STUDY CENTER GERZENSEE

Managing Beliefs about
Monetary Policy under Discretion

Elmar Mertens

Working Paper 08.02

Managing Beliefs about
Monetary Policy under Discretion[*]

Elmar Mertens[†]
Federal Reserve Board

April 2009

[†]For correspondence: Elmar Mertens, Board of Governors of the Federal Reserve System, Mailstop 133, Washington D.C. 20551. email `elmar.mertens@frb.gov`. Tel.: +(202) 452 2916. . The views in this paper and any errors or omissions should be regarded as those solely of the author, and 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.

Managing Beliefs about
Monetary Policy under Discretion

Abstract

Monetary policy is most effective when public beliefs about future policies are actively managed. This is the appeal of policy rules and commitment strategies, typically absent under discretion. But when a policymaker has some private information — as is the case in reality — belief management becomes an integral part of optimal discretion policies, too.

Solving for optimal policy in a simple New Keynesian model, this paper shows how discretionary losses are reduced when the policymaker has private information. Furthermore, disinflations are pursued more vigorously, when the hidden information problem is larger, even when inflation is partly backward-looking.

1 Introduction

There are clear benefits from commitment in monetary policy. Many economic decisions in the private sector are forward-looking and depend on policy expectations. But typically, optimal policy is not time-consistent — for example in the presence of nominal rigidities. The ability to commit to future policies is then crucial for an effective monetary policy. This benefit is well-known from the rules-versus-discretion literature (Kydland and Prescott 1977; Barro and Gordon 1983b). But real-world policymakers generally retain a good deal of discretion in their decision making process,[1] begging the question whether these policymakers forgo substantial benefits from expectations management.

Most of the rules-versus-discretion literature is based on models of perfect information, symmetrically shared between the central bank and the public. In reality however, monetary policy is conducted under imperfect and asymmetric information. Asymmetric information is an inherent feature of delegated management; the conduct of monetary policy by a specialized central bank is no exception.

This paper shows how belief management becomes an integral part of discretionary policies, when the central bank has private information. In this case, the public will make inferences about the hidden information based on observed policy actions such that current policies *directly* affect inflation expectations.[2] The trade-offs faced by a discretionary policymaker resemble then those known from commitment problems.

Investigating the design of optimal policy when the central bank has proprietary information also contributes to the literature on optimal transparency. Morris and Shin (2002) caution against providing the public with too much information. But Woodford (2005) doubts whether their con-

---

[1]This is likely due to a variety of reasons, for example the inability to commit future policymakers, respectively the inability to commit future majorities in a policy making committee, to specific future actions. Then there are also reasons of model uncertainty and the practical impossibility of devising the kind of completely state-contingent plan prescribed by the optimal commitment literature, see for example Clarida, Gali, and Gertler (1999). By and large, these reasons also apply when implementing some form of inflation targeting.

[2]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.

clusions will be relevant in a forward-looking model, where the economy is mostly affected by expectations of future policies. The model analyzed here provides an intriguing counterexample to this conjecture, in that policy losses are lower under hidden information when comparing discretionary policies in a New Keynesian model. Since hidden information gives scope for belief management under discretion, the result occurs precisely for reasons stressed by Woodford (2005).

Attention is limited here to Markov-perfect policies. In the spirit of "bygones are bygones", Markov-perfect state variables equilibrium must be relevant for current payoffs. When the public is imperfectly informed, its prior beliefs matter for public payoffs and they become a distinct, endogenous state variable of the policy problem, which is influenced by policy actions. By managing this state of (public) beliefs, the policymaker indirectly responds to past policies, even when reputational mechanisms via history-dependent strategies, known from Barro and Gordon (1983b) or Chari and Kehoe (1990), are excluded from the analysis.

In Markov-perfect models, a current decision-maker can influence a future decision-maker only via endogenous state variables, such as capital or government debt. In the model presented here, belief management leads to Markov perfect outcomes that share similarities with those from models with commitment respectively reputational mechanisms. Previous research has already recognized how discretionary outcomes can be improved by adding endogenous state variables to the policy problem. Usually, this is done by modifying the central bank's loss function, for example by adding concerns for interest rate smoothing (Woodford 2003b) or by replacing inflation stabilization with price level targeting (Vestin 2006).

What is novel about the present paper, is how beliefs naturally emerge as such an endogenous state variable, without the need for modifying the central bank's loss function or other aspects of the economy. To the extent that hidden information problems are an essential feature of interactions between policymakers and the public, this suggests that the importance of discretionary biases in practice might be different, and likely smaller than what is suggested by full information models.[3]

The problem of "public learns about central bank" studied here is distinct from settings of

---

[3]In a related manner, Blinder's (1998) account of his time as Vice-Chairman of the Federal Reserve Board is dismissive about the relevance of time-inconsistency as a major distortion in real-world decisions at central banks.

"bank learns about economy" studied for example 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. Policy constraints like the Phillips Curve are largely preserved. In the linear quadratic case studied by Svensson and Woodford (2004), certainty equivalence holds and optimal policies are identical to the full information case when actual values are replaced by policymakers' expectations. A key complication for my paper is that the central bank is a strategic, not an atomistic player, who takes the public's inference problem into account when devising its policy. This changes the policy constraints in non-trivial ways.

The framework adopted here exclusively assigns the policymaker, and not the public, with superior information. This is an extreme assumptions. Reality is best described by dispersed information, endowing different bits and pieces of hidden knowledge to the private sector and policymakers. The policy constraints change in dramatic ways when agents are learning about the policymaker, because of his strategic position in the economy. Those strategic effects are the main concern of the paper.

The effects of hidden information on optimal policy are illustrated with a simple New Keynesian model — a model not chosen for its realism, but in order to document the differences with the symmetric information benchmark most clearly within a widely studied setting. The paper solves for the optimal discretion policy in a New Keynesian model where the output target of the policymaker is not directly observed by the public. The public only observes policy actions, but cannot disentangle whether the underlying shock to the output target is persistent or transitory. The policymaker faces a direct feedback from higher inflation expectations when choosing more expansionary policies cautioning him to temporarily boost aggregate activity at the expense of higher inflation. Compared to a full information model, a key difference is how optimal policy contracts the economy in response to inflationary beliefs. Moreover it does so more vigorously, the larger the credibility problems from hidden information. This result has important implications for the conduct of optimal disinflations.

To the best of my knowledge, my paper provides the first analysis of disinflations with an ex-

plicitly optimizing monetary policymaker and unknown policy targets.[4] The results confirm conjectures by Sargent (1982) and Bordo et al. (2007) about the necessity to disinflate more quickly, when credibility is at stake. Other economists, for example Gordon (1982), have rather argued for prolonged and modest disinflation paths when inflation is persistent. Strikingly, my result is shown to carry over also to a setting with a hybrid Phillips Curve, where inflation persistence is partly exogenous. Evidently, disinflation costs are higher in such a setting. However, by bringing down inflation expectations early on a more aggressive disinflation policy still minimizes these costs, since it avoids inflation to persist based on ill-founded beliefs.

The information structure used here is similar to the models of Faust and Svensson (2001, 2002) and Cukierman and Meltzer (1986) who cast their models within similar linear-quadratic settings, but without providing a general framework capable of handling various models with endogenous state variables. Faust and Svensson focus on the welfare effects of credibility with a Lucas-supply curve. Using a forward-looking Phillips Curve, their results can be confirmed and extended here: Policy losses are reduced when output targets are unobservable, such that there is an explicit role for public beliefs. This disciplines the pursuit of persistent output targets, even when time-consistency is imposed on policy.

So far, problems of this kind have mostly been analyzed in highly stylized and often static settings.[5] But the models used for policy analysis are typically dynamic and of larger scale. The technical appendix to this paper presents a flexible, yet tractable way to analyze optimal policy under hidden information, which is applicable to the kind of DSGE models used in policy analysis. The procedure remains tractable and transparent by relying on a linear-quadratic representation of the policy problem driven by Gaussian shocks. A key complication for models with imperfect information is to track the distribution of public beliefs. In a linear, homoscedastic setting, that

---

[4]The closest counterpart to my analysis should be the work of Ireland (1995) who imposes a sluggish response of public beliefs to policy announcements.

[5]See for example the classic contributions by Backus and Driffill (1985a), Canzoneri (1985) and Cukierman and Liviatan (1991), where my definition of static includes also repeated play of one-period games. More recent work includes the papers by Ball (1995) and Walsh (2000). Fully dynamic, but limited in size, are the models of Gaspar, Smets, and Vestin (2006), Faust and Svensson (2001, 2002), Cukierman and Meltzer (1986). A more detailed discussion of the literature can be found in Section 5.

collapses to tracking the evolution of means via the Kalman filter.

The remainder of this paper is structured as follows. Section 2 introduces hidden information in a textbook version of the New Keynesian model and shows how hidden information changes the policy problem. An extension incorporating belief shocks is shown in Section 3. Implications for disinflation strategies are analyzed in Section 4. The related literature is discussed in Section 5. Section 6 concludes the paper. A technical appendix extends the methods used here to a general class of linear quadratic policy problems.

2   A Simple Model of Hidden Information

This section illustrates the issues arising from hidden information with a simple textbook version of the New Keynesian model. The model model is purely forward-looking and the signal extraction problem is univariate. The next section extends this model to a setting where a hybrid Phillips Curve interacts with shocks from a richer information structure.

2.1   New Keynesian Economy

The model is largely identical to the textbook model of optimal policy in a New Keynesian model known from Clarida, Gali, and Gertler (1999), Walsh (2003) or Woodford (2003a). The only difference is a stochastic preference shock to the policymaker's objective function, which is unobservable to the public. Otherwise my model and its notation follow closely Gali (2003) where further details can be found. A key feature of the model is that inflation is determined purely by public expectations of current and future policies. This puts centerstage the concerns of the public about the policymaker's intentions.

Private Sector

As in the textbook model, aggregate decisions of the private sector are represented by the New Keynesian Phillips and IS curves. In this simple model, IS curve and the short term interest rate

are even redundant and the output gap can be used as 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

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

where π_t is inflation and x_t is the output gap[6]. The parameter β is the representative agent's discount factor and κ is a reduced form parameter influenced amongst others by the frequency of price-setting.[7] For any variable z_{t+1}, $z_{t+1|t}$ denotes its private sector forecast. The underlying information set will be explained later.

The output gap measures the difference between actual output and its natural rate. The latter would be the output of the economy if there were no nominal frictions.[8] My discussion will exclusively focus on monetary shocks that leave the natural rate unaffected. Conditional on those shocks, variations in the output gap are thus identical to variations in output and consumption.

Policy Objectives

The policymaker seeks to minimize a present value of expected losses

$$E_t \sum_{k=0}^{\infty} \beta^k \left\{ \pi_{t+k}^2 + \alpha_x (x_{t+k} - \bar{x}_{t+k})^2 \right\} \tag{2}$$

---

[6]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.

[7]Details are given by Gali (2003, p. 159) from whom notation is adopted.

[8]King and Goodfriend (1997) explain how the New Keynesian model can be separated into a core real business-cycle model (RBC), which evolves as if there were no nominal frictions, and a set of "gap" variables that track the difference between the RBC core and the actual economy. This separation has been widely adopted for example in the textbooks of Walsh (2003), Woodford (2003a) and Gali (2008).

with $\alpha_x \geq 0$. The expectations operator E_t reflects the policymaker's information set, to be described later. The non-standard feature of the loss function is the time-varying target for the 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[9]. However, the information structure used below will require that $\bar{x}_t$ is not observed by the private sector. To keep the model close to the NK 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.[10] 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 is capturing a form of heterogeneity otherwise not present in the model. In particular (2) does not necessarily represent a social welfare function. Faust and Svensson (2001) use a similar loss function for the policymaker. Their notion of representative welfare would then be to evaluate (2) at the average output target (here: zero)[11], $L_t^R = \pi_t^2 + \alpha_x x_t^2$. But without specifying the underlying heterogeneity and associated welfare weights this is at best an aggregation with unknown distributional consequences.

In reality, short-term interest rates are the typical instruments of monetary policy. But in this

---

[9]For starters, time-variation in the output target could arise from variations in wedges between the frictionless and the efficient level of output. Time-varying markups would for example shrink distortions from monopolistic competition. There are non-monetary tools to fight such distortions, for example the kind of fiscal tools discussed by Gali (2003). $\bar{x}_t$ could then capture changes in the government's policy of handling these distortions.

[10]In the real world, pressures mounted on central bankers appear to be a recurring, though not necessarily permanent feature. For example, in the short history of the ECB there were the early attempts by German Finance Minister "Red" Oskar Lafontaine and later overtures from the French President Nicolas Sarkozy.

[11]In the results discussed below there will not be a conflict in ranking outcomes under this measure as opposed to the policymaker's objective.

simple model, the short term interest rate can be perfectly substituted by the output gap as policy control. The IS curve is then redundant for determining equilibrium.

Discretionary Policy under Symmetric Information

Before turning to the informational structure of the model, it is helpful to study optimal policy when there is symmetric information. For the time being, let the output target follow a univariate AR(1) process

$$\bar{x}_{t+1} = \rho\, \bar{x}_t + e_{t+1} \quad \text{where} \quad e_{t+1} \sim N(0, \sigma_e^2) \quad \text{and} \quad |\rho| < 1$$

which is mutually observed by the policymaker and the public.[12] Under symmetric information, their expectations coincide such that $z_{t+1|t} = E_t z_{t+1}$ for any variable z_t.

Lacking a commitment technology, the policymaker can always reoptimize his policies and for each optimization he takes his future choices as given. Sine there are only exogenous state variables, he takes the public's inflation expectations as given, too.[13] Only Markov-perfect, discretionary equilibria are considered. This excludes for example trigger strategies to support commitment outcomes.

The solution to this problem is well known. The first order condition balances the inflation cost against the desire to attain the output target:

$$\alpha_x(x_t - \bar{x}_t) + \kappa \pi_t = 0 \tag{3}$$

(Section 2.3 below, will compare this optimality condition against its counterpart under hidden information.) Substitution of (3) into the Phillips Curve yields the following Markov-perfect poli-

---

[12]The process is mean zero and allows also for negative targets. But all variables are in deviation from steady state. By allowing for a (known) average target, this would lead to the classic inflation bias in steady state. (In the context of the present model, details can be found in Woodford (2003a).) To be consistent with non-zero inflation in steady state, the Phillips Curve is then viewed as allowing for indexation to the steady state rate of inflation as in Yun (1996).

[13]In general, the policymaker could not take inflation expectations as given numbers but as a given mapping from expected future state values, where the latter may be partly under his control. This will be the case under hidden information.

9

cies:

$$x_t = \frac{\alpha_x(1-\beta\rho)}{\kappa^2 + \alpha_x(1-\beta\rho)}\,\bar{x}_t \equiv \bar{f}\,\bar{x}_t \quad \text{and} \quad \pi_t = \frac{\kappa}{1-\beta\rho}\,\bar{f}\,\bar{x}_t \tag{4}$$

Inflation and output gap inherit the dynamic properties of the target process. A well known property of optimal policies in a linear quadratic framework is their certainty equivalence, which holds here, too, since $\bar{f}$ does not depend on the volatility σ_e of the target shocks. Under hidden information, this will be different.

Sensibly, $\bar{f}$ is bounded between zero and one. In principle, the policymaker could always attain the output target by choosing $\bar{f} = 1$, but for $\alpha_x < \infty$ this has to be weighed against the inflation resulting from this policy. At the other extreme, there would be no inflation if $\bar{f} = 0$, but only at the cost of missing the target, which matters if $\alpha_x > 0$. Values outside the zero to one range would lead to further target deviations *and* be associated with unnecessary inflation. This will be useful to bear in mind when analyzing policies under hidden information.

Policies with $\bar{f}$ close to unity will be called "bold" and it is instructive to see how policy depends on the preference weight α_x and the persistence of the target process. Inspection of (4) reveals the intuitive property that policies get bolder the higher the preference weight on output, in fact $\bar{f}$ varies between zero and one when α_x is varied between zero and infinity.

Policies are less bold, when the target is persistent. Higher persistence of the target causes higher persistence in policy and thus higher inflation. This is a dynamic version of the inflation bias known from Kydland and Prescott (1977) and Barro and Gordon (1983a) and similar to the stabilization bias known from Svensson (1997).

Under hidden information there will be persistent and transitory shocks to the output target, neither of them being directly observable to the public. As in the full information case, what matters for the inflation response to a policy shock is its perceived persistence. The policymaker will then seek policies that are as bold as possible, while trying to keep perceived persistence as low as possible.

2.2 Hidden Information

Hidden information is introduced by assuming that the public can observe only policy, x_t, but not shocks to the policy target. To make the public's signal extraction interesting, the target is henceforth driven by two components, one persistent, one transitory:

$$\bar{x}_t = \tau_t + \varepsilon_t \qquad\qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2) \tag{5}$$

$$\tau_{t+1} = \rho\, \tau_t + \eta_{t+1} \qquad\qquad \eta_t \sim N(0, \sigma_\eta^2) \quad \text{and} \quad 0 < |\rho| < 1 \tag{6}$$

The private sector has no structural uncertainty about the economy. All parameters are known, including the specification of the target process. The public must however infer the *realizations* of τ_t and ε_t based on the observed history of policies, denoted x^t.[14]

The policymaker observes the complete history of the target components and his expectations are typically different from those of the public. As before, for any variable z_t, the policymaker's expectations are denoted $E_t z_{t+1} = E(z_{t+1}|\tau^t, \varepsilon^t)$ with the obvious property $z_t = E_t z_t$. Public expectations are $z_{t+1|t} = E(z_{t+1}|x^t)$. By construction, $x_{t|t} = x_t$ and $\pi_{t|t} = \pi_t$ (since inflation is a choice variable of the private sector) but typically $\tau_{t|t} \neq \tau_t$ and $\varepsilon_{t|t} \neq \varepsilon_t$.

Surprises in z_t relative to the public's past information will be called "innovations". Formally, they are defined as

$$\tilde{z}_t \equiv z_t - z_{t|t-1}$$

Innovations provide an orthogonal decomposition of the public information set since $\tilde{z}_{t|t-1} = 0$. Even though they are unpredictable from the public's perspective, they may well be predictable based on the complete information set, and typically $E_{t-1}\tilde{z}_t$ will not be identical to zero.

Since the model is linear with Gaussian disturbances, rational expectations of the public can be computed recursively from the Kalman filter. Given prior beliefs $z_{t|t-1}$ and $x_{t|t-1}$, the public

---

[14]In principle, this includes also the history of inflation rates π^t. But as a choice variable of the private sector, inflation merely reflects the private sectors information set, without providing additional information beyond x^t.

observes a realization of policy x_t and updates its beliefs according to

$$z_{t|t} = z_{t|t-1} + K_z \tilde{x}_t \quad \text{with Kalman gain} \quad K_z \equiv \frac{\text{Cov}\,(z_t, \tilde{x}_t)}{\text{Var}\,\tilde{x}_t} \tag{7}$$

A convenient property of the Kalman update is that it preserves the linearity of the model. The difference with adaptive expectations is that the gain coefficient is an endogenous parameter, identical to the least squares slope of projecting z_t on $\tilde{x}_t$. The present model is particularly simple since there is only one observable, x_t, such that K_z is a scalar. (A multivariate setting will be illustrated in Section 3.)

Signal Extraction for Given Policy

As will be verified below, the optimal policy is linear and has the form

$$x_t = f_\tau \tau_t + f_\varepsilon \varepsilon_t + f_b \tau_{t|t-1} \tag{8}$$

for some scalars f_τ, f_ε and f_b. Compared to the symmetric information case, the dependence on $\tau_{t|t-1}$ is novel. It captures policy responses to public beliefs. As will be shown shortly, it influences the persistence of policy shocks, which is a crucial factor in determining inflation.

The public belief system is a straightforward application of the Kalman filter with (6) as state equation and (8) as measurement equation. In the parlance of time-series econometrics, policy poses an unobserved components model to the public. Key for the Kalman filter is the ratio of policy loadings on the *realized* components of the output target, f_ε/f_τ. Only these loadings, and not f_b, are relevant for the Kalman filter. (Details are given in Appendix B.) This "mixing ratio" f_ε/f_τ determines how much a policy innovation reveals about τ_t instead of ε_t. It allows the policymaker to change the signal-to-noise ratio in the public's signal extraction problem.

From the perspective of the public, policy is driven by the iid innovations $\tilde{x}_t$ and it has an

innovations representation in the form of an ARMA(1,1) process:

$$x_t = \rho x_{t-1} + \tilde{x}_t + \rho\psi\tilde{x}_{t-1} \tag{9}$$

$$\text{with} \quad \psi = (f_\tau + f_b) K_\tau - 1$$

For the public, the above innovations representation is observationally equivalent to the hidden components representation of policy (8). Both generate the same variances and autocovariances of policy, whilst implying different impulse responses as will be illustrated below. Via ψ, the persistence of this ARMA depends on the policy coefficients f_τ, f_ε and f_b. For plausible assumptions of the policy coefficients, ψ is bounded between zero and minus one. For $\psi = 0$, persistence is largest as policy follows an AR(1) with auto-correlation equal to ρ. For $\psi = -1$ both roots of the ARMA(1,1) cancel and policy is iid. (Details can be found in Appendix A.)

Together with the Phillips curve (1), the innovations representation of policy is sufficient to determine inflation in a way which crucially depends on the "average persistence" of policy as captured by the ARMA roots ρ and ψ.

$$\pi_t = \kappa \sum_{j=0}^{\infty} \beta^j x_{t+j|t} = \frac{\kappa}{1 - \beta\rho}\left((1 + \beta\rho\psi)\tilde{x}_t + x_{t|t-1}\right) \tag{10}$$

The policy function has two levers to affect the persistence of x_t: First, there is the mixing ratio, which has been discussed above. If policy largely ignores the persistent target, i.e. if f_ε / f_τ is large such that K_τ is close to zero, the MA root gets close to cancel the AR root and policy is (correctly) perceived to be almost iid. In this case, inflation also approaches the solution (4) under symmetric information with $\rho = 0$.

But due to the second lever, f_b, things need not collapse to the AR(1) case, when the mixing ratio tends to zero. In this case, ψ converges to $\rho \cdot f_b$, which is not necessarily zero.[15] f_b represents the marginal reaction to people's prior beliefs and affects the persistence of policy, too. A negative

---

[15]Since $x_t = x_{t|t}$ in this simple model, it follows that $f_\tau K_\tau + f_\varepsilon K_\varepsilon = 1$. When the mixing ratio goes to zero, this collapses to $f_\tau K_\tau = 1$.

f_b counteracts policy persistence induced by τ_t. The marginal reaction to beliefs is likely negative since beliefs $\tau_{t|t-1}$ will be inflationary; this conjecture will be verified in Section 2.4.

To keep inflation low, it is tempting to conclude that the policymaker should better ignore the persistent output target. Alternatively, a high mixing ratio could be chosen, with a higher responsiveness to transitory than persistent shocks, for example $f_\tau = 1$ and $f_\varepsilon = 100$. But neither choice would likely be a sensible policy, since output plays not only an informational role. Attaining the output targets matters, too; calling for $f_\tau = f_\varepsilon = 1$ and $f_b = 0$. For example, ignoring the persistent target by setting $f_\tau = f_b = 0$ alleviates inflationary cost, but it also leads to *persistent* shortfalls from the τ-target. Neither would it appear sensible to overshoot the output target, for example by setting $f_\varepsilon = 100$. The optimal trade-off is the subject of the next sections. But an important restriction imposed by rational expectations has already become clear: at least on average actual policies must match public perceptions.

2.3 The Discretionary Policy Problem

This section sets up the discretionary policy problem for the simple, purely forward-looking New Keynesian model when there is the above structure of hidden information. Extensions of the model, including a hybrid Phillips Curve, will be analyzed in subsequent sections of this paper. The concepts and methods presented here are generalized to a wider class of linear quadratic models in the technical appendix of this paper.

Markov Perfect Equilibria

Attention is limited here to Markov-perfect equilibria, which exclude reputational mechanisms via the kind of history-dependent strategies considered by Barro and Gordon (1983b) or Chari and Kehoe (1990) and avoids the associated multiplicity of equilibria. In the spirit of "bygones are bygones", state variables in a Markov-perfect equilibrium must be relevant for current payoffs.[16] In

---

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

the symmetric information setting shown above, these were the contemporaneous values τ_t and ε_t (but not any elements of their history). Both of these state variables evolve in a purely exogenous fashion which accounts for the myopic behavior of discretionary policy under symmetric information: In Markov-perfect models, a current decision-maker can influence a future decision-maker only via endogenous state variables, like capital or government debt. This channel is however absent in the symmetric information version of the New Keynesian model.

Once hidden information is introduced, an additional state variable becomes relevant: Since the public observes only x_t but neither τ_t not ε_t, it is the public beliefs about the target components which are relevant for public payoffs. Precisely, it is the prior beliefs ($\tau_{t|t-1}$ and $\varepsilon_{t|t-1}$) and not the posteriors ($\tau_{t|t}$ and $\varepsilon_{t|t}$) which qualify as state variables for the time t decision problem, since the latter are already influenced by time t policies. In the present setting, ε_t is iid and $\varepsilon_{t|t-1} = 0$ so only $\tau_{t|t-1}$ needs to be tracked. The vector of Markov-perfect state variables is then

$$S_t = \begin{bmatrix} \tau_t & \varepsilon_t & \tau_{t|t-1} \end{bmatrix}'$$

The transition equation for the new state variable is given by the Kalman Filter. The response of beliefs to policy depends on the Kalman gain K_τ, which reflects how much policy reacts to τ_t.

$$\tau_{t+1|t} = \rho\tau_{t|t} \quad \text{and} \quad \tau_{t|t} = \tau_{t|t-1} + K_\tau \tilde{x}_t \tag{11}$$

The discretionary policymaker retains the freedom to reoptimize his policies at each point in time. On the one hand, this allows a recursive representation of the policy problem as a dynamic program. On the other hand, he does not commit to future policies so these have to be taken as given in the decision problem. To be precise, what is taken as given is how the policymaker reacts to future state variables: Future policies are not given numbers but a given function of future state variables. This distinction is important here, since one of the state variables, $\tau_{t+1|t}$, is under the influence of current policy so that future outcomes can be influenced. The continuation value of his dynamic program is a function of future states, denoted $V^0(S_{t+1})$ and the policy objective is to

minimize

$$\pi_t^2 + \alpha_x(x_t - \tau_t - \varepsilon_t)^2 + E_t V^0(S_{t+1}) \tag{12}$$

The linear quadratic nature of the model allows to guess (and verify) that the value function will

be quadratic and policies linear in the state vector, which simplifies the analysis considerably:

$$V^0(S_{t+1}) = S'_{t+1} V^0 S_{t+1} + v^0$$

$$\Rightarrow \quad E_t V^0(S_{t+1}) = 2v_{13}^0 \rho\, \tau_t\, \tau_{t+1|t} + v_{33}^0 \tau_{t+1|t}^2 + t.i.p.$$

for some positive definite matrix V^0 with elements v_{13}^0, $v_{33}^0 > 0$ and a scalar v^0. Throughout this

paper, a zero superscript "0" indicates coefficients embodying a guess about (future) policy and

"$t.i.p.$" are terms independent of time t policy.

The time-invariant solution to the discretionary policy problem has the linear form anticipated

in (8). In principle, the policymaker is free to deviate from this "rule" at any time. He will just not

find it optimal to do so.

An important constraint on the policy problem is the optimality of beliefs and decisions in the

private sector. Optimality of beliefs are captured by the Kalman filter (7) and the the policymaker

sees himself faced with a fixed Kalman gain K_τ^0 when contemplating his policy problem. Optimal

decisions of the private sector are represented by the Phillips Curve (1) where the policymaker

takes as given how inflation expectations are related to future state variables; $\pi_{t+1|t} = g^0 \tau_{t+1|t}$ for

some scalar g^0. To sum up, the policy problem is to minimize

$$V_t = \min_{x_t, \pi_t, \tau_{t+1|t}} \pi_t^2 + \alpha_x(x_t - \tau_t - \varepsilon_t)^2 + 2v_{13}^0 \rho\, \tau_t\, \tau_{t+1|t} + v_{33}^0 \tau_{t+1|t}^2 + t.i.p. \tag{13}$$

$$\text{s.t.} \quad \pi_t = \beta g^0 \tau_{t+1|t} + \kappa x_t \tag{14}$$

$$\tau_{t+1|t} = \rho(1 - K_\tau^0(f_\tau^0 + f_b^0))\tau_{t|t-1} + \rho K_\tau^0 x_t \tag{15}$$

whose solution is indeed of the form anticipated in (8). Whilst beliefs embodied in g^0, V^0 and

16

K_τ^0 are taken as given in the policy problem, in equilibrium they must be consistent with the solution to the policy problem. This poses an intricate fixed point problem. Fixed points between current expectations and future policy as in g^0 and V^0 are common in Markov-perfect models under symmetric information. What is new is the fixed point between current policy and beliefs about the systematic relationship between current policy and states contained in the Kalman gain K_τ^0.

Changed Policy Trade-Offs with Belief Management

The first-order conditions of (13) require optimal policy to satisfy

$$\alpha_x(x_t - \bar{x}_t) + \kappa\pi_t + \rho K_\tau^0 \mu_t = 0 \tag{16}$$

where μ_t is the multiplier on the belief constraint (15). It is the term involving μ_t which distinguishes the optimality condition (16) from its counterpart under symmetric information (3) discussed above.

To shed some light on the fixed point considerations behind the solution to (13), suppose that the output target is positive and the policymaker must balance an increase in output against its inflationary costs. The marginal value of relaxing the belief constraint is likely positive, owing to the positive autocorrelation in the persistent component of the target. Likewise, the Kalman gain K_τ will be positive, since policy will co-move positively with the target. The new "belief term" $\rho K_\tau^0 \mu_t$ in (16) will then caution the policymaker against pursuing the output target too aggressively. As will be seen in the numerical analysis below, optimal policy will be less bold under hidden information — except when shocks to τ_t are so rare that the Kalman gain K_τ is very small.

The change in policy trade-offs under hidden information can be nicely illustrated with a picture similar to Kydland and Prescott (1977). Under symmetric information, the policymaker's indifference curves over output and inflation are concentric around $\pi_t = 0$ and $x_t = \bar{x}_t = \tau_t + \varepsilon_t$. The optimality condition (3) seeks the tangency point between the indifference curves and the pol-

17

icy constraint. The latter being the Phillips Curve with intercept $\beta\pi_{t+1|t} = \beta\bar{g}^0\rho\tau_t$ for some $\bar{g}^0$. This is depicted by the dashed lines in Figure 1. In equilibrium, $\bar{g}^0$ must be identical to the optimal policy coefficient computed in (4), which is a positive number. That is, the larger policy responds to a given level of the persistent target, the higher the intercept it faces in equilibrium.

[Figure 1 about here.]

Under hidden information belief management comes into play and changes the picture. To reach some substantive conclusions, I am willing to make the following assumptions about the policy coefficients. Apart from being plausible, they will be verified to be true in the computations below for a wide range of calibrations. First, policy should react positively to target shocks, $f_\tau^0 > 0$. Second, policy seeks to counteract belief $f_b^0 < 0$. But third, it still seeks to accommodate a target, even when its realization coincides with public beliefs: $f_\tau^0 + f_b^0 > 0$.[17] These imply that K_τ and g^0 are positive.

A key result is that hidden information steepens the slope of the Phillips Curve when compared against the symmetric information case. Substituting the belief dynamics (11), the Phillips Curve becomes

$$\pi_t = \beta g^0 \rho \left(1 - K_\tau^0(f_\tau^0 + f_b^0)\right) \tau_{t|t-1} + \left(\kappa + \beta g^0 \rho K_\tau^0\right) x_t \tag{17}$$

The steepening of the Phillips Curve worsens the policy trade-off and makes policies less bold with respect to both target components. Underlining the importance of beliefs, the intercept of the Phillips Curve depends now on the public's prior beliefs, $\tau_{t|t-1}$, instead of the actual value of τ_t. Coming out of steady state with $\tau_{t|t-1} = 0$, this alone makes policies bolder than otherwise. An important aspect for the fixed point computations is that, via K_τ, the slope of the Phillips Curve becomes ever steeper the bolder policies are with respect to τ_t, which again tames the boldness of equilibrium policies.

---

[17]This assumption implies that the public expects policies to be expansionary, $x_{t|t-1} > 0$, when $\tau_{t|t-1} > 0$.

Belief management changes the indifference curves as well. Most importantly, output acts as a signal about the persistence of policy targets which again influences the evaluation of future losses in the policy problem. This shifts output preferences, such that they are not centered around $\bar{x}_t = \tau_t + \varepsilon_t$ anymore. Substituting again the belief dynamics, the indifference curves can be computed from

$$\pi_t^2 + \alpha_x(x_t - \tau_t - \varepsilon_t)^2 + \gamma_0\, x_t^2 + \gamma_1\, \tau_t\, x_t + \gamma_2\, \tau_{t|t-1}\, x_t \tag{18}$$

where the scalars $\gamma_0 > 0$, $\gamma_1 \lessgtr 0$ and $\gamma_2 > 0$ depend on the coefficients of (18). [18] These indifference curves are centered around $\pi_t = 0$ and

$$x_t^* = \frac{\alpha}{\alpha + \gamma_0}\underbrace{(\tau_t + \varepsilon_t)}_{\bar{x}_t} - \frac{1}{2}\gamma_1\, \tau_t - \frac{1}{2}\gamma_2\, \tau_{t|t-1}$$

Regardless of slope and intercept of the Phillips Curve, x_t^* is the "maximally desirable" level of output. Apart from its dependence on the original target term $\bar{x}_t$, it shifts both with the actual and perceived level of the persistent target component τ_t. But for starters consider a transitory shock to the output target, say $\varepsilon_t = 1$ whilst $\tau_t = \tau_{t|t-1} = 0$: Any policy response will partly be attributed to a persistent shock and thus increase $\tau_{t+1|t}$ causing future inflation. The associated losses to the policymaker are captured by the γ_0 term of the indifference curves. Independently of the Phillips Curve, the policymaker does then not even desire to attain that transitory target but only a fraction $\alpha/(\alpha + \gamma_0)$ thereof.

Since $\gamma_2 > 0$, public beliefs $\tau_{t|t-1}$ shift the indifference curves towards lower output levels. While γ_1 cannot be signed analytically, it happens to be positive over the range of calibrations considered below and this contributes to making policy less bold. All in all, prior beliefs of the public $\tau_{t|t-1} > 0$ caution policy in two ways: First they increase inflation immediately (the intercept of the Phillips Curve) and — if not counteracted by current policy — they herald future inflation and

---

[18]It is straightforward to show that $\gamma_0 > 0$ follows from the positive definiteness of the value function, and $\gamma_2 > 0$ from the aforementioned assumptions on the policy coefficients. Analytically, γ_1 cannot be signed, but for the variety of calibrations considered in the numerical simulations below it turns out to be positive.

shrink the "maximally desirable" level of output, x_t^*, towards zero.

2.4 Optimal Policy in the Simple Model

This section presents results for the optimal policy. Calibration values are taken from Gali (2003) with equally weighted policy preferences ($\alpha_x = 1$) and equal-probable shocks to the target components ($\sigma_\eta = \sigma_\varepsilon = 1$), see Table 1.[19] The solution algorithm for the underlying fixed point problem is discussed in the technical appendix.

[Table 1 about here.]

Optimal Mixing Ratio and Belief Responses

Key statistics of the policy function are the mixing ratio f_ε / f_τ, which governs the Kalman gains, and f_b via which policy responds to prior beliefs. As anticipated, f_b is negative. The policy response to $\tau_{t|t-1}$ is synonymous with counteracting inflation expectations of the public formed in the past. There is a one-to-one correspondence between the public's prior beliefs of the hidden state and the public's inflation expectations in this simple model:

$$\pi_{t|t-1} = \frac{\kappa}{1 - \beta\rho} \left(f_\tau + f_b \right) \tau_{t|t-1}$$

How optimal policy seeks to quell past beliefs can be seen from the impulse response shown in Figure 2. The first two columns show responses to shocks in τ_t and ε_t. The third column documents responses to initial conditions $\tau_t = 0, \varepsilon_t = 0, \tau_{t|t-1} = 1$. This corresponds to a situation where the policymaker is faced with erroneous beliefs about his inflationary output preferences. The optimal response is a prolonged contraction until beliefs and outcomes have settled back in steady state after about four periods. Given that the New Keynesian model generally lacks endogenous persistence, the length of this learning process is a remarkable outcome echoing the results of Erceg

---

[19]Given the limited range of shocks considered, the calibration is not designed to match the level of variations observed in the data.

and Levin (2003). Moreover, the effect of fighting past beliefs is also present in the other impulse responses. When the true target shock is iid, this leads to a contractionary policy one period after the shock. This pattern is similar (though not fully identical) to commitment policies under full information. In both cases, a credible promise to undo expansionary shocks in the future lowers inflation expectations; similar to the disciplinary channel emphasized by Walsh (2000), Faust and Svensson (2001) and Gaspar, Smets, and Vestin (2006).

[Figure 2 about here.]

The other lever of policy is the mixing ratio, which is higher compared to the full information case.[20] Under hidden information, policy is less bold in its pursuit of persistent output targets. This lowers the signal-to-noise ratio in the public's signal extraction problem and the public (correctly) places a lower probability on a policy innovation $\tilde{x}_t$ being caused by a persistent target shock.

Innovation Responses

The model with hidden policy components is observationally equivalent to a symmetric information model where the policy target follows a univariate ARMA(1,1). Both yield the same second moments and have identical likelihoods. But there is an important difference: The hidden components model distinguishes different sets of impulses responses, which can be associated with different episodes in monetary policy.

[Figure 3 about here.]

The differences between true impulse responses and public beliefs are illustrated in Figure 3. For output and inflation, the figure shows two sets of impulse response: First, the expected responses computed by the public, after observing a unit innovation in policy, $\tilde{x}_t$, at time zero. After its initial upwards jump, output remains expanded at about half its impact value and decays persistently thereafter. The inflation path is equally equally elevated and persistent.

---

[20]For the baseline calibration, the mixing ratio is 1.2340 under symmetric information and 1.2866 under hidden information.

Secondly, the figure shows the true responses to the structural shocks τ_t and ε_t, computed under the full information measure spanned by (τ^t, ε^t). They are scaled such as to yield a unit innovation in output as well. After a shock to the persistent target, τ_t, policy is persistently more expansive than originally expected by the public. The difference between these two sets of impulse response represents the errors of public forecasts made in the initial period. As the structural responses unfold, the public learns about the true nature of the shock. The figure also shows how public beliefs are updated in subsequent periods, leading to persistent upwards, respectively downwards revisions of beliefs. The innovations responses are rational and on average correct. Persistently positive forecast errors to a shock in τ_t are offset by persistently negative forecast errors when a shock to ε_t occurs.

When particular periods are supposed to have been dominated by one set of shocks rather than another, 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. Their model uses a Gaussion 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.

Sensitivity Analysis of Policy Coefficients

Policy trade-offs are particularly affected by two parameters: The relative variance of transitory to persistent target shocks and the preference weight α_x, whereas increases in the slope of the Phillips Curve, κ, affect policy trade-offs similarly to decreases in α_x.[21] When considering changes in the importance of the target components ε_t and τ_t, the overall variance of the output target will be fixed at some level $\sigma_{\bar{x}}^2$. Denoting the weight on τ_t by $\omega \in [0; 1]$ this translates into

$$\sigma_\varepsilon^2 = (1 - \omega)\sigma_{\bar{x}}^2 \quad \text{and} \quad \sigma_\eta^2 = \omega(1 - \rho^2)\sigma_{\bar{x}}^2$$

---

[21]The loss function can be written as $L_t = \kappa^2 \bar{\pi}_t^2 + \alpha_x(x_t - \bar{x}_t)^2$, where $\bar{\pi}_t = \sum_{j=0}^{\infty} \beta^j x_{t+j|t}$.

Figure 4 documents changes in the policy coefficients f_τ, f_ε, f_b as well as the mixing ratio f_ε/f_τ due to variations in ω and α_x. The upper panels also show the corresponding values of f_τ and f_ε under symmetric information. Because of certainty equivalence, their surfaces are flat along the ω-axis. When there is hidden information, f_ε is uniformly smaller than under symmetric information. This is caused by the public's inability to distinguish between realizations in the two target components. Any innovation in x_t will be partly attributed to have been caused by the persistent component τ_t. This has two adverse effects in the first-order condition (16): First, if the true shock was to the iid component ε_t, inflation will be higher compared to the full information case. Second, since the public expects the target change to have some persistence, there will also be inflationary costs in the future, tracked by μ_t. Both effects caution the policymaker and lower f_ε compared to the full information case.

[Figure 4 about here.]

Because of the second effect, f_τ is mostly smaller under hidden information as well. As the public underestimates the persistence of policy after a shock to τ_t (see Figure 3), inflation is lower than under full information. This would give the policymaker some slack in pursuing the output target, if it were not mostly outweighed by the marginal effect of policy on beliefs, which are represented by the term $\rho K_\tau \mu_t$ in the first-order condition (16). However, when the probability of a persistent shock is very small (ω close to zero) or when the policymaker is known not to care much about attaining it (α_x small), the Kalman gain K_τ will be small and public beliefs $\tau_{t|t}$ will be very insensitive to policy and their importance vanishes in (16). In those cases, persistent shocks are very hard to detect for the public. But when they occur, the policy response can be bolder than under symmetric information as is shown in Panel (a) of Figure 4.

It is worth recalling that a higher mixing ratio and a more negative reaction to prior beliefs increase the persistence of the policy process, causing higher inflation. Indeed, as can be seen from Panel (c) of the figure, f_b is negative everywhere. The belief reaction is strongly negative, when there is more weight on inflation in the loss function (smaller values of α_x) and when persistent

shocks are more prevalent (ω close to one). Both cases make it more important, respectively more likely, that inflationary beliefs are kept in check.

Under discretion, the policymaker takes the public's belief system as given, without actively seeking influence it, whereas commitment policymaker would have to consider the systematic effects of his actions on the Kalman gain for example. Still it is instructive to see how policy affects the public's signal-to-noise ratio via the mixing ratio. As shown in Panel (d), this ratio increases when policy preferences place more weight on output than inflation. As inflation becomes more and more costly for the policymaker, f_τ is decreasing faster than f_ε, which makes it ever harder for the public to detect persistent policy changes. When changing ω, the mixing ratio is largest for intermediary values, typically above $\omega \geq 0.5$. In this range, hidden information problem is most prevalent and a high mixing ratio helps to lessen the sensitivity of beliefs to policy. When ω approaches unity, the public can expect any policy to be caused by a persistent shock with near certainty and the mixing ratio drops to its full information level. When the target is almost exclusively driven by iid shocks ($\omega \to 0$) expected future inflation drops towards zero and the mixing ratio drops to one as f_τ approaches f_ε.

Policy Losses

Comparing policies under hidden information against outcomes under full information begs the question what would be the preferred setting. Considering the loss function of the policymaker, it turns out that the *ex-ante expectation* of the policy loss, $E(V_t)$ is improved under hidden information over the wide range of calibrations discussed above. Figure 5 reports how the improvement in policy losses under hidden information are large enough that an average inflation rate corresponding to about one-standard deviation unit would have to be added to inflation under hidden information for policy losses to be equal; the compensating inflation is somewhat smaller when the volatility weight on persistent shocks, ω, is very small; and it can be considerably larger when persistent shocks are very prevalent and the weight on inflation stabilization is large.[22]

---

[22]Additional details are given in Appendix C.

[Figure 5 about here.]

Moreover, the same holds when considering the notion of "representative" loss discussed in Section 2.1. The reason is simply that the improvement in outcomes is due to the policymaker's restraint in pursuing the output targets. By lowering inflation and output gap, this is clearly beneficial for the "representative" loss, which would be minimized by keeping the output gap at zero anyway.

The benefits of reduced inflation also outweigh the *policy* losses from staying away from the targets, at least from an ex-ante perspective considering both persistent and transitory shocks as well as their respective likelihoods.[23] Conditional on the occurence of an iid shock ε_t, inflation is of course higher, and the policymaker misses the target by more than he would under full information, see Figure 2. On average, this is however outweighed by the benefits incurred when a persistent shock occurs.

Considering the different levers present in the policy function (8) under hidden information, a quantitative decomposition of the reduction in policy loss looks as follows: First, there is the optimal policy under symmetric information. Feeding this same policy through the system but under the hidden information, losses drop by an amount, correspond to a compensating rate of average inflation of about one third of the standard deviation of inflation in the new equilibrium. The optimal policy under hidden information then seeks to improve upon this by changing the mixing ratio and by reacting to past beliefs. Using the optimal mixing ratio, but neglecting the response to prior beliefs ($f_b = 0$) makes expected losses drop further; compared to full information the compensating average inflation amounts to about one standard deviation of inflation. In addition to this, the optimal policy reacts also negatively to prior beliefs and the average inflation compensating for the improvement in losses over full information equals almost two standard deviations of inflation. For comparison, the difference in full information losses of discretion and commitment corresponds to a compensating average inflation of about two-and-a-half standard deviations of inflation under discretion.

---

[23]Expected loss, $E(V_t)$, is the unconditional expectation of the policymaker's value function across states of nature. (See the Technical Appendix for computational details.) Optimal policy is of course defined on a state-by-state basis.

3   Belief Shocks

The simple New Keynesian model analyzed so far has only one communication channel between policymaker and public: Policy actions themselves. Since policy is driven by more shocks than there are communication channels, the public cannot perfectly infer the drivers of policy, not even in equilibrium. 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 richer setting, nesting the cases of full and hidden information considered before.

In addition to observing policy, the public is now assumed to receive a noisy signal about the persistent output target. The target signal is contaminated by noise shocks n_t, which will be called "belief shocks". They are iid and the public's measurement vector is

$$Z_t = \begin{bmatrix} x_t & (\tau_t + n_t) \end{bmatrix}' \quad \text{where} \quad n_t \sim N(0, \sigma_n^2)$$

and the notation for public beliefs of a variable z_t is now adapted to

$$z_{t|t} \equiv E(z_t|Z^t)$$

The presence of two correlated observables in the public's inference problem requires to extend the univariate filtering methods discussed in the previous section. Also, the state vector needs to be augmented by n_t. (Notice that $n_{t|t-1} = 0$.) A detailed presentation of handling this and larger settings has been relegated to the technical appendix.

The belief shocks are uncorrelated with fundamentals (here: τ_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 funda-

mental" shocks.[24] Inflation is affected by belief shocks via the forward-looking Phillips Curve (1), making it suboptimal for policy to ignore belief shocks. As with given prior beliefs about the output target ($\tau_{t|t-1}$), they will raise inflation and optimal policy should want to fight their effects by contracting output. In the present model, economic responses to noise shocks will exhibit patterns similar to cost-push shocks, echoing results of Angeletos and La'O (2008b).

By changing the volatility of noise shocks, the extended model also nests the cases of hidden and full information analyzed in the previous section. The scope for hidden information increases with the volatility of belief shocks. For $\sigma_n = 0$, the model is identical to the full information model, since τ_t is perfectly observable. The opposite occurs when σ_n is very large. In this case, the signal becomes useless and the model converges to the hidden information setting from the previous section where policy is the only observable.

[Figure 6 about here.]

Impulse responses to a noise shock are shown in Figures 6, again using the baseline calibration from Table 1. Under this configuration, each of the three shocks in this model occurs with the same probability. When the target signal $\tau_t + n_t$ 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 τ_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 (1). Second, it signals that the target τ_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. Figure 6 also documents that public beliefs of future output and inflation are *both* elevated during the entire episode, which distinguishes belief shock induced dynamics from cost-push behavior, since the latter would typically be accompanied by an *expected* recession as well.

---

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

How does policy change with the volatility of belief shocks? To answer this question, Figure 7 shows how policy coefficients and expected losses change when σ_n is varied between zero and infinity. As discussed above, 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_τ, f_ε and f_b vary smoothly and monotonically between the comparative statics of hidden vs full information studied before. They are all smaller and policy losses are reduced as the extent of hidden information increases with σ_n.

[Figure 7 about here.]

The policy response to noise shocks is always negative. The reasons are similar to what has been discussed in the previous section for the negative response to *prior* beliefs f_b. 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 7 shows how the policy reaction to one-standard deviation shock, $f_n \cdot \sigma_n$ changes with the shock variance. The noise response peaks at an intermediary level of the noise variance, where the public places roughly equal weight on policy and the target signal in its updating of beliefs.

$$\tau_{t|t} = \tau_{t|t-1} + K_x \tilde{x}_t + K_s(\tilde{\tau}_t + n_t) \tag{19}$$

Changes in the Kalman gains K_x and K_s for various noise levels are shown in the bottom left panel of Figure 7. In the extremes the response of policy to noise shocks is zero, as either the size of the shock shrinks to zero (and there are no erroneous beliefs to fight) or the public pays no attention to a signal with infinite noise.

A natural interpretation of variations in noise variance is to view these as changes in 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 τ_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 target is perfectly observable, as is the case above, the current policymaker cannot influence beliefs, since the Kalman gain K_x in (19) is zero.[25] In contrast, when targets are perfectly inferable from observed policies, this link is retained causing the difference in outcomes.

4 (In-)Credible Disinflations and Exogenous Persistence

A pertinent question in monetary policy is whether to conduct disinflations quickly or gradually. The answer involves a minimization of the economic costs incurred by the necessary output contractions along the disinflation path. These costs hinge on the persistence of inflation. If persistence is large, a larger or more protracted contraction might be necessary. A pertinent policy question is then whether to chose the "cold turkey" approach of a quick disinflation, involving a large initial contraction, or whether to chose a more gradual approach, implementing a longer sequence of smaller contractions.

Academic research has offered different advice on these issues, see for example the discussion between Gordon (1982) and Sargent (1982). Arguments for or against either approach differ in whether credibility is assumed to have an effect on inflation persistence or not. A quick disinflation could enhance the credibility of the policymaker's intention to disinflate and help reducing the inflation rate by itself. Taking this view, Sargent (1982) favors the "cold turkey" approach. Being more concerned with exogenous sources of inflation persistence makes Gordon (1982) lean

---

[25]This is similar to what Faust and Svensson, p. 374 call the regime "OG: observable goal and intention", for which they find results corresponding to what has been found in this paper.

towards advocating more gradual disinflation paths.

The framework presented in this paper allows to address these questions in a fully dynamic framework with an optimizing policymaker. The linear quadratic approach allows to handle multiple, endogenous state variables, including those arising from partially backward-looking inflation dynamics. To the best of my knowledge, this is the first explicitly optimizing analysis of disinflation strategies when policy goals are unobserved.[26]

The following disinflation experiment is considered: How should policy react to a surge in inflation due to unfounded public beliefs? In an admittedly stylized way, this resembles the initial conditions of the Volcker disinflation as discussed by Erceg and Levin (2003) and Goodfriend and King (2005). Such inflation beliefs may be caused by a belief shock, n_t, or inherited via $\pi_{t|t-1} = g^0 \tau_{t|t-1}$. Qualitatively, results are similar in either case and the discussion below will focus on responses to belief shocks.

In the model of the previous sections, the cost of disinflation depends largely on the policymaker's capability to lower policy expectations quickly. In the belief shock model of the previous section, policy induces a stronger contraction of the economy in response to beliefs $\tau_{t|t-1}$ when credibility problems are larger, see the middle right panel of Figure 7.[27] While this suggests that disinflations should be more aggressive, it does not yet speak to concerns about the trade-offs under exogenous inflation persistence.

To see how policy changes in the presence of exogenous persistence, the Phillips Curve is

---

[26]Related is the work of Ireland (1995) who finds similar results when imposing a sluggish response of public beliefs on policy announcements. As an alternative, Ireland (1997) seeks to reconcile the conjectures of Sargent (1982) and Gordon (1982) by differentiating between disinflations at high or low levels of inflation.

[27]Similarly, the negative response to a belief shock is stronger for larger values of σ_n, up to the point where the growing noise variance beliefs react less and less to these shocks as .

augmented with a backward-looking term, representing price indexation at the rate γ.[28]

$$\pi_t = \frac{\gamma}{1+\beta\gamma}\pi_{t-1} + \frac{\beta}{1+\beta\gamma}\pi_{t+1|t} + \frac{\kappa}{1+\beta\gamma}x_t \tag{20}$$

$$= \gamma\pi_{t-1} + \kappa\sum_{k=0}^{\infty}\beta^k x_{t+k|t}$$

In this hybrid Phillips Curve (20), inflation is not only determined by the expected path of future policies known from (1), but also by lagged inflation. Policy innovations are still the ultimate driver of inflation, but they carry less weight in changing current inflation.

[Figure 8 about here.]

Figure 8 compares impulse responses to a belief shock n_t when varying the indexation rate $\gamma \in \{0; 0.5; 1\}$; for $\gamma = 0$ the model is identical to what has been studied above.[29] For better comparison of the disinflation policies, the belief shocks have been scaled such as to yield a unit innovation in inflation on impact. As higher indexation rates increase exogenous inflation persistence, optimal policy contracts the economy ever more aggressively to a belief shock — bolstering the case for the "cold turkey" approach. The lower panel of Figure 8 confirms this also over a wider range of values for the noise variance σ_n^2.

Even though policy contracts the economy more vigorously when exogenous persistence is larger, disinflations are not necessarily quicker. Due to the higher exogenous persistence in inflation it takes longer for inflation to fall when γ is larger. Policy cannot avoid the higher degree of backward-lookingness in inflation. But this is precisely why an aggressive initial contraction is warranted. It does not only fight beliefs as in the previous section. By reducing current inflation, it reduces also the amount of future inflation caused by ill-founded beliefs to be carried forward via the backward-looking term in the Phillips Curve.

---

[28] As in Woodford (2003a) or Christiano, Eichenbaum, and Evans (2005) this can be derived from the optimizing behavior of Firms under Calvo pricing. Firms who do not optimize their prices are supposed to change prices at the rate Π_{t-1}^γ where Π_{t-1} is last period's level (not log) of inflation. As shown by Woodford (2003a), this changes also welfare functions such as (2) to be concerned with quasi-differenced inflation $\pi_t - \gamma\pi_{t-1}$ instead of inflation. The point of the experiment is here is however to consider how exogenous persistence changes policies whilst keeping the objective function constant. The policymaker's loss function is thus kept unchanged.

[29] Other parameters are calibrated at the values shown in Table 1.

5 Related Literature

Since asymmetric information is such a pertinent issue in policymaking, it is no wonder, that there is a wide body of related literature. General surveys can be found in Rogoff (1989), Walsh (2003, Chapter 8) and Persson and Tabellini (2000, Chapter 15). The literature can roughly be classified by answering the following questions: *Who* learns about *what* and *how*? *How* is policy described, as an explicit optimization problems or by a behavioral policy rule? In this paper, policy is optimized while the public solves a signal extraction problem about hidden policy targets.

The tractability of the solution method presented here stems from the unobservable states following smooth, Gaussian processes as opposed to regime switches. Discrete regime switches are attractive 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, Lu, and Pasten (2008). Unobserved regime switches lead to important non-linearities in the public's inference problem, which complicate the constraints in an optimal policy problem considerably. The aforementioned literature has correspondingly focused on very small state spaces and/or finite horizon problems, since unobserved regimes switches are hard to incorporate into the kind of general dynamic settings commonly used for policy analysis. Svensson and Williams (2006) discuss the resulting difficulties in more detail.

Learning about regime-switches does not pose such problem when it is the central bank who learns about economic conditions as in Sargent (1999). This is because of the strategic behavior of the policymaker when facing agents learning about, respectively from, him as opposed to the non-strategic behavior of atomistic private agents.

Svensson and Woodford (2003, 2004), Aoki (2006) study optimal policy with an imperfectly informed central bank in linear quadratic settings similar to mine. A convenient feature of this approach is that the public's "learning" reduces to a time-invariant signal extraction problem. This is different from the kind of evolutionary belief system studied in the learning literature represented for example by Evans and Honkapohja (2001). Adaptive learning leads to interesting dynamics where past data drives changes in regression coefficients, but is so far hard to capture in optimal

policy problem. For fixed policy rules, the issue is analyzed by Orphanides and Williams (2005, 2006). An exception is the work of Gaspar, Smets, and Vestin (2006) who endow the public with a time-varying, adaptive learning rule. They derive a Markov-perfect policy with history dependence induced similarly as here via the reaction to people's beliefs. Their policies generate data with low inflation persistence to influence people's constant gain learning. Thanks to the lower complexity of the inference problem adopted here, their results can be corroborated in a very transparent way.

Hidden information is modeled here as a signal extraction problem where the private sector does not observe the realization of shocks, but where the structure of the economy and its parameter values are mutually known. In a rational expectations equilibrium, the private sector then knows the correct policy function but can only imperfectly infer the nature of shocks. This equilibrium notion is stronger than the self-confirming equilibria considered by Fudenberg and Levine (1993) and Sargent (1999) or the recursive learning schemes studied for example by Evans and Honkapohja (2001) or Orphanides and Williams (2005). In those cases, the public beliefs about structural relations may be erroneous as long as they are justified by the data generated from the model. In the rational expectations equilibrium pursued here, the public knows the true policy function — but not the states driving it. This serves as a useful, non-trivial benchmark for evaluating the consequences of a superiorly informed policymaker in dynamic economies.

Closest to the simple model studied in Section 2 are the studies by Cukierman and Meltzer (1986) and Faust and Svensson (2001, 2002). This paper shares with them not only the linear framework and the Kalman filtering of the public, but also that it casts the policy problem around unobserved policy goals. New is the general framework capable of handling various models with endogenous state variables. Faust and Svensson focus on the welfare effects of credibility. Within a slightly different economic structure (forward-looking Phillips Curve instead of Lucas-supply curve),[30] their results are broadly confirmed here: Outcomes are improved when output targets are unobservable. (See also the discussion on varying transparency in Section 2.4.) The common force at work is that the updating of public beliefs depends directly on observed policy. Similar

---

[30]A further difference is that their analogue to the iid shock ε_t is not a target component but a control error of policy.

to mechanisms discussed by Walsh (2000), this disciplines policy while retaining Markov-perfect time-consistency.

6 Conclusions

When a policymaker is better informed than the public, public beliefs about his hidden information become a distinct 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. Since public beliefs are a natural state variable under imperfect information, managing this state of beliefs is Markov-perfect and time consistent.[31]

This paper solves for the optimal discretion policy in a New Keynesian model with unknown output targets and finds that policy contracts the economy in response to inflationary beliefs. In addition, the pursuit of output targets is scaled back, because of their inflationary effects on public beliefs. This policy, in particular its history dependence, shares some similarities with commitment policies. To the extent that hidden information is a realistic feature of actual policymaking, this suggests much smaller costs for real-world policymakers from retaining some degree of discretion as long as they keep public beliefs about their intentions in check.

The stylized model analyzed here implies that intransparency of policy targets is preferable — at least under discretion. However, an important caveat is that public beliefs matter here only for linking economic activity to pricing decisions in the New Keynesian Phillips Curve. The welfare effects of transparency might be subject trade-offs, leading to more differentiated results, when expectations of future activity were to affect both investment and pricing decisions.

The model gives also rise to belief shocks as a source of business cycle fluctuations. Similar to the work of Lorenzoni (2006), such shocks shift public perceptions about economics fundamentals, whilst the actual fundamentals remain unchanged. Under imperfect information, these shifts in public beliefs are rational since the belief shocks are correlated with informative signals about

---

[31]This excludes the explicit reputational mechanisms based on history dependent strategies known from by Barro and Gordon (1983b) or Chari and Kehoe (1990) are excluded from the analysis.

fundamentals. The optimal discretion policy seeks to quell the erroneous beliefs arising from these shocks. In the New Keynesian model studied here, belief shocks induce dynamics similar to cost push shocks, which is similar to belief shock dynamics found by Angeletos and La'O (2008b) based on higher-order dynamics.

Apart from illustrating the optimal policy under imperfect information within a widely studied New Keynesian model, the technical appendix to this paper provides a general solution method which allows to extend the analysis to larger settings, relevant for practical policy analysis. By relying on linear quadratic approximations and Gaussian uncertainty, the optimal policy problem becomes tractable without losing its economic intricacy.

A fruitful area for future research would be to extend the analysis to a policymaker's proprietary information about economic fundamentals, for example determinants of potential output. Such information would be widely dispersed amongst different members of the public as well as the central bank, suggesting to combine the policy concepts studied here with the dispersed information settings studied for example by Lorenzoni (2006) or Angeletos and La'O (2008a).

Appendix

A   Innovation Representation in the Simple Model

This section derives the ARMA(1,1) innovations process (9) for policy in the simple model of Section 2. First, policy can be separated into innovation and public expectations

$$x_t = \underbrace{f_\tau \tilde{\tau}_t + f_\varepsilon \tilde{\varepsilon}_t}_{\tilde{x}_t} + \underbrace{(f_\tau + f_b)\tau_{t|t-1}}_{x_{t|t-1}} \tag{21}$$

The ARMA representation follows from using $\tau_{t+1|t} = \rho\, \tau_{t|t}$ and the Kalman filter to express the evolution of prior beliefs as

$$x_{t+1|t} = \rho x_t + \rho\left(\left(f_\tau + f_b\right) K_\tau - 1\right)\tilde{x}_t$$

B   Kalman Filter in the Simple Model

In the simple model of Section 2, the signal extraction problem of the public uses the policy function (8) as measurement equation and the AR(1) transition of τ_t (6) as state equation. The Kalman gains are

$$K_\tau = \frac{1}{f_\tau}\frac{\Sigma_\tau}{\Sigma_\tau + \bar{\sigma}^2} \quad \text{and} \quad K_\varepsilon = \frac{1}{f_\varepsilon}\frac{\bar{\sigma}^2}{\Sigma_\tau + \bar{\sigma}^2}$$

where $\bar{\sigma}^2 \equiv f_\varepsilon^2 / f_\tau^2 \cdot \sigma_\varepsilon^2$ and Σ_τ solves the Riccati equation

$$\Sigma_\tau = \sigma_\eta^2 + \rho^2\,\bar{\sigma}^2\,\frac{\Sigma_\tau}{\Sigma_\tau + \bar{\sigma}^2} = \frac{\sigma_\eta^2}{1 - \rho^2\,\frac{\bar{\sigma}^2}{\Sigma_\tau + \bar{\sigma}^2}}$$

As discussed on Section 2.3, it is plausible to assume that $0 \le f_\tau \le 1$, $0 \le f_\varepsilon \le 1$, $f_b \le 0$, and $f_\tau + f_b \ge 0$. It is then straightforward to verify that $-1 \le \psi \le 0$, since the Kalman gains are positive and $f_\tau K_\tau + f_\varepsilon K_\varepsilon = 1$.

C Compensating Rate of Average Inflation

Section 2.4 evaluates ex-ante policy losses based on the policymaker's objective function

$$E(V_t) = \frac{E\left(\pi_t^2 + \alpha_x(x_t - \bar{x}_t)^2\right)}{1 - \beta}$$

The difference in losses under hidden versus full information, can be expressed as a compensating rate of average inflation, $\bar{\pi}$, which would equalize policy losses in both equilibria when added to the dynamics under hidden information.

$$E((1-\beta)V_t|\text{Full Info}) = E\left((\pi_t + \bar{\pi})^2 + \alpha_x(x_t - \bar{x}_t)^2|\text{Hidden Info}\right)$$

As an alternative measure of policy losses, this "compensating average inflation" abstracts from the validity of the linear quadratic framework for non-zero (or non-indexed) inflation rates in steady state.

Since this paper looks only at shocks to output targets, the calibration does not try to match a level of second moments observed in the data and the compensating rate of average inflation is scaled by the standard deviation of inflation under hidden information.

References

Abrams, Burton A. 2006. "How Richard Nixon Pressured Arthur Burns: Evidence from the Nixon Tapes." *Journal of Economic Perspectives* 20 (4): 177–188 (Fall).

Angeletos, George-Marios, and Jennifer La'O. 2008a, October. "Dispersed Information over the Business Cycle: Optimal Fiscal and Monetary Policy." *mimeo*, MIT.

———. 2008b, December. "Incomplete Information, Higher-Order Beliefs and Price Inertia." *mimeo*, MIT.

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

———. 2006. "Optimal commitment policy under noisy information." *Journal of Economic Dynamics and Control* 30 (1): 81–109 (January).

Backus, David, and John Driffill. 1985a. "Inflation and Reputation." *The American Economic Review* 75 (3): 530–538 (June).

———. 1985b. "Rational Expectations and Policy Credibility Following a Change in Regime." *The Review of Economic Studies* 52 (2): 211–221 (April).

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

Barro, Robert J., and David B. Gordon. 1983a. "A Positive Theory of Monetary Policy in a Natural Rate Model." *The Journal of Political Economy* 91 (4): 589–610 (August).

———. 1983b. "Rules, Discretion and Reputation in a Model of Monetary Policy." *Journal of Monetary Economics* 12 (July): 101–121.

Blinder, Alan S. 1998. *Central Banking in Theory and Practice*. 1st MIT Press paperback edition, 1999. The Lionel Robbins Lectures. Cambridge, MA: The MIT Press.

Bordo, Michael D., Christopher Erceg, Andrew Levin, and Ryan Michaels. 2007, March. "Three Great American Disinflations." NBER Working Papers 12982, National Bureau of Economic Research, Inc.

Calvo, G.A. 1983. "Staggerd Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–398.

Canzoneri, Matthew B. 1985. "Monetary Policy Games and the Role of Private Information." *The American Economic Review* 75 (5): 1056–1070 (December).

Chari, V. V., and Patrick J. Kehoe. 1990. "Sustainable Plans." *The Journal of Political Economy* 98 (4): 783–802 (August).

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1): 1–45 (February).

Clarida, Richard, Jordi Gali, and Mark Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37 (December): 1661–1707.

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

Cukierman, Alex, and Allan H. Meltzer. 1986. "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information." *Econometrica* 54 (5): 1099–1128 (September).

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

Evans, George W., and Seppo Honkapohja. 2001. *Learning and Expectations in Marcoeconomics*. Princeton, NJ: Princeton University Press.

Faust, Jon, and Lars E O Svensson. 2001. "Transparency and Credibility: Monetary Policy with Unobservable Goals." *International Economic Review* 42 (2): 369–97 (May).

———. 2002. "The Equilibrium Degree of Transparency and Control in Monetary Policy." *Journal of Money, Credit and Banking* 34 (2): 520–39 (May).

Fudenberg, Drew, and David K. Levine. 1993. "Self-Confirming Equilibrium." *Econometrica* 61 (3): 523–545 (May).

Gali, Jordi. 2003. "New Perspectives on Monetary Policy, Inflation, and the Business Cycle." Chapter 6 of *Advances in Economics and Econometrics: Theory and Applications*, edited by Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, Eighth World Congress of the Econometric Society, 151–197. Cambridge, UK: Cambridge University Press.

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

Gaspar, Vitor, Frank Smets, and David Vestin. 2006. "Adaptive Learning, Persistence, and Optimal Monetary Policy." *Journal of the European Economic Association* 4 (2-3): 376–385 (04-05).

Goodfriend, Marvin, and Robert G. King. 2005. "The Incredible Volcker Disinflation." *Journal of Monetary Economics* 52 (5): 981–1015 (July).

Gordon, Robert J. 1982. In Hall 1982, Chapter 2, 11 – 40. National Bureau of Economic Research Project Report.

Hall, Robert E., ed. 1982. *Inflation: Causes and Effects*. Chicago: The University of Chicago Press. National Bureau of Economic Research Project Report.

Ireland, Peter N. 1995. "Optimal disinflationary paths." *Journal of Economic Dynamics and Control* 19 (8): 1429–1448 (November).

Ireland, Peter N. 1997. "Stopping Inflations, Big and Small." *Journal of Money, Credit and Banking* 29 (4): 759–775 (November).

King, Robert G., and Marvin Goodfriend. 1997. "The New Neoclassical Synthesis and the Role of Monetary Policy." In *NBER Macroeconomics Annual 1997*, edited by Ben S. Bernanke and Julio J. Rotemberg, 231–283. Cambridge (MA): The MIT Press.

King, Robert G., Yang K. Lu, and Ernesto S. Pasten. 2008. "Managing Expectations." *Journal of Money, Credit and Banking* 40 (8): 1625–1666.

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

Kydland, Finn E., and Edward C. Prescott. 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *The Journal of Political Economy* 85 (3): 473–492 (June).

Lorenzoni, Guido. 2006, August. "A Theory of Demand Shocks." NBER Working Papers 12477, National Bureau of Economic Research, Inc. *forthcoming in the American Economic Review*.

Morris, Stephen, and Hyun Song Shin. 2002. "Social Value of Public Information." *American Economic Review* 92 (5): 1521–1534 (December).

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

———. 2006. "Monetary Policy with Imperfect Knowledge." *Journal of the European Economic Association* 4 (2-3): 366–375 (04-05).

Persson, Torsten, and Guido Tabellini. 2000. *Political Economics*. Zeuthen lecture book series. Cambridge, MA: The MIT Press.

Rogoff, Kenneth. 1989. "Reputation, Coordination and Monetary Policy." Chapter 6 of *Modern Business Cycle Theory*, edited by Robert J. Barro. Cambridge, MA: Harvard University Press.

Sargent, Thomas J. 1982. In Hall 1982, Chapter 3, 41 – 97. National Bureau of Economic Research Project Report.

———. 1999. *The Conquest of American Inflation*. Princeton (NJ): Princeton University Press.

Svensson, Lars E. O. 1997. "Optimal Inflation Targets, "Conservative" Central Banks, and Linear Inflation Contracts." *The American Economic Review* 87 (1): 98–114.

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

———. 2004. "Indicator variables for optimal policy under asymmetric information." *Journal of Economic Dynamics and Control* 28 (4): 661–690 (January).

Svensson, Lars E.O., and Noah Williams. 2006, November. Bayesian and Adaptive Optimal Policy under Model Uncertainty. *mimeo* Princeton University.

Vestin, David. 2006. "Price-level versus inflation targeting." *Journal of Monetary Economics* 53 (7): 1361 – 1376.

Walsh, Carl E. 2000. "Market Discipline and Monetary Policy." *Oxford Economic Papers* 52 (2): 249–271 (April).

———. 2003. *Monetary Theory and Policy*. 2nd edition. Cambridge, MA: The MIT Press.

Woodford, Michael. 2003a. *Interest and Prices*. Princeton, NJ: Princeton University Press.

———. 2003b. "Optimal Interest-Rate Smoothing." *Review of Economic Studies* 70 (4): 861–886 (October).

———. 2005, August. "Central bank communication and policy effectiveness." *The Greenspan Era: Lessons for the Future*. FRB Kansas City Symposium, Jackson Hole, Wyoming, 399–474.

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

Figure 1: Policy Trade-Offs in the Simple Model
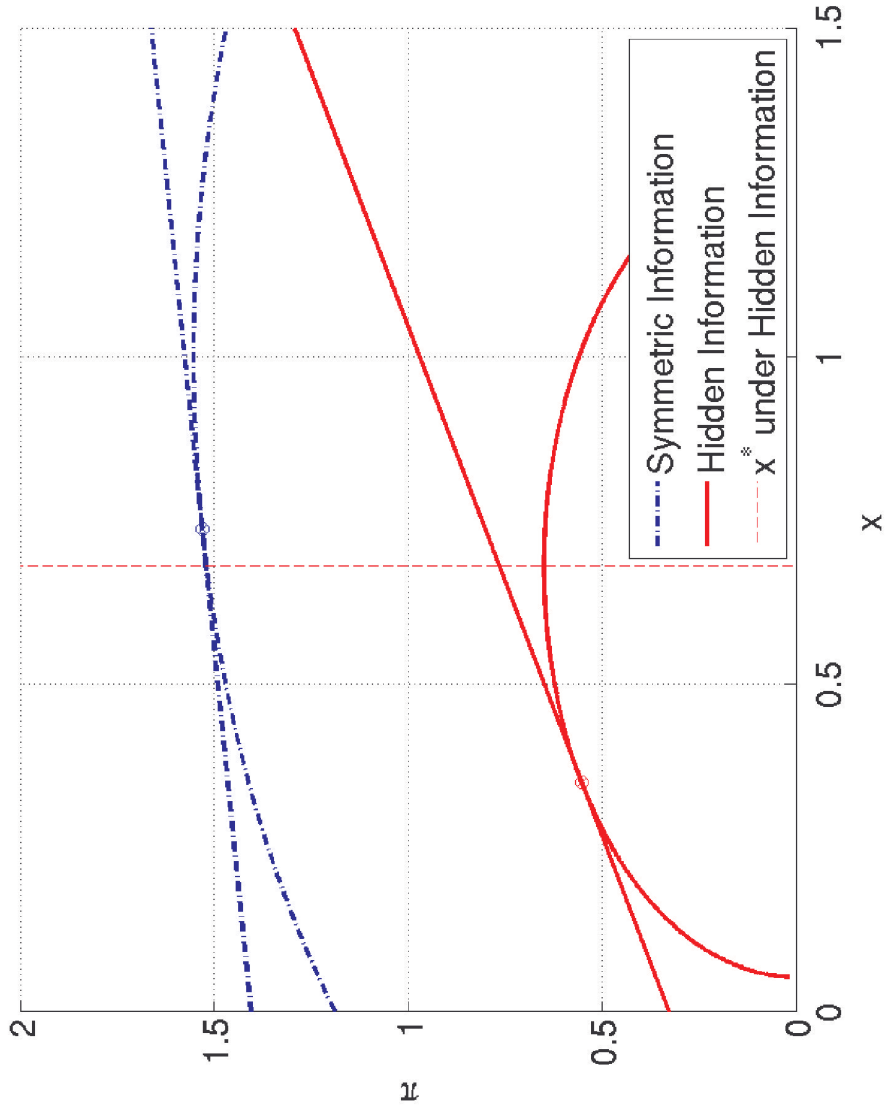
Note: Phillips Curve and indifference curves for policy problem in simple New Keynesian model. Computed for the case of a publicly anticipated, persistent target level: $\tau_t = 1$, $\tau_{t|t-1} = 1$ and $\varepsilon_t = 0$. Dashed lines show symmetric information case with $\pi_{t+1|t} = \bar{g}\rho\tau_t$ and $\bar{g}$ computed from the equilibrium policy (4). Solid lines depict Phillips Curve (17) and indifference curves (18) under hidden information. Optima occur at the circled tangency points.

Figure 2: Optimal Policy in the Simple Model
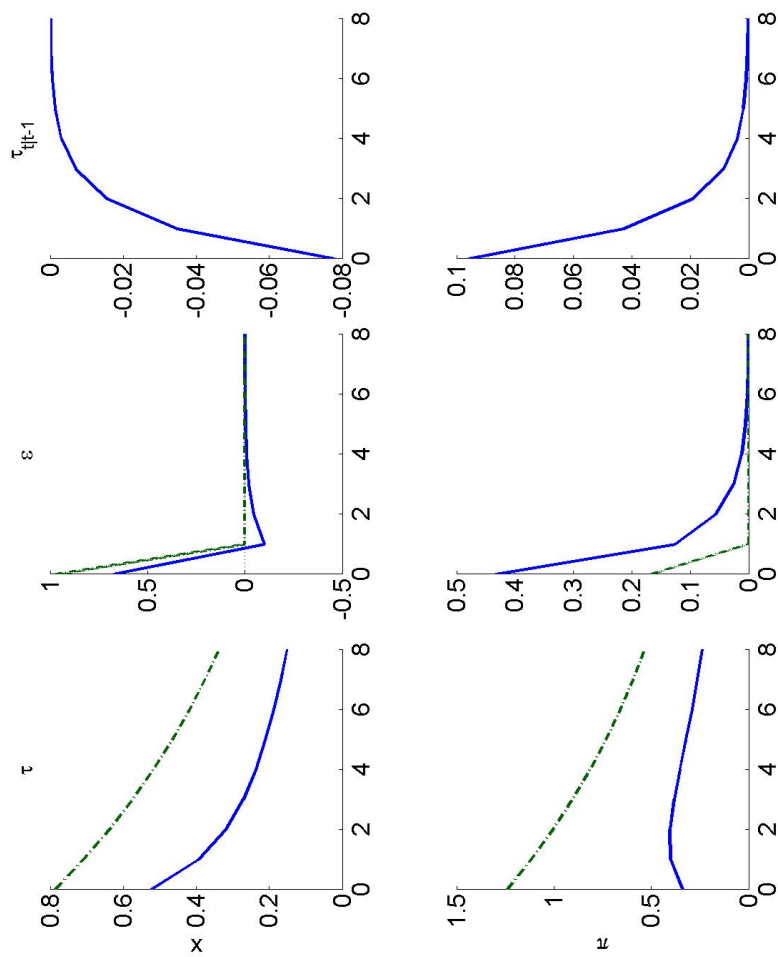
Note: Impulse responses of output gap (x_t) and inflation (π_t) under hidden (straight lines) and full information (dashed) to unit standard deviation shocks τ_t and ε_t. The column labeled $\tau_{t|t-1}$ shows responses to initial conditions $\tau_t = 0$, $\varepsilon_t = 0$, $\tau_{t|t-1} = 1$.

Figure 3: Innovation Responses in the Simple Model
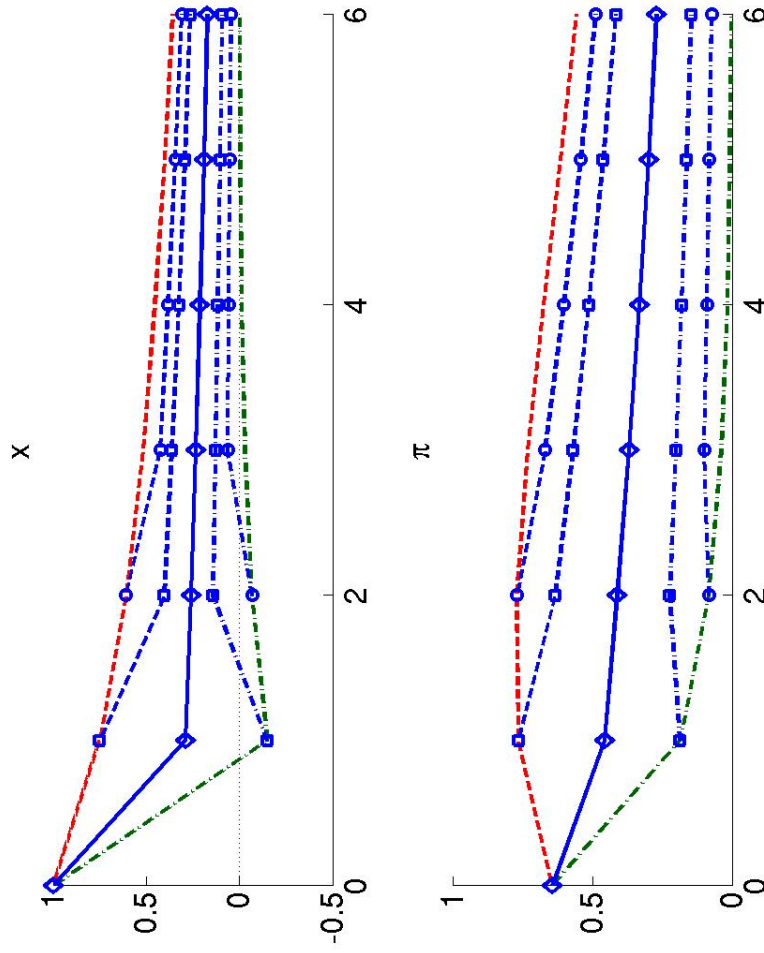
Note: This picture documents the updating of public beliefs about the evolution of output gap and inflation after a unit innovation in policy ($\tilde{x}_t$) at time zero. Dashed lines are true impulses under the full information measure to structural shocks τ_t (top line on each graph, red) respectively ε_t (bottom, green). These structural responses are scaled to yield a unit innovation in output. Dashed blue lines document update impulse responses formed by the public at times 1 (with square marker) respectively time 2 (circles). Top panel shows the output gap (x_t), inflation (π_t) is shown in the lower panel.

Figure 4: Sensitivity Analysis

(a) f_τ

(b) f_ε

(c) Belief Reaction f_b

(d) Mixing Ratio f_ε / f_τ
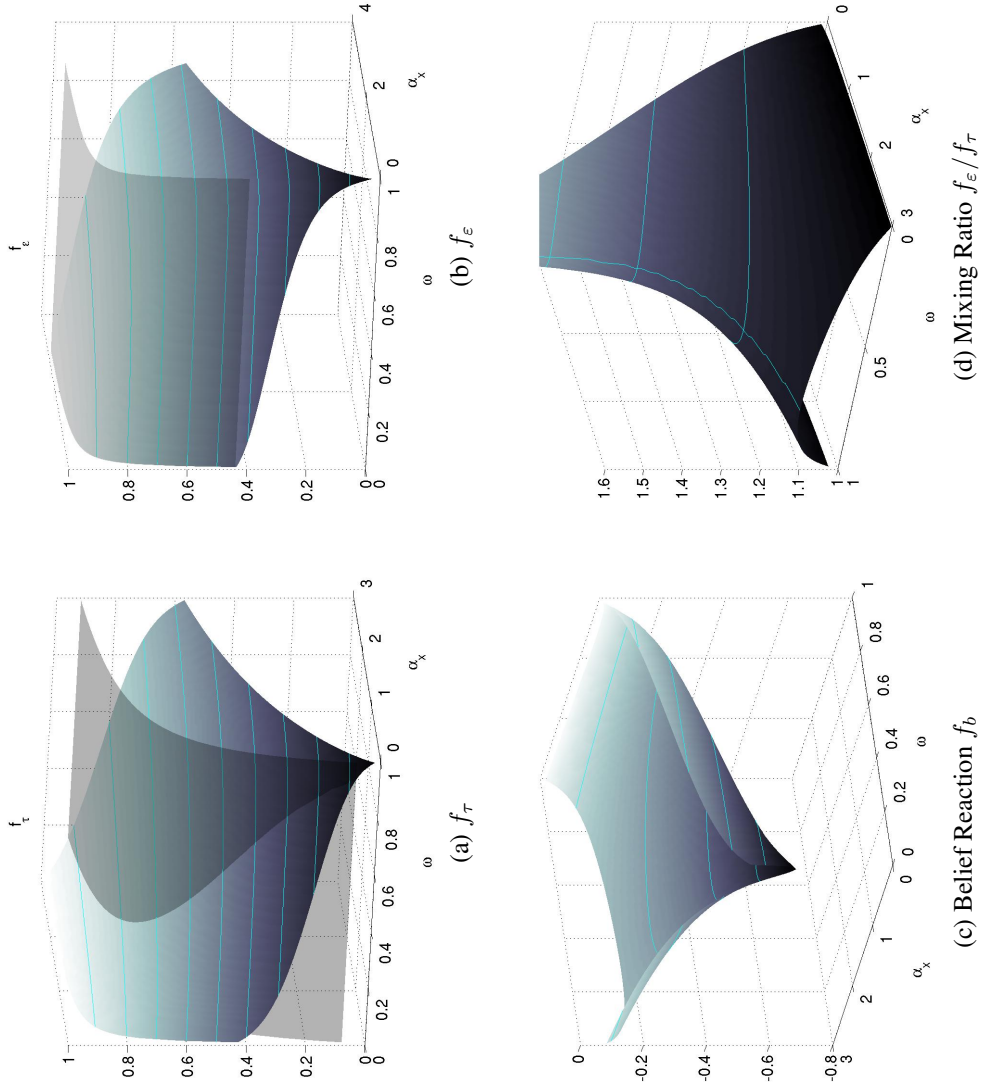
Note: Sensitivity of policy function due to variations in preference weight α_x and variance weight ω in the simple model of Section 2, keeping other parameters at their baseline values from Table 1. In order to provide the best perspective on each surface, axes are rotated differently in each panel. Transparent surfaces in Panels (a) and (b) depict corresponding variations in f_τ, respectively f_ε in the symmetric information model.

44

Figure 5: Improved Policy Losses under Hidden Information

(a) Compensating average inflation

(b) Alternative Hidden Info Policies
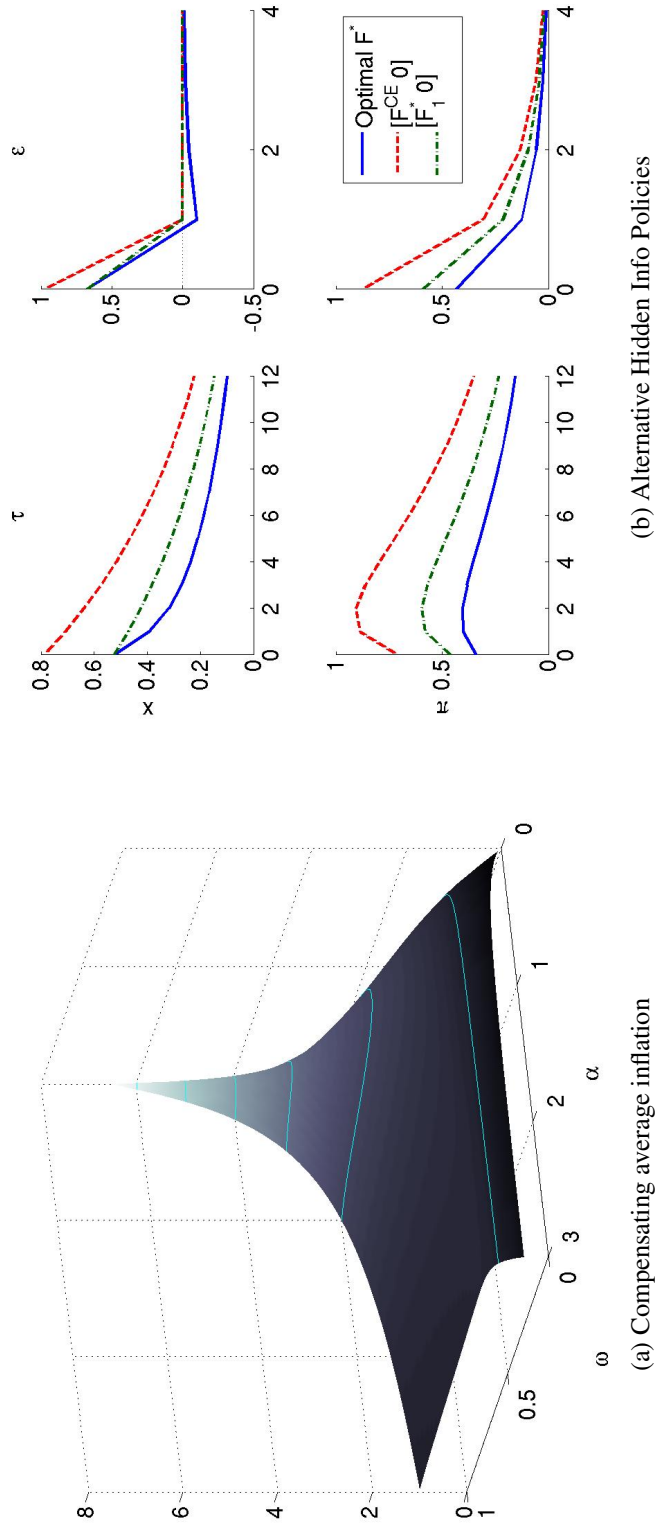
Note: Panel (a) shows differences in expected policy losses between full and hidden information settings. A positive number means that policy losses are larger under full information, implying that hidden information outcomes are preferred. The difference in losses is expressed as a compensating increase in average inflation under hidden information, scaled in standard deviation units. Baseline calibration as shown in Table 1 wth variations in prefernce weight α_x and variance weight ω as indicated in the figure. Panel (b) shows impulse responses of output and inflation for different policies under under hidden information with. Straight lines show the optimal policy. Dash-dotted lines use the optimal coefficients f_ε and f_τ but neglect prior beliefs ($f_b = 0$). Dashed lines correspond to using the symmetric information policy when there is hidden information. Compared with discretion under full information, the compensating rates of average inflation correspond to 0.37, 1.11 and 1.82 units of the standard deviation of inflation in each equilibrium.

45

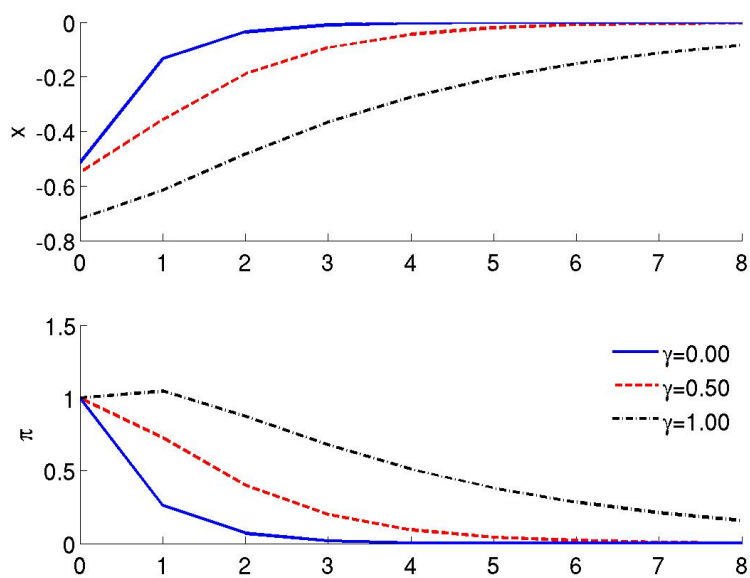Figure 6: Impulse Responses with Belief Shocks
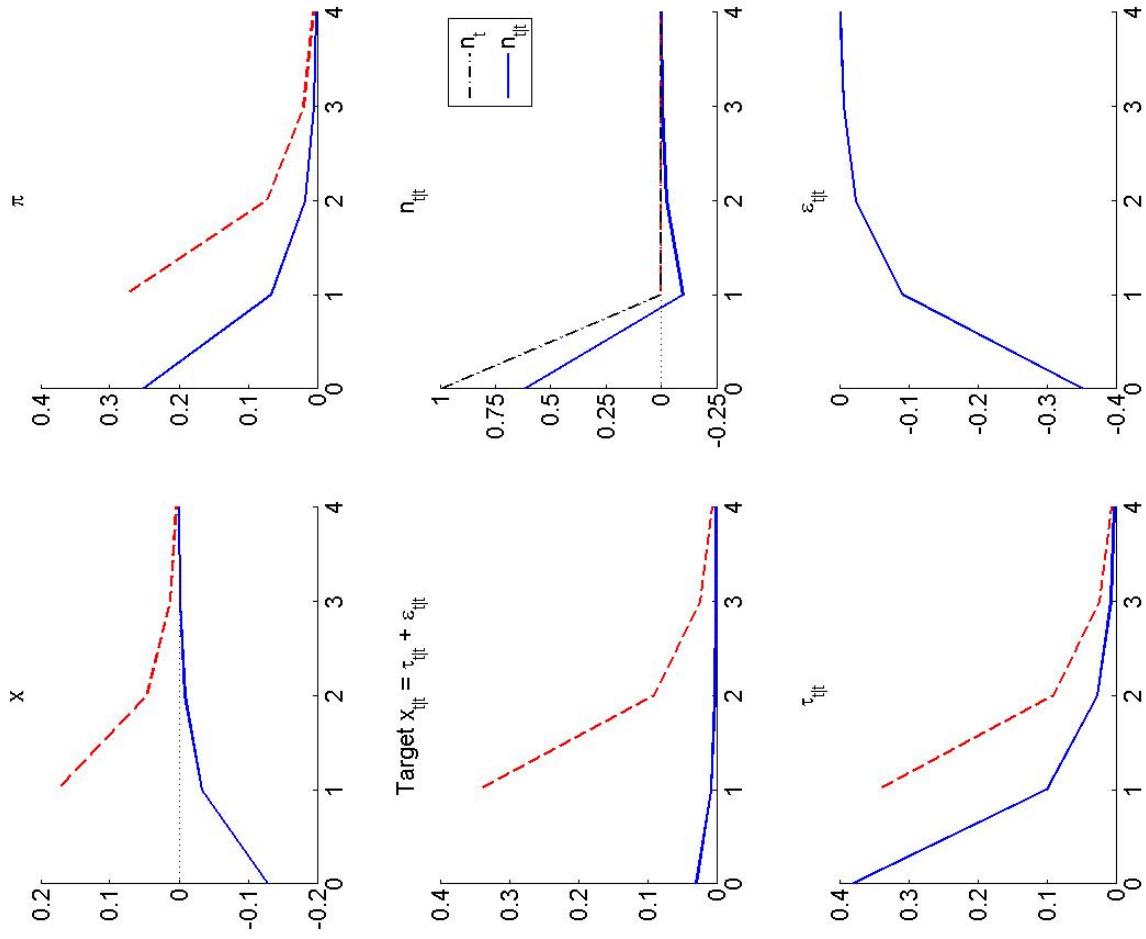
Note: Straight lines are impulse responses to a unit belief shocks n_t. For any variable z_t, dashed lines (red) depict the evolution of prior beliefs $z_{t|t-1}$. Calibration values from Table 1. Dash-dotted line in middle right panel, "n_t", depicts impulse response of the belief shock.

Figure 7: Varying Transparency
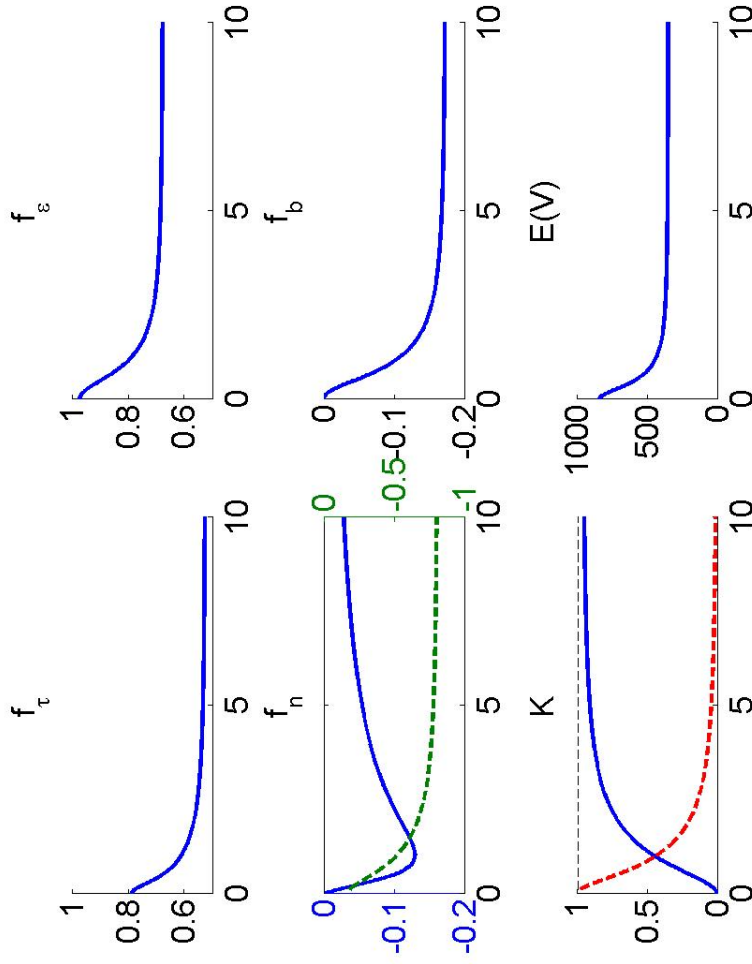
Note: Change in policy coefficients, expected policy loss and Kalman gains while varying as the volatility of belief shocks, σ_n, is varied along the x-axis. Computations are based on the belief shock model of Section 3. The middle left panel shows scaled response coefficients, f_n, to a belief shock, since the size of belief is changed on the x-axis. The straight line shows responses to a belief shock with unit-standard deviation, $f_n \cdot \sigma_n$. The dashed line scales the coefficient such that belief shock generates a unit increase in inflation on impact, f_n/π_0 (with units indicated on the right hand axis). (There is no inflation response to a belief shock when, $\sigma_n = 0$, and the plot shows no data for this point.) The bottom left panel shows Kalman gains K_x (solid line) and K_s (dashed line) from the public's updating equation (19). (Notice: The solid line, K_x, converges to 0.94 not one as the variance of belief shocks grows to infinity.)

47

Figure 8: Disinflation with Exogenous Persistence

(a) Impulse response
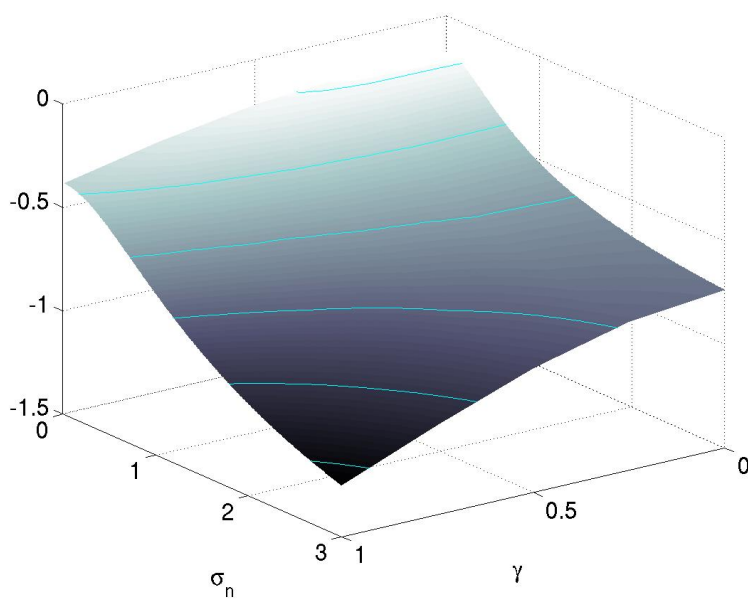
(b) Impact Coefficients of x_t

Note: Panel (a): Impulse responses of output and inflation to a belief shock with varying degrees of price indexation ($\gamma = \{0; 0.5; 1\}$) in the Phillips Curve (20). Belief shocks are normalized such that they produce a unit response in inflation on impact. Panel (b) reports the impact coefficients of output in response to such normalized belief shocks for different values of the indexation rate γ and noise level σ_n^2. (Other parameters calibrated as in Table 1.)

Table 1: Model Calibration

		Private Sector Parameters
β	0.99	Time preference
σ	1.00	Risk Aversion / Inverse EIS
θ	0.75	Calvo Probability of not repricing
ϕ	1.00	Inverse Frisch Labor Elasticity
κ	0.1717	Slope of Phillips Curve: $\kappa = (1 - \theta) \cdot (1 - \beta \cdot \theta)/\theta \cdot (\sigma + \phi)$
		Policy Preferences
α_x	1.00	Policymaker's preference for output stabilization $L_t + \pi_t^2 + \alpha_x(x_t + \bar{x}_t)^2$
		Driving Processes
σ_ε	1.00	Volatility of iid target component, $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$ from $\bar{x}_t = \tau_t + \varepsilon_t$
ρ	0.90	Persistence of target component $\tau_{t+1} = \rho\tau_t + \eta_{t+1}$
σ_η	1.00	$\eta_t \sim N(0, \sigma_\eta^2)$
$\sigma_{\bar{x}}$	1.00	Total volatility of output target used for sensitivity analysis. $\sigma_{\bar{x}} \equiv \sigma_\varepsilon^2 + \sigma_\tau^2/(1 - \rho^2)$
σ_n	1.00	Volatility of belief shocks $n_t \sim N(0, \sigma_n^2)$ (Section 3)

Notes: Private-sector parameters taken from Gali (2003)'s calibration to quarterly U.S. data. Innovation variances are each normalized to unity and not intended to match the scale of any second moments. The sensitivity of results to ρ and α_x is discussed in Section 2.4. As shown there, variations in α_x are isomorphic to varying κ. As a measure of credibility, σ_n is varied in Section 3.

MANAGING BELIEFS ABOUT MONETARY POLICY UNDER DISCRETION

TECHNICAL APPENDIX

April 2009

A   A Class of Linear Quadratic Models

The mechanisms of the hidden information policy problem extend beyond the simple model of the previous sections and are applicable to a general class of linear quadratic environments. It is thus beneficial to cast the exposition around this more general class of models. The applications presented in Section 3 (noisy signals, backward looking inflation) have already relied on this general framework.

Again, attention is limited to a Markov perfect, discretionary policy problem. In the spirit of "bygones are bygones", state variables in a Markov-perfect equilibrium must be relevant for current payoffs. The public's prior beliefs are part of these Markov states since they matter for public payoffs. A current decision-maker can influence future decisions only by manipulating the state of beliefs as well as other endogenous state variables, for example capital, carried forward into future decision problems. There is no commitment to future policies.

In general, 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 well tractable: The model is cast in a Gaussian framework with constant variances. Tracking entire distributions then collapses to tracking only their means and can be handled with the Kalman filter. It is the public's prior, not posterior, beliefs which enter the state vector, since the latter will be formed after observing current data which is influenced by current policy.

This section defines a rational expectations equilibrium where 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 that he finds it ex-post optimal not to deviate from the policy function assumed in people's Kalman filter.

There are four types of variables: 1) Backward-looking variables, X_t, corresponding for example to the policy targets τ_t and ε_t in the model of the previous sections. 2) Policy controls, U_t, for example like the output gap above. 3) Publicly observable variables, Z_t, coinciding with the output gap in the simple model. 4) Forward-looking decision variables of the private sector, Y_t, like

inflation and the interest rate in the above model.[1] They will be treated as vectors of dimensions N_x, N_u, N_z, and N_y respectively.

The backward looking variables need not only capture exogenous forcing variables like τ_t and ε_t but also endogenous states like capital, habits, or lagged variables, for example inflation in a model with price indexation. They 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 $Ew_tw_t' = I$.[2]

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:

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

$$Z^t = \{Z_t, Z_{t-1}, Z_{t-2}, \ldots\}$$

The history Z^t spans the public information set.[3] A sufficient condition to ensure superior information of the policymaker is that $N_z < N_w$. For any variable z_t,

$$z_{t|t} \equiv E(z_t|Z^t)$$

denotes the expectation of z_t on the private sector's information set. Synonymously these expectations will be called public beliefs. In particular, $X_{t|t-1}$ are the prior beliefs about X_t before observing Z_t. By construction, $Y_t = Y_{t|t}$ always holds since public decisions are based on public information. In principle, Y_t could also be added to the measurement vector, but without adding

---

[1] Except for such simple models, the interest rate is typically modeled as the policy control and the output gap is a forward-looking variable of the private sector.

[2] Without loss of generality, X_t is constructed such that $N_x \geq N_w$.

[3] In addition, there is no uncertainty about the structure of the economy 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).

new information.

The optimality conditions of private sector behavior are represented by an expectational linear difference equation involving only publicly observable variables and public sector expectations:[4]

$$A^1_{yy}Y_{t+1|t} = A_{yy}Y_{t|t} + A_{yx}X_{t|t} + B_yU_{t|t} \tag{3}$$

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

$$V_t = E_t\sum_{k=0}^{\infty}\beta^k L_{t+k} \tag{4}$$

$$L_t = \begin{bmatrix} X_t \\ Y_t \\ U_t \end{bmatrix}' Q \begin{bmatrix} X_t \\ Y_t \\ U_t \end{bmatrix} \tag{5}$$

where the per period loss function L_t is quadratic in X_t, Y_t and U_t, Q is assumed to be a positive definite matrix, and the expectation operator is conditional on the history of w^t.

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.[5] In the current form, the loss function (5) depends on public beliefs via $Y_t = Y_{t|t}$.

---

[4]Notice that the policy control or parts of X_t are not precluded from entering directly in this forward looking constraint. This will be the case when, for example, the policy control is publicly observable such that $U_{t|t} = U_t$. A more general way to set up (3) would be to write

$$A^1_{yy}Y_{t+1|t} = A_{yy}Y_{t|t} + A^2_{yx}X_{t|t} + B^2_yU_{t|t} + A^3_{yx}X_t + B^3_yU_t$$

with the understanding that the measurement equation (2) implies $A^3_{yx}X_t + B^3_yU_t = A^3_{yx}X_{t|t} + B^3_yU_{t|t}$. This reduces then to (3) with $A_{yx} = A^2_{yx} + A^3_{yx}$ and $B_y = B^2_y + B^3_y$.

[5]Likewise, linear terms in $X_{t|t}$ and $U_{t|t}$ could be added to the transition equation for the backward looking variables.

The Simple NK Model

Equations (1), (2), (3) and (5) describe the class of LQ models for which we seek a solution to the optimal policy problem under discretion and hidden information. The simple NK model of Section 2 in the paper can be represented in the general framework as follows: The output gap equals the policy control, $U_t = x_t$ and is also identical to the measurement vector $Z_t = U_t$ such that $C_u = 1$ and $C_x = 0$. Furthermore, the backward and forward looking variables are $X_t = [\tau_t \, \varepsilon_t]'$, respectively $Y_t = \pi_t$.

In this model, the backward looking variables are purely exogenous, $A_{xy} = 0$ and $B_x = 0$, which considerably simplifies the solution under symmetric information (Svensson 2007, p. 24). However, in the hidden information problem the state vector will be augmented by public beliefs and the state vector will be endogenous. So no additional complication arises from allowing the backward looking variables in (1) to be partly endogenous, too.

The backward looking variables are

$$X_{t+1} = \begin{bmatrix} \tau_{t+1} \\ \varepsilon_{t+1} \end{bmatrix} = \underbrace{\begin{bmatrix} \rho & 0 \\ 0 & 0 \end{bmatrix}}_{A_{xx}} X_t + \underbrace{\begin{bmatrix} \sigma_\eta & 0 \\ 0 & \sigma_\varepsilon \end{bmatrix}}_{D} w_{t+1}$$

with $A_{xy} = 0$ and $B_x = 0$.

Inflation is the only forward looking variable of the private sector, $Y_t = \pi_t$, and the Phillips Curve corresponds to the associated forward looking constraint with $A_{yy}^1 = \beta$, $A_{yy} = 1$, $A_{yx} = 0$, and $B_y = -\kappa$.

B   Private Sector Equilibrium

The policymaker is constraint by the beliefs and the behavior of the private sector. The private sector is atomistic and takes policies as given. Before turning to optimal policy, it is useful to consider notions of private sector equilibrium for a given policy. This generalizes the discussion in

Section 2.2 on determining inflation for a given policy function.

Attention is limited to time-invariant, Markov-perfect equilibria. Policies will depend only on current levels of backward-looking variables and prior beliefs about those. In equilibrium, policy is a function of the Markov states:

$$U_t = F_1^0 X_t + F_2^0 X_{t|t-1} \tag{6}$$

for some F_1^0, F_2^0. Notice that this does not presuppose a commitment of the policymaker to such a rule. Discretion will rather require that this policy is ex-post optimal, such that the policymaker has no incentive to deviate once the private sector has formed beliefs consistent with the policy.

For the time being, the discussion adopts now the perspective of the private sector who takes the policy (6) as given when forming beliefs and making choices. This gives rise to a fairly strong notion of private sector equilibrium which can be applied to the simple NK model in the paper. As will be seen shortly, such an equilibrium need not always be unique. As will be shown below, a weaker notion of "temporary equilibrium" will in general be sufficient to constrain the discretionary policy problem.

Definition (Private Sector Equilibrium). *Given the policy in (6), the private sector equilibrium is a sequence of observations $\{Z_t\}$, perceived states $\{X_{t|t}\}$, perceived policies $\{U_{t|t}\}$ and private sector choices $\{Y_t\}$ such that*

- *Expectations and beliefs are rational. In this linear framework, they are formed using the Kalman filter with measurements Z_t.*

- *Choices are optimal, that is they satisfy the forward looking constraint (3).*

Using the Kalman filter, beliefs then evolve as

$$X_{t|t} = X_{t|t-1} + K^0(Z_t - Z_{t|t-1}) \tag{7}$$

$$U_{t|t} = F_1^0 X_{t|t} + F_2^0 X_{t|t-1} \tag{8}$$

Amongst others, the Kalman gain K^0 depends on the policy coefficients $\boldsymbol{F^0} = [F_1^0 \ F_2^0]$ in (6). Before turning to conditions for existence and uniqueness of the private sector equilibrium, some details are presented for the Kalman Filter.

Kalman Filter

For the policy given in (6), the private sector's Kalman filter combines (6) with (1) and (2) to obtain the state and measurement equations

$$X_{t+1} = \underbrace{\left(A_{xx} + B_x F_1^0\right)}_{\equiv A} X_t + A_{xy} Y_{t|t} + B_x F_2^0 X_{t|t-1} + D w_{t+1} \tag{9}$$

$$Z_t = \underbrace{\left(C_x + C_u F_1^0\right)}_{\equiv C} X_t + C_u F_2^0 X_{t|t-1} \tag{10}$$

and beliefs evolve as

$$X_{t|t} = X_{t|t-1} + K^0(Z_t - Z_{t|t-1}) \tag{11}$$

with Kalman gain K

$$K \equiv \mathrm{Cov}\left(X_t, Z_t - Z_{t|t-1}\right) \mathrm{Var}\left(Z_t - Z_{t|t-1}\right)^{-1} \tag{12}$$

The Kalman gain is identical to the coefficients of a least squares projection of X_t on $Z_t - Z_{t|t-1}$.

$$K^0 = \Sigma C' \left(C \Sigma C'\right)^{-1} \tag{13}$$

where Σ solves the Riccati equation

$$\Sigma = A\Sigma A' + DD' - AC' \left(C\Sigma C'\right)^{-1} CA' \tag{14}$$

The Kalman filter depends only on the policy coefficients F_1^0, via which policy reacts to X_t, and is independent of the reaction coefficients associated with the predetermined state variable $X_{t|t-1}$. The presence of private sector controls $Y_{t|t}$ and predetermined variables $X_{t|t-1}$ does not affect the Kaman gain.

The above assumes that the $N_z \times N_x$ matrix C has full row rank.[6] In principle (and also in practice) it can happen that C is collinear for some F_1^0. Numerically it is already critical if C is nearly collinear. This corresponds to situations when there are multiple observables[7] which are (almost) perfectly correlated such that $\text{Var}\,\tilde{Z}_t = C\Sigma C'$ is ill-conditioned. Economically, this means that a candidate policy $\boldsymbol{F^0}$ tries to mimic other signals in Z_t. I have never observed such mimicking strategies in equilibrium, but depending on initial conditions it can occur along the path of the policy improvement algorithm. In these cases, the Kalman filter is implemented by pruning the redundancies in the set of observable variables via a singular value decomposition of C. To obtain numerically stable solution, this is done for singular values of C smaller than 10^{-8}.

Conditions for Existence and Uniqueness

Optimal choices of the private sector solve the forward-looking constraint (3) given the policy (6) and private sector beliefs about X_t. Based on the Kalman filter, (3) and (1), this can be written as a system of expectational difference equations driven by the iid disturbance $\tilde{Z}_t$.[8]:

$$X_{t+1|t} = (A_{xx} + B_x \hat{F}^0)X_{t|t-1} + A_{xy}Y_{t|t} + (A_{xx} + B_x F_1^0)\,K^0\,\tilde{Z}_t$$

$$A_{yy}^1 Y_{t+1|t} = (A_{yx} + B_y \hat{F}^0)X_{t|t-1} + A_{yy}Y_{t|t} + (A_{yx} + B_y F_1^0)\,K^0\,\tilde{Z}_t$$

---

[6]Recall that $N_z < N_w \leq N_x$.

[7]This is the case in the model with belief shocks in Section 3, but not in the simple model of Section 2.

[8]Innovations $\tilde{Z}_t$ are defined relative to the public's prior belief $\tilde{Z}_t \equiv Z_t - Z_{t|t-1}$. By construction they are orthogonal to prior information of the private sector and are iid under *the public's probability measure*.

where $\hat{F}^0 \equiv F_1^0 + F_2^0$. The matrices

$$\bar{A} = \begin{bmatrix} I & 0 \\ 0 & A_{yy}^1 \end{bmatrix} \quad \text{and} \quad \bar{B} = \begin{bmatrix} (A_{xx} + B_x\hat{F}^0) & A_{xy} \\ (A_{yx} + B_y\hat{F}^0) & A_{yy} \end{bmatrix}$$

collect the coefficients on the endogenous variables.

This is the kind of linear systems studied by King and Watson (1998) and Klein (2000), where A_{yy}^1 is allowed to be singular. And their "counting rules" for stable and unstable roots can be applied to derive conditions for existence and uniqueness of the private sector equilibrium.

Proposition 1 (Existence and Uniqueness of Private Sector Equilibrium)**.** *Existence and unique-ness of a private sector equilibrium depend on the roots z of $|\bar{A}z - \bar{B}| = 0$ for matrices $\bar{A}$ and $\bar{B}$ defined above. A unique equilibrium exists only if there are N_x roots inside the unit circle and N_y outside. The matrices $\bar{A}$ and $\bar{B}$, and thus also the condition for existence and uniqueness, depend on the policy rule (6) but not on the Kalman gain K^0. This is an instance of certainty equivalence in linear rational expectations systems.*

Proof. The result follows from applying the solution methods of King and Watson (1998) or Klein (2000) to the linear rational expectations system above. Applying their methods yields the counting rule in the proposition and the solution has the form

$$Y_{t|t} = \bar{G}X_{t|t-1} + H_y\tilde{Z}_t$$

$$X_{t+1|t} = \bar{A}X_{t|t-1} + H_x\tilde{Z}_t$$

where $\bar{G}$ and $\bar{A}$ depend only on $\bar{A}$ and $\bar{B}$ but not on K^0 (for given policies, F_1^0 and F_2^0.) $\qquad\square$

In the simple NK model of Section 2 in the paper, the condition is trivially met[9] but in general this needs not be the case. A pertinent example is the nominal indeterminacy of Sargent and

---

[9]It is straightforward to check that there are two stable roots (ρ and 0) associated with the exogenous target variables and one unstable root ($1/\beta$) associated with inflation, which is the only forward looking variable.

Wallace (1975), which holds for any exogenous policy like (6) when the interest rate is the control variable. This applies also to the New Keynesian model, as discussed for example by Gali (2008).

If a unique solution exists, the construction of the private sector equilibrium is useful to analyze outcomes under different candidate policies as in Section 2. But for the purpose of constraining the discretionary policy problem, the above equilibrium notion is actually too strong. In this equilibrium, private sector expectations treat (6) as a time-invariant policy rule, carried out forever. And even though this will resemble the equilibrium outcome, it misrepresents the nature of the discretion problem where the policymaker can reoptimize his plans at each period. Therefore, non-uniqueness of a private sector equilibrium does not foreclose uniqueness of a discretionary equilibrium. To constrain the discretion problem, a weaker form of private sector equilibrium is sufficient. It is a temporary equilibrium in the spirit of Grandmont (1977):

Definition (Temporary Private Sector Equilibrium). *At a given point in time, the private sector has given beliefs about current policy according to (6). They are embodied in a Kalman gain K^0 used to update beliefs about X_t as in (7). Furthermore, people hold possibly different beliefs about future policies. They are embodied in a mapping G^0 which leads to expectations about future private decisions:*

$$Y_{t+1|t} = G^0 X_{t+1|t} \tag{15}$$

The temporary equilibrium then reduces to optimal choices which satisfy the forward looking constraint (3) given the beliefs in (15).

In a temporary equilibrium, private sector expectations of future choices are given. It is then straightforward to substitute the forward-looking variables by a linear combination of publicly

perceived policies and states:

$$Y_{t|t} = G_x^0 X_{t|t} + G_u^0 U_{t|t} \tag{16}$$

$$\text{where} \quad G_x^0 = (A_{yy}^1 G^0 A_{xy} - A_{yy})^{-1}(A_{yx} - A_{yy}^1 G^0 A_{xx})$$

$$\text{and} \quad G_u^0 = (A_{yy}^1 G^0 A_{xy} - A_{yy})^{-1}(B_y - A_{yy}^1 G^0 B_x)$$

The construction his temporary equilibrium is not a special feature of this hidden information setup. Similar computations are performed for example by Söderlind (1999) in his derivation of optimal Markov perfect policies under symmetric information.

C   Discretion Policy and Equilibrium

Discretionary policy is time-consistent. At each point in time the policymaker can reoptimize while taking his future optimizations as given. This leads to a recursive representation of the optimization problem as a dynamic program. The state variables of the policy problem are the backward looking variables and prior beliefs, there is no further history dependence. Furthermore, the policymaker must account for the rational expectations and optimal choices of the private sector. This is summarized in the following definition:

Definition (Discretionary Policy). *At each point in time, for given private beliefs embodied in F^0 and G^0, the policymaker chooses U_t to minimize*

$$V_t = \min_{U_t, Y_t, X_{t+1}} \left\{ L_t + \beta E_t V_{t+1}^0 \right\}$$

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

$$Y_{t|t} = G_x^0 X_{t|t} + G_u^0 U_{t|t}$$

where G_x^0 and G_u^0 are defined as in (16) above. The constraints correspond to the transition equation for X_t (1), and the private sector's temporary equilibrium (16). The continuation value

of this dynamic optimization problem, V_{t+1}^0, is a given function of future, Markov perfect state variables

$$S_{t+1} \equiv \begin{bmatrix} X_{t+1} \\ X_{t+1|t} \end{bmatrix}$$

Since the problem is linear quadratic, the value function can be taken to be linear quadratic as well (Bertsekas 2005):

$$V_{t+1}^0 = S_{t+1}' \boldsymbol{V^0} S_{t+1} + v^0 \tag{17}$$

The solution is then based on iterating between a conventional linear regulator problem and the Kalman Filter.[10] The regulator problem has the following form and is described in more detail in Appendix D:

$$S_t' \boldsymbol{V^*} S_t + v^* = \min_{U_t} \left\{ S_t' \boldsymbol{Q^0} S_t + 2 S_t' \boldsymbol{N^0} U_t + U_t' \boldsymbol{R^0} U_t + \beta E_t S_{t+1}' \boldsymbol{V^0} S_{t+1} + v^0 \right\} \tag{18}$$

$$\text{s.t.} \quad S_{t+1} = \boldsymbol{A^0} S_t + \boldsymbol{B^0} U_t + \boldsymbol{D} w_{t+1} \tag{19}$$

for given $\boldsymbol{F^0}$, G^0, a positive definite V^0 and a scalar v^0. The matrices $\boldsymbol{Q^0}$, $\boldsymbol{N^0}$, $\boldsymbol{R^0}$, $\boldsymbol{B^0}$ and $\boldsymbol{D}$ are derived in the next section.[11] The optimal policy is

$$U_t = -(\boldsymbol{R^0} + \beta \boldsymbol{B^{0\prime}} \boldsymbol{V^0} \boldsymbol{B^0})^{-1} (\boldsymbol{N^0} + \beta \boldsymbol{B^{0\prime}} \boldsymbol{V^0} \boldsymbol{A^0}) S_t \tag{20}$$

$$\equiv \boldsymbol{F^*} S_t$$

The optimal policy is linear as has been anticipated in (6). The policy appears certainty equivalent since it is independent of the shock loadings $\boldsymbol{D}$.[12] But in fact, the setup of the regulator itself is

---

[10]The definition of the discretion problem takes the matrix $\boldsymbol{V^0}$ and the scalar v^0 as given. In the policy improvement algorithm used to implement the solution, they will be calculated such as to be consistent with continuing the policy $\boldsymbol{F^0}$ and the beliefs G^0 forever. This is shown at the end of this section.

[11]Except for V^0 and v^0 matrices with superscript "0" depend on $\boldsymbol{F^0}$ and G^0. As will be see below, also V^0 and v^0 can be computed to be consistent with carrying out policies $\boldsymbol{F^0}$ and G^0 forever.

[12]Certainty equivalence is a well-known result of linear regulator problems (Bertsekas 2005).

not certainty equivalent since it depends on the private sector's Kalman filter. Policies are thus *not certainty equivalent.*

Definition (Equilibrium under Discretion). *Equilibrium under discretionary policymaking consists of sequences $\{U_t\}$, $\{X_t\}$, $\{Y_t\}$ and $\{Z_t\}$ such that each*

- *U_t solves the policymaker's problem*

- *Y_t is the solution to a temporary equilibrium whose underlying beliefs are consistent with the optimally chosen policies U_t*

- *X_t and Z_t evolve according to (1) and (2)*

where policies are a time-invariant function of the states.

Formally, this requires that $\boldsymbol{F^0} = \boldsymbol{F^*}$, and $G^0 = G^* = G_x^0 + G_u^0(F_1^* + F_2^*)$, where F_1^* and F_2^* partition $\boldsymbol{F^*}$ conformably with X_t and $X_{t|t-1}$. (K^0 is then consistent with $F_1^0 = F_1^*$.) Furthermore, the value function satisfies

$$\boldsymbol{V^0} = \boldsymbol{V^*} = \boldsymbol{Q^0} + \boldsymbol{N^0 F^*} + \boldsymbol{F^{*\prime} R^0 F^*} + \beta(\boldsymbol{A^0} + \boldsymbol{B^0 F^*})'\boldsymbol{V^0}(\boldsymbol{A^0} + \boldsymbol{B^0 F^*}) \qquad (21)$$

This equilibrium concept is similar to the self-confirming equilibria of Fudenberg and Levine (1993) and Sargent (1999) in that both are a fixed point of mutual beliefs and actions in multi-player games. However, in a self-confirming equilibrium, players hold erroneous beliefs about the structure of the economy, which are justified by observable outcomes. A similar fixed point of beliefs and outcomes is used in the limited-information rational expectations equilibria of Marcet and Sargent (1989a, 1989b) and Sargent (1991). This is different here, where the public completely knows and understands the structure of the economy.

D   Regulator for Discretion Problem

To set up the linear regulator problem shown in (18) and (19), the temporary equilibrium (16) and the Kalman filter (7) can be used to substitute $Y_{t|t}$ out of the loss function (5) and the transition equation (1) for X_t. The Kalman filter yields the transition equation for $X_{t|t-1}$.

The derivation proceeds by using the temporary equilibrium (16) and the Kalman filter (7) to substitute $Y_{t|t}$ out of the loss function (5) and transition equation (1) for X_t. The Kalman filter also yields the transition equation for $X_{t|t-1}$. The Kalman filter also depends on a prior belief about observables $Z_{t|t-1} = C_x X_{t|t-1} + C_u U_{t|t-1}$ and thus on a prior belief on policy. To simplify the regulator, it is assumed that this belief is consistent with $\boldsymbol{F^0}$ (as it will be in equilibrium), such that

$$Z_{t|t-1} = \underbrace{(C_x + C_u(F_1^0 + F_2^0))}_{\hat{C}} X_{t|t-1}$$

The Kalman update can be written as

$$X_{t|t} = KC_x X_t + (I - K\hat{C})X_{t|t-1} + KC_u U_t$$

Together with the temporary equilibrium (16) this yields

$$Y_{t|t} = \Gamma_x^0 X_t + \hat{\Gamma}_x^0 X_{t|t-1} + \Gamma_u^0 U_t$$

$$\text{with} \quad \Gamma_x^0 = (G_x^0 + G_u^0 F_1^0)KC_x$$

$$\hat{\Gamma}_x^0 = (G_x^0 + G_u^0 F_1^0)(I - K\hat{C}) + G_u^0 F_2^0$$

$$\Gamma_u^0 = (G_x^0 + G_u^0 F_1^0)KC_u$$

Loss Function

The loss function (5) can be rewritten in terms of the regulator's states and control using

$$
\begin{bmatrix} X_t \\ Y_{t|t} \\ U_t \end{bmatrix} = \underbrace{\begin{bmatrix} I & 0 & 0 \\ \Gamma_x^0 & \hat{\Gamma}_x^0 & \Gamma_u^0 \end{bmatrix}}_{\boldsymbol{H^0}} \begin{bmatrix} X_t \\ X_{t|t-1} \\ U_t \end{bmatrix}
$$

such that

$$
L_t = \begin{bmatrix} X_t \\ Y_t \\ U_t \end{bmatrix}' Q \begin{bmatrix} X_t \\ Y_t \\ U_t \end{bmatrix} = \begin{bmatrix} S_t \\ U_t \end{bmatrix}' \boldsymbol{H^{0\prime}} Q \boldsymbol{H^0} \begin{bmatrix} S_t \\ U_t \end{bmatrix}
$$

$$
= S_t' \boldsymbol{Q^0} S_t + 2 S_t' \boldsymbol{N^0} U_t + U_t' \boldsymbol{R^0} U_t
$$

where $\boldsymbol{Q^0}$, $\boldsymbol{N^0}$ and $\boldsymbol{R^0}$ conformably partition the above quadratic form as:

$$
\boldsymbol{H^{0\prime}} Q \boldsymbol{H^0} = \begin{bmatrix} \boldsymbol{Q^0} & \boldsymbol{N^0} \\ \boldsymbol{N^{0\prime}} & \boldsymbol{R^0} \end{bmatrix}
$$

State Transition

Likewise, the state transitions for X_t and $X_{t|t-1}$ can be derived as

$$
X_{t+1} = (A_{xx} + A_{xy}\Gamma_x^0)X_t + A_{xy}\hat{\Gamma}_x^0 X_{t|t-1} + (A_{xy}\Gamma_u^0 + B_x)U_t + Dw_{t+1}
$$

$$
X_{t+1|t} = \boldsymbol{A_{xx}} KC_x X_t + \left(\boldsymbol{A_{xx}}(I - K\hat{C}) + (A_{xy}G_u^0 + B_y)F_2^0 \right) X_{t|t-1} + \boldsymbol{A_{xx}} KC_u U_t
$$

where $\quad \boldsymbol{A_{xx}} = A_{xx} + A_{xy}(G_x^0 + G_u^0 F_1^0) + B_y F_1^0$

The matrices $\boldsymbol{A^0}$, $\boldsymbol{B^0}$ and $\boldsymbol{D}$ in (19) are thus given by:

$$\boldsymbol{A^0} = \begin{bmatrix} (A_{xx} + A_{xy}\Gamma_x^0) & A_{xy}\hat{\Gamma}_x^0 \\ \boldsymbol{A_{xx}}KC_x & \left(\boldsymbol{A_{xx}}(I - K\hat{C}) + (A_{xy}G_u^0 + B_y)F_2^0\right) \end{bmatrix}$$

$$\boldsymbol{B^0} = \begin{bmatrix} A_{xy}\Gamma_u^0 + B_x \\ \boldsymbol{A_{xx}}KC_u \end{bmatrix} \quad \text{and} \quad \boldsymbol{D} = \begin{bmatrix} D \\ 0 \end{bmatrix}$$

Value Function consistent with $\boldsymbol{F^0}$ and G^0

The policy improvement algorithm described in Section E uses a continuation value consistent with carrying out the policy $\boldsymbol{F^0}$ forever. The continuation value is linear quadratic in S_t as in (17). It is computed from the closed loop representation of the regulator obtained by plugging the policy $\boldsymbol{F^0}$ into (18) and (19). $\boldsymbol{V^0}$ solves the Lyapunov equation

$$\boldsymbol{V^0} = \left\{\boldsymbol{Q^0} + \boldsymbol{N^0}\boldsymbol{F^0} + \boldsymbol{F^{0\prime}}\boldsymbol{R^0}\boldsymbol{F^0}\right\} + \beta(\boldsymbol{A^0} + \boldsymbol{B^0}\boldsymbol{F^0})'\boldsymbol{V^0}(\boldsymbol{A^0} + \boldsymbol{B^0}\boldsymbol{F^0})$$

The equation has a unique solution if the matrix in curly braces is positive definite and if the closed loop transition matrix $(\boldsymbol{A^0} + \boldsymbol{B^0}\boldsymbol{F^0})$ has all eigenvalues inside the unit circle. The former is assured by the form of the original loss function[13] and the latter holds if a stationary equilibrium exists.[14]

Optimal policies are certainty equivalent (for given $\boldsymbol{F^0}$) and do not depend on v^0.[15] Still, the scalar v^0 can be computed from:

$$v^0 = \frac{\beta}{1 - \beta} \operatorname{tr}\left(\boldsymbol{V^0}\boldsymbol{D}\boldsymbol{D'}\right)$$

where tr is the trace operator.

---

[13]Please recall that Q in (5) is assumed to be positive definite.

[14]Efficient methods for solving Lyapunov equations are available for example via the LAPACK routines encoded in MATLAB or by using the doubling algorithms of Anderson et al. (1995).

[15]Sargent and Ljungqvist (2004) or Svensson (2007) give further details.

Unconditionally expected losses are computed from the unconditional variance covariance matrix of the states:

$$E(V_t^0) = \text{tr}\left(\boldsymbol{V^0} ES_t S_t'\right) + v^0$$

$$ES_t S_t' = \left(\boldsymbol{A^0} + \boldsymbol{B^0 F^0}\right)\left(ES_t S_t'\right)\left(\boldsymbol{A^0} + \boldsymbol{B^0 F^0}\right)' + DD'$$

E   Policy Improvement Algorithm

The equilibrium is a fixed point of public beliefs and policy actions and maps $(\boldsymbol{F^*}, G^*, \boldsymbol{V^*})$ into itself. An intuitive and efficient way to compute this fixed point is the following policy improvement algorithm. It is efficient, since policy improvement methods converge faster than value function iterations (Whittle 1996; Bertsekas 2005).[16] It is intuitive, since the algorithm uses the regulator (18) to seek for a one-period deviation from a candidate equilibrium. Non-existence of such a deviation is the defining property of equilibrium.

Formally, the algorithm starts with a candidate policy $\boldsymbol{F^0}$ and beliefs G^0 and computes the Kalman gain K^0 and continuation value $\boldsymbol{V^0}$ associated with continuing this policy forever. If the conditions for a private sector equilibrium are met (Proposition 1), one can even compute the G^0 consistent with F^0. The solution (20) to the above regulator problem then yields the optimal one-period deviation. As long as $\boldsymbol{F^0} \neq \boldsymbol{F^*}$ and $G^* \neq G^0$ there is no equilibrium. In this case, a new iteration starts using (F^*, G^*) as new candidate policies.

The difference with a value function iteration is that at each step, the regulator uses a continuation value consistent with carrying out the candidate policy forever whereas a value function iteration would update $\boldsymbol{V^0}_{j+1} = \boldsymbol{V^*}_j$ at the j-th step. In contrast, the policy improvement algorithm solves at each step an *infinite* horizon problem, where Kalman gain K^0 and continuation value $\boldsymbol{V^0}$, and if possible also G^0, are consistent with the candidate policy.

---

[16]Söderlind (1999) solves for optimal discretionary policies under symmetric information with value function iterations and comments on the slow performance of the algorithm.

Uniqueness of Equilibrium

Above I argued for uniqueness of the equilibrium in steady state, since the model then collapses to a full information setting with a unique steady state under simultaneous move timing. But off-steady state, the above equilibrium is an intricate fixed point between optimal one-period policies ($\boldsymbol{F}^*$), and public beliefs ($\boldsymbol{F^0}$, G^0). Formally, it is a fixed point between two Riccati equations, one from the policymaker's regulator problem, combining (20) and (21), the other associated with the public's Kalman Filter, see equation (14). Under suitable regularity conditions (Bertsekas 2005), both solve well-defined problems with unique solutions given the other's solution. However, to the best of my knowledge there exist no results on the existence and uniqueness of such nested systems. This is also the conclusion of Hansen and Sargent (2007, Chapter 15) who solve multi-player equilibria with similarly stacked Riccati equations.

However, in my practical experience, the algorithm typically converges, and if so always to the same equilibrium from arbitrary starting values for $(\boldsymbol{F^0}, G^0)$. In particular, over a wide range of calibrations (see Figure 4 in the paper), each equilibrium has been checked by drawing 50 times initial values from a mean zero Normal distribution with variance 10. Given equilibrium coefficients between zero and one, this is basically a flat prior. Each time, when the algorithm converges it converges to the same equilibrium.[17]

---

[17]Occasionally an equilibrium might not be found for a particular initial guess. In this case, another draw is made until the algorithm has converged 50 times.

References

Anderson, Evan W., Lars Peter Hansen, Ellen R. McGrattan, and Thomas J. Sargent. 1995. "On the mechanics of forming and estimating dynamic linear economies." Staff Report 198, Federal Reserve Bank of Minneapolis. Also published as Chapter in the Handbook of Computational Economics (1996).

Bertsekas, Dimitri P. 2005. *Dynamic Programming and Control*. 3rd Edition. Volume I. Belmont, MA: Athena Scientific.

Fudenberg, Drew, and David K. Levine. 1993. "Self-Confirming Equilibrium." *Econometrica* 61 (3): 523–545 (May).

Gali, Jordi. 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton, New Jersey: Princeton University Press.

Grandmont, Jean Michel. 1977. "Temporary General Equilibrium Theory." *Econometrica* 45 (3): 535–572 (apr).

Hansen, Lars P., and Thomas J. Sargent. 2007, January. "Robustness." unpublished Manuscript.

King, Robert G., and Mark W. Watson. 1998. "The Solution of Singular Linear Difference Systems under Rational Expectations." *Internatinal Economic Review* 39 (4): 1015–1026 (November).

Klein, Paul. 2000. "Using the generalized Schur form to solve a multivariate linear rational expectations model." *Journal of Economic Dynamics and Control* 24 (10): 1405–1423 (September).

Marcet, Albert, and Thomas J. Sargent. 1989a. "Convergence of least squares learning mechanisms in self-referential linear stochastic models." *Journal of Economic Theory* 48 (2): 337–368 (August).

Marcet, Albert, and Thomas J Sargent. 1989b. "Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information." *Journal of Political Economy* 97 (6): 1306–22 (December).

Sargent, Thomas J. 1991. "Equilibrium with signal extraction from endogenous variables." *Journal of Economic Dynamics and Control* 15 (2): 245–273 (April).

———. 1999. *The Conquest of American Inflation*. Princeton (NJ): Princeton University Press.

Sargent, Thomas J., and Lars Ljungqvist. 2004. *Recursive Macroeconomic Theory*. 2nd. Cambridge, MA: The MIT Press.

Sargent, Thomas J., and Neil Wallace. 1975. ""Rational" Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule." *The Journal of Political Economy* 83 (2): 241–254 (April).

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

Svensson, Lars E.O. 2007, June. "Optimization under Commitment and Discretion, the Recursive Saddlepoint Method, and Targeting Rules and Instrument Rules: Lecture Notes." *mimeo*.

Whittle, Peter. 1996. *Optimal Control, Basics and Beyond*. Wiley-Interscience Series in Systems and Optimization. Chichester: John Wiley & Sons.