varepsilon_0 = 1/\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}}$

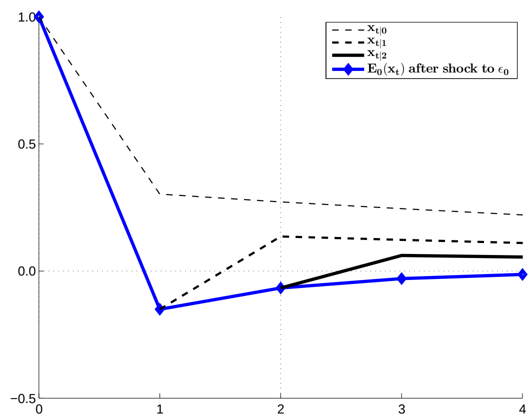

Note: All figures display impulse responses for the output gap in the simple New Keynesian model generated either by a positive shock to γ_t (left-hand column) or ε_t (right-hand column) at time $t = 0$. All outcomes are generated from the general case of hidden information when the public observes only x^t, and all impulses have been scaled to obtain a unit innovation to the public's information set. Each panel compares impulse responses derived from the policymaker's information set at time $t = 0$ that spans the true shocks, against impulse responses constructed at different points in time under the public's limited information set. Panels (a) and (b) display impulse responses constructed at time $t = 0$. Assuming the absence of further shocks, Panels (c) and (d) depict public impulse responses updated at time $t = 2$ (solid black lines) to reflect the actual evolution of the output gap under the true impulse response (solid with diamond markers). For comparison, public impulse responses constructed at times $t = 0$ and $t = 1$ (dashed black) are also shown.

Figure 3: Impulse Responses after an Information Shock

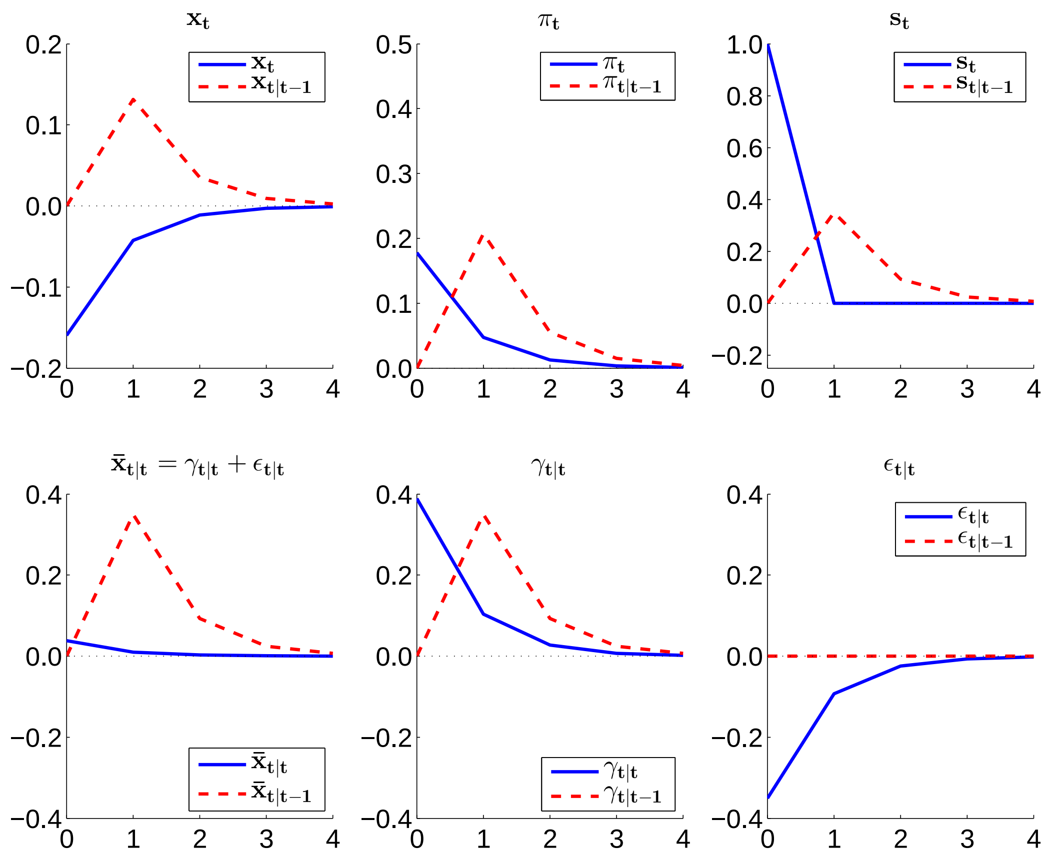

Note: Straight blue lines are impulse responses to a unit shock n_t, when the public observes an exogenous signal ($s_t = \gamma_t + n_t$) in addition to observing the policymakers output gap policy x_t. In these simulations, the true target components are held at zero ($\gamma_t = \varepsilon_t = 0$, the target $\bar{x}_t = \gamma_t + \varepsilon_t$ is denoted x_t^* above). Dashed red lines depict the evolution of prior beliefs.

Figure 4: Variations in Transparency

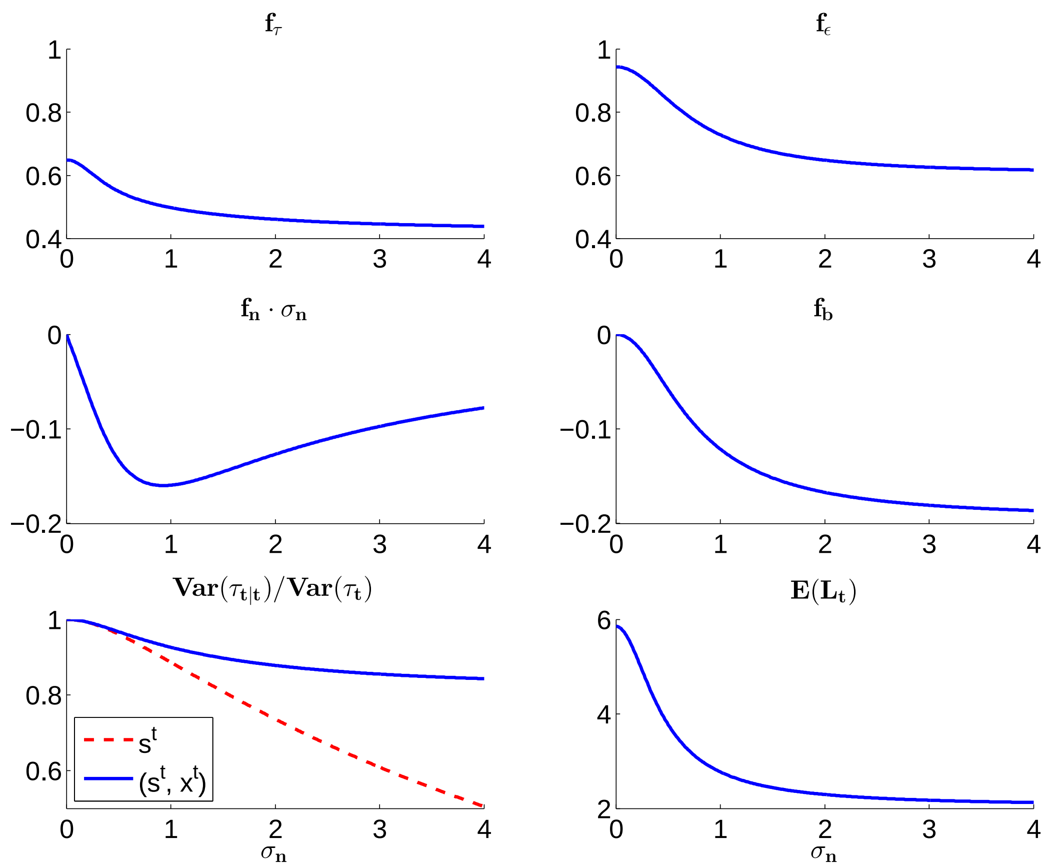

Note: Each panel depicts different model coefficients from the information shock model of Section 4 as the volatility of information shocks, σ_n is varied along the x-axis. All other model parameters are held at their baseline values described in the main text. Panels in the top and middle rows depict variations in response coefficients of optimal policy, which takes the form $x_t = f_\gamma \gamma_t + f_\varepsilon \varepsilon_t + f_n n_t + f_b \gamma_{t|t-1}$. (The middle left panel shows scaled response coefficients, $f_n \cdot \sigma_n$, to account for the size of the information shocks which is varied along the x-axis.) The bottom right panel depicts unconditionally expected policy losses $E(L_t)$. The bottom left panel compares the share of variations in the persistent target component explained by the exogenous signal s_t (red-dashed line), $\text{Var}\,(E(\gamma_t|s^t))/\,\text{Var}\,(\gamma_t)$, against the corresponding share explained by all information that is available to the public (solid blue line), $\text{Var}\,(E(\gamma_t|s^t, x^t))/\,\text{Var}\,(\gamma_t)$.