Discreet Commitments and Discretion of Policymakers with Private Information*

Elmar Mertens†
Federal Reserve Board

First draft: 02/2009
Current draft: 01/2010

*I would like to thank Philippe Bacchetta, Jean-Pierre Danthine, Robert G. King, Peter Kugler and Yang Lu for comments at an earlier stage of this project. This research has been carried out within the National Center of Competence in Research "Financial Valuation and Risk Management" (NCCR FINRISK). NCCR FINRISK is a research program supported by the Swiss National Science Foundation.

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

Discreet Commitments and Discretion of Policymakers with Private Information

Abstract

This papers presents general methods to compute optimal commitment and discretion policies, when a policymaker is better informed about the realization of some shocks than the public. In this situation, public beliefs about the hidden information emerge as additional state variables, managed by the policymaker.

Under commitment, policy is additive in two components: The optimal policy, as if the government shared the public's information set and the systematic manipulation of that information set. Even under discretion, belief management imparts history dependence.

Illustrated in a New Keynesian economy with time-varying output targets of the policymaker, belief management improves outcomes compared to symmetric information. At the margin, the policymaker tries to be intransparent about policy objectives by engineering disturbances which lower public beliefs about the persistence of output targets.

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

Keywords: Optimal Monetary Policy, Commitment, Discretion, Incomplete Information, Kalman Filter

1. INTRODUCTION

Monetary policy is regularly conducted in situations of imperfect information. On the one hand, policymakers face uncertainties about the current state of the economy and the transmission mechanisms of shocks. On the other hand, policymakers are privy to confidential information, for example arising from internal staff forecasts, supervisory activities or their own biases and preferences. Policy decisions may reveal this information only imperfectly to the public. As a consequence, policy decisions affect the economy not only via conventional channels studied in settings of symmetric information, but also via the reaction of public beliefs to policy actions.

The issue of "bank learns about economy" has received widespread attention in the literature, for example by Sargent (1999), Orphanides (2001), Aoki (2003) or Svensson and Woodford 2003a; 2004. This problem has the simplifying property that 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.

In contrast, research has made comparably less advances into problems of the type "economy learns about bank", in particular in the context of dynamic general equilibrium models. This paper studies optimal policies under commitment and discretion in the case, where the policymaker is better informed. A key complication of this setting 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 asymmetric information facing both the private sector and policymakers. When agents are learning about the policymaker, policy constraints change in dramatic ways because of his strategic position in the economy. Those strategic effects are the main concern of this paper.

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 general, the entire distribution of public beliefs needs to be tracked by this kind of policy problem. The framework presented is very tractable by casting its model around a Gaussian framework with linear transition dynamics and constant variances. Tracking entire distributions then collapses to tracking only their means and can be handled with the Kalman filter. General methods are provided for linear quadratic economies.

In the rational expectations equilibria studied here, the private sector 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.

The methods are illustrated with a simple New Keynesian model. A striking result in this model is that hidden information policies lead to better outcomes than under full information. This may not be too surprising under discretion, since belief management imparts additional history dependence to policies.[1] But how could it be preferable for a policymaker to withhold information under commitment? To illustrate the forces at work here, assume a policy goal of minimizing inflation variance when inflation π_t is determined by the forward-looking Phillips Curve $\pi_t = \beta\pi_{t+1|t} + \kappa x_t$. Policy is committed to a rule, which implements an output gap x_t with mean zero and variance σ_x^2. For simplicity, suppose the output gap is iid. In this case, the policy loss would be equal to $\kappa^2\sigma_x^2$.

The information structures considered below, will allow the policymaker to forecast future output gaps beyond what is anticipated by the private sector.[2] In the extreme, suppose he could

[1] A companion paper (Mertens, 2009) studies the discretion policy in an extension of this model also in more detail.

[2] In the settings studied below, such effects arise from a signal extraction problem faced by the private sector, which

perfectly foresee the one-step ahead output gap. If this information was shared with the private sector, inflation variance would be equal to $\kappa^2(1+\beta^2)\sigma_x^2$ which is unambiguously higher and the economy would clearly be worse of by releasing this information. In the application considered below, this example will be generalized.

These benefits from "discreet" information policies stem from two features: First, the Philips Curve is the only expectational equation in this simple model. Second, policy losses are convex in inflation. The economic motivation for minimizing inflation variance is to reduce welfare losses associated with price dispersion. With sticky prices, price dispersion will be larger the more inflation rates fluctuate.

Closest to the methods of this paper are Aoki (2006) and Svensson and Woodford 2003a; 2004 who solve policy problems with the opposite informational asymmetry, namely when the policymaker is partially informed. In their case, there is a neat separation between the policy problem and the signal extraction *of the policymaker*. The policymaker's measurement equation is actually irrelevant for formulating the policy problem. As opposed to their work, in this paper it is the policymaker who manipulates the measurement equation of the partially informed public. What the public perceives as an innovation or not depends endogenously on policy and this complicates the analysis.

Related is also the work of Pearlman et al. (1986) who solve a class of expectational difference equations under partial information by combining rational expectation solutions based on the principle of "solving unstable roots forward, stable roots backward"[3] with the Kalman Filter. Their methods are substantially extended here to handle the presence of different information sets and the interdependence between policy and public beliefs.

The commitment policy analyzed here can be broken down into a part which reacts to public beliefs and where certainty equivalence is retained, and a non-certainty equivalent part which designs the optimal impulses to the public's information set. This aspect of the commitment algorithm shares some similarities with the robust control policies in forward-looking models studied

is endogenously affected by the systematic behavior of the policymaker.

[3]See Blanchard and Kahn (1980), Sargent (1987), King and Watson (1998) and Klein (2000).

by Hansen and Sargent (2007).

The remainder of this paper is structured as follows: Section 2 describes the informational assumptions in a class of linear quadratic models used throughout the paper. Section 3 studies the optimal commitment policy and Section 3 provides the corresponding analysis under discretion. The methods are illustrated with a very simple New Keynesian model in Section 5. Section 6 surveys the related literature before Section 7 concludes the paper.

2. INFORMATION SETS AND THE PRIVATE SECTOR

The methods presented in this paper apply to a wide class of linear-quadratic models, similar to what has been studied before by Söderlind (1999), Svensson and Woodford 2003a; 2004 and Hansen and Sargent (2007) to name but a few. The economy consists of a homogeneous private sector and a policymaker and there are four types of variables: Backward-looking variables, X_t. Policy controls, U_t. Decision variables of the private sector, Y_t. And finally, publicly observable variables, Z_t. These variables are vectors of dimensions N_x, N_u, N_z, and N_y respectively.

The backward looking variables can capture exogenous forcing variables but also endogenous states like capital, habits, or lagged variables, for example inflation in a model with price indexation. Driven by an exogenous N_w-dimensional white noise process they evolve as[4]

$$X_{t+1} = A_{xx}X_t + A_{xy}Y_t + B_xU_t + Dw_{t+1} \quad \text{where} \quad w_{t+1} \sim N(0, I) \tag{1}$$

2.1. *Information Sets*

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

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

only a linear combination of policy controls and backward looking variables:

$$Z_t = C_x X_t + C_u U_t \tag{2}$$

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

The history Z^t spans the public information set. 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). A sufficient condition to ensure superior information of the policymaker is that $N_z < N_w$. For any variable v_t, $v_{t|t} \equiv E(v_t|Z^t)$ denotes the expectation of v_t on the private sector's information set. Synonymously these expectations will be called public beliefs. For example, $X_{t|t-1}$ are the prior beliefs about X_t before observing Z_t.

The optimality conditions of private sector behavior are represented by an expectational linear difference equation involving only publicly observable variables and public sector expectations:[5]

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

By construction, $Y_t = Y_{t|t}$ always holds since public decisions are based on public information. In principle, Y_t could also be included to the measurement vector, but this would not add any new information. For the remainder of this paper, the forward-looking variables will be denoted $Y_{t|t}$ to reflect the constraint that they lie in Z^t.

Surprises in a variable v_t relative to the public's past information will be called "innovations". Formally, they are defined as $\tilde{v}_t \equiv v_t - v_{t|t-1}$. By construction, innovations are unpredictable from the public's perspective, $\tilde{v}_{t|t-1} = 0$. But typically they will be predictable based on the

[5]Notice that neither the policy control nor parts of X_t are 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

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

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

complete information set, and $E_{t-1}\tilde{v}_t$ need not be identical to zero. Innovations of the measurement variables, $\tilde{Z}_t$, provide an orthogonal decomposition of the public information set. The public information set is thus equivalently described by Z^t or $\tilde{Z}^t$.

Due to the linear-quadratic nature of the problem, attention is limited to linear policy functions. Since variables evolve linearly with Gaussian disturbances, rational expectations of the public can then be computed recursively from the Kalman filter. Given prior beliefs $z_{t|t-1}$ and $Z_{t|t-1}$, the public observes a realization of the measurement variables Z_t and updates its beliefs according to

$$v_{t|t} = v_{t|t-1} + K_v \tilde{Z}_t \quad \text{with Kalman gain} \quad K_v \equiv \text{Cov}\,(v_t, \tilde{Z}_t)(\text{Var}\,\tilde{Z}_t)^{-1} \tag{4}$$

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 coefficients are endogenous parameters. This poses a fixed point problem for the optimal policy setups considered below. Conditional on the Kalman gain, the problem will be linear-quadratic and amenable to standard solution methods. Its solution is linear, and in equilibrium it must be consistent with the coefficients from the Kalman gain.

2.2. *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 and optimal allocations for a given policy. As an example of such a setting, Erceg and Levin (2003) studied the Volcker disinflation as an imperfect information problem where the public learns about the inflation target of the Federal Reserve. In their model, policy follows a fixed Taylor rule whose intercept is given by the unobserved inflation target. The remainder of this paper provide methods which allow to solve for the *optimal* policy in such a setting.

Definition (Private Sector Allocation)**.** *A private sector allocation is a given sequence of observa-*

tions $\{Z_t\}$, perceived states $\{X_{t|t}\}$ and perceived policies $\{U_{t|t}\}$ as well as optimal private sector choices $\{Y_{t|t}\}$. The private sector choices are optimal, that is they satisfy the forward looking constraint (3), for a given evolution of the backward-looking variables (1) *and a given (linear) policy function.*

Definition (Private Sector Equilibrium)**.** *A private sector equilibrium combines an optimal allocation as defined above with optimal beliefs. Optimal beliefs of the public are rational expectations, conditional on the history of variables observed by the public, Z^t. In this linear framework, they are formed using the Kalman filter with measurements Z_t.*

A well known feature of linear quadratic models is the certainty equivalence of their solutions. To some extent, certainty equivalence is also preserved in models with imperfect information.[6] As will be shown next, certainty equivalence applies to the relationship between optimal allocations of the private sector and public beliefs about the state of the economy. For concreteness, a policy function needs to be assumed, for example

$$U_t = F_x X_t + F_y Y_{t|t} \tag{5}$$

(The results can easily be generalized to a larger state vector.) Given a linear policy function, optimal choices of the private sector are the solution to a set of expectational difference equations.

$$\underbrace{\begin{bmatrix} I & 0 \\ 0 & A^1_{yy} \end{bmatrix}}_{\bar{\boldsymbol{A}}} \begin{bmatrix} X_{t+1|t} \\ Y_{t+1|t} \end{bmatrix} = \underbrace{\begin{bmatrix} (A_{xx} + B_x F_x) & (A_{xy} + B_x F_y) \\ (A_{yx} + B_y F_x) & (A_{yy} + B_y F_y) \end{bmatrix}}_{\bar{\boldsymbol{B}}} \begin{bmatrix} X_{t|t} \\ Y_{t|t} \end{bmatrix} \tag{6}$$

Proposition 1 (Existence and Uniqueness of Private Sector Allocation (Pearlman et al., 1986))**.** *Existence and uniqueness of a private sector allocation depend on the roots z of $|\bar{\boldsymbol{A}}z - \bar{\boldsymbol{B}}| = 0$ for matrices $\bar{\boldsymbol{A}}$ and $\bar{\boldsymbol{B}}$ defined in* (5). *A unique equilibrium exists only if there are N_x roots inside the unit circle and N_y outside. The matrices $\bar{\boldsymbol{A}}$ and $\bar{\boldsymbol{B}}$, and thus also the condition for existence and*

[6] See for example Svensson and Woodford 2003a; 2004.

uniqueness, depend on the policy rule (5) but not on the Kalman filter or the volatility of shocks. 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 $Y_{t|t} = \bar{G}X_{t|t}$, $X_{t+1|t} = \bar{P}X_{t|t}$ where $\bar{G}$ and $\bar{P}$ depend only on $\bar{\boldsymbol{A}}$ and $\bar{\boldsymbol{B}}$ but not on K (for given policies, F_x and F_y.) □

A private sector equilibrium is defined relative to the information set of the public. In the analysis below, this information set will be endogenized and at least partly controlled by the policymaker. Optimal policies will reflect the policymaker's manipulations of the public's information set, which in turn depends on the stochastic structure of the information set. Policies will not be certainty equivalent and the above results of certainty equivalence of private sector allocations hold only for a given policy.

In order to compute a sequence of optimal allocations, it is sufficient to be given a sequence of innovations $\{\tilde{Z}_t\}$. State projections can then be computed from the filtering equation $X_{t+1|t+1} = X_{t+1|t} + K\tilde{Z}_t$. For beliefs to be optimal, the Kalman gain K can be computed using (1) as state and (5) as measurement equation.[7]

3. OPTIMAL POLICY UNDER COMMITMENT

The policymaker seeks to minimize the expected present value of current and future losses, where Q is assumed to be a positive definite matrix, and the expectation operator is conditional on the history of w^t.

$$V_t = E_t \sum_{k=0}^{\infty} \beta^k L_{t+k} \tag{7}$$

[7]Only the relationship between the *projections* of policies and states onto the public information set needs to be known in order to compute optimal allocations of the private sector; $U_{t|t} = \bar{F}_x X_{t|t} + F_y Y_{t|t}$. If the imperfect information problem is not degenerate, only some linear combinations of the backward looking variables can be perfectly observed. The coefficients F_x and $\bar{F}_x$ may differ as long as their difference lies in the nullspace which is orthogonal to perfectly observable combinations of X_t, such that $(\bar{F}_x - F_x)(X_t - X_{t|t}) = 0$.

$$\text{with} \quad L_t = \frac{1}{2} \begin{bmatrix} X_t \\ Y_{t|t} \\ U_t \end{bmatrix}' Q \begin{bmatrix} X_t \\ Y_{t|t} \\ U_t \end{bmatrix} \quad \text{with} \quad Q = \begin{bmatrix} Q_{xx} & Q_{xy} & N_x \\ Q'_{xy} & Q_{yy} & N_y \\ N'_x & N'_y & R \end{bmatrix} \tag{8}$$

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.[8] In its current form, the loss function (8) depends on public beliefs via $Y_t = Y_{t|t}$.

Under commitment, policy is a state-contingent plan, formulated once and for all at an initial time period. It specifies actions for all possible future events, but the plan is not to be re-optimized at later periods and will typically be time-inconsistent. Formally, the commitment policy of a privately informed policymaker is defined as follows:

Definition (Commitment Policy)**.** *The commitment policy is a state-contingent plan $U_t = U(w^t)$, devised once and for all by the policymaker at an initial date t_0 to minimize (7). The plan is not re-optimized at subsequent dates. It takes into account the natural constraints and the private sector behavior imposed by (1)–(3) as well as the partial information of the private sector, in particular their Kalman Filtering with the measurement equation (2). When implementing the policy at time t, the policymaker will be informed about the history of shocks until that period, w^t. The commitment is made from an ex-ante perspective at $t = t_0$, before the first shock is realized.*

At first glance it may appear surprising that commitment policy is only contingent on w^t and not Z^t. Due to its endogeneity, Z^t depends on w^t. Since the mapping from w^t to Z^t is known, it is sufficient for policies to be contingent on w^t. Equilibrium is a combination of the commitment policy and the private sector equilibrium defined in the previous section:

Definition (Commitment Equilibrium)**.** *The commitment equilibrium with private information consists of a policy plan U_t and sequences of forward and backward looking variables Y_t and X_t such that*

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

- *The policy plan is an optimal commitment policy.*

- *The private sector is in equilibrium. Forward looking variables Y_t are consistent with an optimal private sector allocation satisfying (3) for given $U_{t|t}$ and beliefs are rational expectations according to (4) with measurements following (2).*

- *The backward looking variables follow the state equation (1).*

As in Svensson and Woodford 2003a; 2004, a stationary commitment equilibrium is considered from a "timeless perspective" corresponding to a commitment made far in the past $t_0 \to -\infty$.

The "timeless perspective" implies that initial expectations are identical to unconditional expectations $E_{t_0}(v_t) = E(v_t)$ for any variables, v_t. Equivalently, the economy could be initialized with initial beliefs of private agents $X_{t_0|t_{-1}}$ consistent with the stationary policy plan, and innovations $\tilde{X}_{t_0}$ drawn from a time-invariant distribution of innovations $\tilde{X}_t$ consistent with the Kalman Filter to be derived below.

3.1. *Lagrangian of Policy Problem*

Formally, the policy problem is characterized by the Lagrangian

$$\mathcal{L}_{t_0} = E_{t_0} \sum_{t=t_0}^{\infty} \Big\{ \beta^t L_t + \beta^{t+1} \mu'_{t+1} \left[X_{t+1} - A_{xx} X_t - A_{xy} Y_{t|t} - B_x U_t - D w_{t+1} \right] \tag{9}$$
$$+ \beta^t \lambda'_t \left[A^1_{yy} Y_{t+1|t} - A_{yy} Y_{t|t} - A_{yx} X_{t|t} - B_y U_{t|t} \right] + \beta^t \zeta'_t \left[Z_t - C_x X_t - C_u U_t \right] \Big\}$$

The Lagrangian is minimized with respect to U_t, X_t and $Y_t = Y_{t|t}$ and maximized with respect to the multipliers μ_t, λ_t, and ζ_t. Variables involved in the constraints (1) and (2) lie in w^t and the associated multipliers μ_t and ζ_t are measurable with respect to w^t. The nature of λ_t will be discussed next.

Optimal policy commitment is typically not time-consistent (Kydland and Prescott, 1977) and the multiplier on the forward looking constraint, λ_t, represents the marginal benefits to the policymaker from adhering to past commitments at time $t+1$. They are an essential state variable for a recursive representation of policy following Kydland and Prescott (1980) and Marcet and Marimon (1998), which will be explored below.

The forward looking constraint describes optimal behavior of the private sector. By construction it only contains publicly observable variables and projections on the public information set. Its multiplier λ_t is thus measurable with respect to the public information set Z^t, $\lambda_t = \lambda_{t|t}$. Henceforth, this restriction will be reflected by denoting the multiplier as $\lambda_{t|t}$. As usual, the multiplier $\lambda_{t|t}$ can be interpreted as the policymaker's *marginal* value of relaxing the forward-looking constraint. The emphasis on *marginal* value implies in particular the restriction that Z^t is fixed and that all variables in (3) remain in Z^t. This excludes thought experiments asking for the shadow value of relaxing the forward looking constraint in some particular state of nature, $\bar{w}^t$, without applying the same change to the constraint in other states of nature, which are indistinguishable from $\bar{w}^t$ based on information contained in Z^t.

3.2. *Perturbing Conditional Expectations*

Optimal policy must satisfy the first-order conditions of the policy problem. When considering policy perturbations to (9), care needs to be taken with respect to the effects on information sets. In particular, two types of derivatives are of concern: Derivatives of a projection, like $X_{t|t}$, with respect to the underlying variable, X_t, and the derivative of a projection, like $X_{t|t}$ with respect to measurement variables Z_t. These derivatives need to consider perturbations on a state-by-state basis, keeping other variables constant.

Following Svensson and Woodford 2003a; 2004 and Aoki (2006), this is formalized with functional derivatives. Reflecting the ex-ante nature of the policymaker's commitments, the functional derivatives are normalized using the probabilities assigned to each state in the initial time period, which are assumed to be equal to their unconditional probabilities. A detailed discussion has been relegated to Appendix B.3. Representative for these two cases above, the appendix derives the following expressions

$$\frac{\partial}{\partial Y}\ \left(X_t'\,Q_{xy}\,Y_{t|t}\right) = Q_{xy}'\,X_{t|t} \tag{10}$$

$$\frac{\partial}{\partial Z}\ \left(X_t'\,Q_{xy}\,Y_{t|t}\right) = K_y'\,Q_{xy}'\,(X_t - X_{t|t}) + K_x'\,Q_{xy}\,\underbrace{(Y_t - Y_{t|t})}_{=0} \tag{11}$$

where K_x and K_y are the Kalman gains involved in projecting X_t, respectively Y_t, onto the public information set.[9]. Similar expressions hold for other terms in (9) involving expectations conditional on Z^t.[10]

In (10) perturbations of Y_t are projected onto Z^t, keeping Z^t and X_t constant. The perturbations are evaluated from an ex-ante perspective. As a shorthand, one could think of $E(X_t\ Y_{t|t}') = E(X_{t|t}\ Y_t')$ and then differentiate with respect to Y_t.

Perturbations like (11) consider purely informational effects from changing $Z(w^t)$, keeping $Y(Z^t)$ and $X(w^t)$ constant. This affects every probability density involving Z_t, be it marginal, joint or conditional, with direct consequences for the Kalman gains. (The probability of exogenous events, w^t, is of course fixed.) Ultimately, public observations depend on policies and backward looking variables via (2). Keeping policies fixed, while considering changes in the public information set is similar to finding the optimum in a consumption-savings problem, by perturbing savings, keeping consumption constant.

The Kalman filter implies possibly long-lasting effects of a current measurement variable on future projections $v_{t|t} = \sum_{j=0}^{\infty} K_j \tilde{Z}_{t-j}$ where $K_j = \mathrm{Cov}\,(v_{t+j}, \tilde{Z}_t)\,\mathrm{Var}\,(\tilde{Z}_t)^{-1}$. The first-order condition of the Lagrangian (9) with respect to Z_t would thus have to track an infinite number of forward looking terms similar to (11). To avoid this, it will be convenient to cast the policy problem in recursive form using the saddlepoint formulation of Marcet and Marimon (1998).

[9] As in (4) this means $K_y = \mathrm{Cov}\,(Y_t, \tilde{Z}_t)\ \mathrm{Var}\,(\tilde{Z}_t)^{-1}$ and $K_x = \mathrm{Cov}\,(X_t, \tilde{Z}_t)\ \mathrm{Var}\,(\tilde{Z}_t)^{-1}$.

[10] A particular property of the two examples given here is that $Y_t = Y_{t|t}$. The derivatives hold however also in the general case where none of the variables lies in Z^t. Such a situation could occur if for example projections $X_{t|t}$ would directly enter the loss function (8) or the transition equation (1).

3.3. *Recursive Formulation and a Separation Result*

Following Marcet and Marimon (1998), the policy problem can be converted into a recursive saddlepoint problem by treating the lagged multiplier $\lambda_{t-1|t-1}$ as a fictional state variable, initialized with $\lambda_{t_0-1|t_0-1} = 0$, and defining the modified Lagrangian

$$\mathcal{L}_{t_0}^{MM} = \mathcal{L}_{t_0} + \beta^{-1}\lambda_{t_0-1|t_0-1}A_{yy}^1 Y_{t_0}$$

Using, $L_t^{MM} \equiv L_t + \beta^{-1}\lambda_{t-1|t-1}A_{yy}^1 Y_{t|t-1}$, the associated Bellman equation is

$$\begin{aligned} V_t^{MM} = \min_{U_t, X_{t+1}, Y_{t|t}, Z_t} \ \max_{\mu_{t+1}, \lambda_{t|t}, \zeta_t} \quad & L_t^{MM} + \beta E_t V_{t+1}^{MM} \\ & +\beta\mu_{t+1}' \left[X_{t+1} - A_{xx}X_t - A_{xy}Y_{t|t} - B_x U_t - Dw_{t+1}\right] \\ & -\lambda_{t|t}' \left[A_{yy}Y_{t|t} + A_{yx}X_{t|t} + B_y U_{t|t}\right] \\ & +\zeta_t' \left[Z_t - C_x X_t - C_u U_t\right] \end{aligned} \tag{12}$$

In a full information model, the state variables of the recursive problem are X_t and $\lambda_{t-1|t-1}$ (Svensson, 2007). Under imperfect information, the public's prior beliefs $X_{t|t-1}$ must be added to the list of state variables.[11] Since $\tilde{X}_t$ is orthogonal to $X_{t|t-1}$ and $\lambda_{t-1|t-1}$, it will be convenient to write the state vector as

$$S_t = \begin{bmatrix} \tilde{X}_t \\ X_{t|t-1} \\ \lambda_{t-1|t-1} \end{bmatrix} \tag{13}$$

Given the linear quadratic structure of the problem, it will be guessed (and verified) that the

[11]It is the public's prior, beliefs which enter the state vector, not the posterior beliefs $X_{t|t}$. The latter will be formed after observing current data which is influenced by current policy. The pertinent history of beliefs leading to this policy is described by $X_{t|t-1}$.

value function is linear quadratic and policies are linear in the state vector.

$$U_t = F_x \tilde{X}_t + \bar{F}_x X_{t|t-1} + \bar{F}_\lambda \lambda_{t-1|t-1} \tag{14}$$

The first-order conditions to (12) form a set of linear difference equations, involving expectations with respect to both w^t and Z^t. Details are given in Appendix B. The coefficients in some equations depend on the yet to be determined Kalman gains and value function, which in turn depend on the optimal policy, leading to a fixed point problem.

As a key result, the policy problem can neatly be separated into two steps. First, the FOC are conditioned down onto the public information set. This leads to a solution relating the projection of policy, $U_{t|t}$, to the projections of the state variables $X_{t|t}$ and $\lambda_{t-1|t-1}$, which will be called *public information policy*. In the second step, the remaining terms, which are orthogonal to the public information set, can be solved to derive the optimal policy innovations $\tilde{U}_t$ which will determine the evolution of the public information set. This will be called *innovation policy*.

3.4. *Public Information Policy*

Conditioning down the FOC onto the public information set eliminates all terms involving the unknown Kalman gains. In fact, it leads to the same solution known from the commitment problem under symmetric information.

Proposition 2 (Certainty Equivalence of Public Information Policy)**.** *The first-order conditions to the commitment problem (12) imply a solution for the projections of policy onto the public information set which is linear in $X_{t|t}$ and $\lambda_{t-1|t-1}$*

$$U_{t|t} = \bar{F}_x X_{t|t} + \bar{F}_\lambda \lambda_{t-1|t-1} \tag{15}$$

$$Y_{t|t} = \bar{G}_x X_{t|t} + \bar{G}_\lambda \lambda_{t-1|t-1} \tag{16}$$

$$\lambda_{t|t} = \bar{H}_x X_{t|t} + \bar{H}_\lambda \lambda_{t-1|t-1} \tag{17}$$

$$X_{t+1|t} = \bar{P}_x X_{t|t} + \bar{P}_\lambda \lambda_{t-1|t-1} \tag{18}$$

This solution is certainty equivalent. It does not depend on the shock loadings D of the backward looking variables in (1), nor on the measurement equation (2) and the Kalman Filter. It is identical to the commitment policy under symmetric information.

Proof. The result and its proofs correspond to the certainty equivalence found by Svensson and Woodford 2003a; 2004 for policies which are restricted to lie in the public information set. Details can be found in Appendix B. □

As is usual for commitment problems, the co-state variable $\lambda_{t-1|t-1}$ imparts history dependence on policies. In the public information policy, it is the history of public beliefs of state variables not necessarily their actual values which are tracked by the co-state variables.

$$\lambda_{t|t} = \sum_{k=0}^{\infty} (\bar{H}_\lambda)^k \bar{H}_x X_{t-k|t-k}$$

While the evolution of the information set Z^t remains to be determined, the public information policy identifies the coefficients $\bar{F}_x$ and $\bar{F}_\lambda$ in (14).[12] In addition, the Kalman gains are determined relative to K_x, the gain associated with the backward looking variables X_t. For example, the policy for $U_{t|t}$ in (16) yields

$$K_u = \text{Cov}\,(U_t, \tilde{Z}_t)(\text{Var}\,(\tilde{Z}_t))^{-1} = \text{Cov}\,(U_{t|t}, \tilde{Z}_t)(\text{Var}\,(\tilde{Z}_t))^{-1} = \bar{F}_x K_x$$

The measurement equation (2) and the FOC associated with Z_t play no role in the public information policy. The reason is that derivatives like (11), are perpendicular to the public information set. As a consequence the multiplier on the measurement equation is not observable by the public either. Intuitively, this means that if private information is to remain hidden from the public, so must be the value he attaches to that information.[13]

[12] This can be seen from condition down (15) on Z^{t-1}.

[13] In the words of Sir Humphrey Appleby: "Thou who keeps a secret shall keep it a secret, that he keeps a secret."

Proposition 3 (Marginal Value of Information is Unobservable to the Public). *$\zeta_t = \zeta(w^t)$ is the shadow price attached to a particular realization of the measurement vector Z_t when the realized history equals w^t. The first-order condition of (9) with respect to Z_t implies $\zeta_{t|t} = 0$.*

Proof. The proof is straightforward based on the first order condition of (9) for Z_t. (See equation (41) in Appendix 3.) □

Proposition 3 extends a result from Aoki (2006) and Svensson and Woodford 2003a; 2004, obtained for the case when Z^t describes the policymaker's information set. In this case, the measurement equation can be dropped from the Lagrangian. The intuition is simple: For a given history of past innovations $\tilde{Z}^{t-1}$, perturbing Z_t considers alternative plans for the innovation $\tilde{Z}_t$. By construction, these are unforecastable from Z^{t-1} and one cannot plan how to surprise oneself. In the present paper, it is however the policymaker who has superior information and Z^t represents the public's information set. ζ_t then measures the marginal value of private information to the policymaker.

Given the public information policy, it remains to determine the innovation policy.

$$\tilde{U}_t = F_x \tilde{X}_t \tag{19}$$

Together with $U_{t|t-1} = \bar{F}_x X_{t|t-1} + \bar{F}_\lambda \lambda_{t-1|t-1}$, this will complete the optimal policy solution.

3.5. *Innovation Policy*

By committing to a particular policy rule like (19), the policymaker can influence which combination of state variables is observable to the public and which not. Substituting the policy rule (19) into the measurement equation (2) yields $(C_x + C_u F_x)(X_t - X_{t|t}) = 0$. Unobservable combinations of X_t lie in the nullspace of $C_x + C_u F_x$. While the range of this nullspace is fixed by the number of variables ($N_X - N_Z$), the policymaker can determine its location through the choice of F_x. Not all the $N_u \times N_x$ coefficients of F_x are in fact free, since public information policy also restricts the projection of policy innovations on the public's current information set,

$$\tilde{U}_{t|t} = F_x \tilde{X}_{t|t} = \bar{F}_x \tilde{X}_{t|t} \quad \Rightarrow \quad (F_x - \bar{F}_x) K_x = 0$$

Due to the dependence of the Kalman gain K_x on F_x this imposes $N_u \times N_z$ *non-linear* restrictions on the innovation policy. Section 5 provides a neat geometric illustration for the bivariate case, where innovation policy chooses its location as a line in a plane.

Having solved for the public information policy, the remainder of the FOC characterizes the optimal nullspace. Appendix B.2 derives the following condition:

$$\Phi_x (X_t - X_{t|t}) + \Phi_u (U_t - U_{t|t}) = 0 \tag{20}$$

for some coefficient matrices Φ_x and Φ_u. Together with the innovation representation of the Kalman Filter this can be solved for the optimal innovation policy

$$\tilde{U}_t = -\left(\Phi_u - \Phi_z C_u\right)^{-1} \left(\Phi_x - \Phi_z C_x\right) \tilde{X}_t = F_x \tilde{X}_t \tag{21}$$

where $\Phi_z = (\Phi_x + \Phi_u \bar{F}_x) K_x$.

The coefficients in (20) and (21) depend on the Kalman gain K_x and the value matrix $\tilde{V}_{xx}$ which in turn depend on the optimal policy F_x. The solution seeks for a fixed point mapping F_x in itself. A convenient and intuitive algorithm for doing this is the policy improvement algorithm described in Appendix E.

4. OPTIMAL POLICY UNDER DISCRETION

Commitment policies are optimal from an ex-ante perspective, but they need not be time-consistent. This section derives the optimal time-consistent policy, chosen by a discretionary policymaker who is free to re-optimize his choices at each point in time. Attention is limited to Markov perfect equilibria, which avoids the multiplicity of equilibria associated with the kind of reputational mechanisms considered by Barro and Gordon (1983) or Chari and Kehoe (1990). In the spirit of "bygones are bygones", state variables in a Markov-perfect equilibrium must be relevant for current payoffs. A companion paper, Mertens (2009), applies the general framework presented here to a New Keynesian model, which is also used to illustrate the methods of this paper in Section 5.

Under discretion, the policymaker takes the history of beliefs, summarized by $X_{t|t-1}$ as well as the coefficients of the public's Kalman filter (4) as given. In equilibrium, the beliefs and the public's Kalman filter depend on his policies in ways which are similar to the commitment problem. But a discretionary policymaker considers only the effect of his policy choice on the current value of the measurements Z_t observed by the public. In contrast to the commitment problem, the *systematic* effects of his policy on the Kalman filter, embodied in the gain coefficients, are neglected. In a nutshell, the policymaker considers the effect his policy has on the current "data" used by the public's Kalman regressions, but not on the "regression slopes" induced by his systematic behavior in equilibrium.

The policymaker must account for the optimal choices of the private sector. Due to the forward looking nature of (3), these depend on private sector beliefs of future choices and future policies, to which the current policymaker has not committed. He rather takes private sector expectations of future choices $Y_{t+1|t}$ as a given function of future states. In principle, the private sector equilibrium described in Section 2 could be used to relate these beliefs to given future policies using the private sector allocation defined in Section 2.

If a unique allocation exists, the construction of the private sector equilibrium is useful to analyze outcomes under different candidate policies. But the absence of a private sector equilibrium need not be a problem, since for the purpose of constraining the discretionary policy problem, this equilibrium notion is stronger than necessary.[14] Only a temporary equilibrium in the spirit of Grandmont (1977) is necessary:

[14] The private sector allocation defined in Section 2 forms private sector expectations based on a a time-invariant policy rule, carried out forever. Even though this eventually resembles 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.

Definition (Temporary Private Sector Equilibrium)**.** *At a given point in time, the private sector has given beliefs about current policy. They are embodied in a Kalman gain K_x^0 used to update beliefs about X_t as in (4). Furthermore, people hold beliefs about future policies. These beliefs are embodied in a mapping G^0 between future states and future private decisions:*

$$Y_{t+1|t} = G^0 X_{t+1|t} \tag{22}$$

Given the beliefs and expectations in (22), the temporary equilibrium then reduces to optimal (static) choices of the private sector satisfying the forward looking constraint (3).

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

$$\begin{aligned} & Y_{t|t} = G_x^0 X_{t|t} + G_u^0 U_{t|t} \qquad (23) \\ \text{where} \quad & G_x^0 = (A_{yy}^1\, G^0\, A_{xy} - A_{yy})^{-1}(A_{yx} - A_{yy}^1\, G^0\, A_{xx}) \\ \text{and} \quad & G_u^0 = (A_{yy}^1\, G^0\, A_{xy} - A_{yy})^{-1}(B_y - A_{yy}^1\, G^0\, B_x) \end{aligned}$$

Constructing this temporary equilibrium is not a special feature of the hidden information setup. When there are endogenous state variables, similar computations are necessary for computing optimal Markov perfect policies under symmetric information (Söderlind, 1999).

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 X_t and prior beliefs $X_{t|t-1}$, there is no further history dependence.

Equilibrium policy is a function of the Markov states:

$$U_t = F_x^0 X_t + F_b^0 X_{t|t-1} = F^0 S_t \qquad \text{with} \quad S_t \equiv \begin{bmatrix} X_{t+1} \\ X_{t+1|t} \end{bmatrix} \tag{24}$$

for some F_x^0, F_b^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.

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*

$$\begin{aligned} V_t &= \min_{U_t, Y_{t|t}, X_{t+1}} \left\{ L_t + \beta E_t V_{t+1}^0 \right\} \\ \textit{s.t.} \quad X_{t+1} &= A_{xx} X_t + A_{xy} Y_{t|t} + B_x U_t + D w_{t+1} \\ Y_{t|t} &= G_x^0 X_{t|t} + G_u^0 U_{t|t} \end{aligned}$$

where G_x^0 and G_u^0 are defined as in (23) above. The constraints correspond to the transition equation for X_t (1), and the private sector's temporary equilibrium (23). The continuation value of this dynamic optimization problem, V_{t+1}^0, is a given function of future values of the Markov-perfect state variables S_{t+1}.

Since the problem is linear quadratic, the value function will be linear quadratic as well (Bertsekas, 2005). The solution is then based on iterating between a conventional linear regulator problem and the Kalman Filter. The regulator problem has the following form:

$$V_t = S_t' \boldsymbol{V}^* S_t + v^* = \min_{U_t} \left\{ S_t' \boldsymbol{Q^0} S_t + 2 S_t' \boldsymbol{N^0} U_t + U_t' \boldsymbol{R^0} U_t + \beta E_t \left(S_{t+1}' \boldsymbol{V^0} S_{t+1} + v^0 \right) \right\} \tag{25}$$

$$\text{s.t.} \quad S_{t+1} = \boldsymbol{A^0} S_t + \boldsymbol{B^0} U_t + \boldsymbol{D} w_{t+1} \tag{26}$$

for given 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 Appendix D.[15] The optimal policy is

$$U_t = -(\boldsymbol{R^0} + \beta \boldsymbol{B^0}'\boldsymbol{V^0}\boldsymbol{B^0})^{-1}(\boldsymbol{N^0} + \beta \boldsymbol{B^0}'\boldsymbol{V^0}\boldsymbol{A^0})\ S_t \tag{27}$$
$$\equiv F^* S_t = F_x X_t + F_b X_{t|t-1}$$

The optimal policy is linear which confirms the guess embodied in using the Kalman filter to model conditional expectations. The policy appears certainty equivalent since it is independent of the shock loadings $\boldsymbol{D}$.[16] But the setup of the regulator 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 $F^0 = F^*$, and $G^0 = G^* = G^0_x + G^0_u(F^*_x + F^*_b)$, where F^*_x and F^*_b partition F^* conformably with X_t and $X_{t|t-1}$. (K^0_x is then consistent with $F^0_x = F_x$.) Furthermore, the value function satisfies

$$\boldsymbol{V^0} = \boldsymbol{V^*} = \boldsymbol{Q^0} + \boldsymbol{N^0}F^* + F^{*\prime}\boldsymbol{R^0}F^* + \beta(\boldsymbol{A^0} + \boldsymbol{B^0}F^*)'\boldsymbol{V^0}(\boldsymbol{A^0} + \boldsymbol{B^0}F^*) \tag{28}$$

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-

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

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

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.

5. ILLUSTRATION WITH A NEW KEYNESIAN MODEL

This section applies the optimal policies under commitment and discretion to a simple New Keynesian economy. The model has been chosen for its simplicity and low dimensionality, not for its realism, in order to better illustrate the concepts of the previous sections. A companion paper provides a more detailed discussion of the discretion policy under imperfect information in an extended version of the model (Mertens, 2009).

5.1. *New Keynesian Economy*

The model is largely identical to the textbook model of optimal policy in a New Keynesian model known from Clarida et al. (1999), Walsh (2003) or Woodford (2003). The only difference is a stochastic preference shock to the policymaker's objective function which is unobservable to the public. Otherwise the 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.

5.1.1. *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. Readers familiar with this setting can skip this part and move over to the discussion of policy objectives and hidden information.

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

$$\pi_t = \beta \pi_{t+1|t} + \kappa\, x_t \tag{29}$$

where π_t is inflation and x_t is the output gap[17]. 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.[18] As before, $v_{t+1|t}$ denotes the private sector forecast for some variable v_{t+1}.

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.[19] The discussion will exclusively focus on monetary shocks which leave the natural rate unaffected. Conditional on those shocks, variations in the output gap are thus identical to variations in output and consumption.

5.1.2. *Policy Objectives*

The policymaker seeks to minimize a present value of expected losses

$$E_t \sum_{k=0}^{\infty} \beta^k \left\{ \pi_{t+k}^2 + \alpha_x (x_{t+k} - \bar{x}_{t+k})^2 \right\} \tag{30}$$

with $\alpha_x \geq 0$. The expectations operator E_t reflects the policymaker's information set. 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.

Mechanically, the target is necessary to introduce a policy trade-off between inflation and out-

[17]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.

[18]Details can be found in Gali (2003, p. 159) from whom notation is adopted.

[19]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 which track the difference between the RBC core and the actual economy. This separation has been widely adopted in the textbook treatments of Gali (2003), Walsh (2003) and Woodford (2003).

put gap stabilization in the model. Typically this is done by adding a cost-push disturbance to the Phillips Curve as in Clarida et al. (1999). In a full information setting, an exogenous cost-push shock would be isomorphic to the output target introduced here.[20] In this imperfect information setting it does however matter, whether this disturbance enters the Phillips Curve and where it would be directly observable to the public, or whether it enters via the policymakers loss function, without being directly in the public information set. The latter approach has been chosen here, since it allows to keep the number of shocks and observables smaller.

The output target can be viewed as arising from time-varying preferences of the policymaker as in Cukierman and Meltzer (1986). 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.[21] 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.[23]

Alternatively, 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. A key restriction of the information structure introduced here is however, that $\bar{x}_t$ cannot be observed by the representative private agent. This is assumption is most easily satisfied by interpreting the output target as arising from policymaker's preferences.

[20]To see this isomorphism, the target could be treated as a cost-push shock and the deviation of output gap from target, $x_t - \bar{x}_t$ could be re-labeled as an output gap, corresponding to the model of Clarida et al. (1999).

[21]Under either interpretation, the output target is capturing a form of heterogeneity otherwise not present in the model. In particular (30) 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 (30) at the average output target (here: zero)[22], $L_t^R = \pi_t^2 + \alpha_x x_t^2$. But without specifying the underlying heterogeneity and associated welfare weights this is at best an aggregation with unknown distributional consequences.

[23]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. Abrams (2006) gives a striking account of hidden but forceful policy influences in the U.S.. His study documents how President Nixon covertly pressured the then Chairman of the Federal Reserve, Arthur Burns to ease policy in the run-up to the Great Inflation.

5.1.3. *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:

$$\bar{x}_t = \tau_t + \varepsilon_t \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2) \tag{31}$$

$$\tau_{t+1} = \rho\, \tau_t + \eta_{t+1} \qquad \eta_t \sim N(0, \sigma_\eta^2) \quad \text{and} \quad 0 < |\rho| < 1 \tag{32}$$

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, $Z^t = x^t$. 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 $\varepsilon_{t|t} \neq \varepsilon_t$.

5.1.4. *LQ Representation*

The simple NK model laid out here can be represented in the general framework of Section 2 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 $Y_t = \pi_t$. The Phillips Curve (29) corresponds to the forward looking constraint (3) with $A^1_{yy} = \beta$, $A_{yy} = 1$, $A_{yx} = 0$, and $B_y = -\kappa$. The evolution of the target components corresponds to (1) with $A_{xy} = 0$, $B_x = 0$ and

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

5.2. *Signal Extraction for Given Policy*

This section studies the public's signal extraction problem for a given policy. In particular, it documents the key role played by the relative sensitivity of policies to realizations in the two target components in shaping the response of public beliefs to observed policies.

Optimal policies are linear. Coefficient values will differ between commitment and discretion, but in general policies have the form[24]

$$x_t = f_\tau \tau_t + f_\varepsilon \varepsilon_t + \bar{f}_\tau \tau_{t|t-1} + \bar{f}_\lambda \lambda_{t-1|t-1} \tag{33}$$

Since $\varepsilon_{t|t-1} = 0$, it need not be tracked as a state variable.

There are two unobserved states, τ_t and ε_t. Since ε_t is white noise, it is sufficient to consider (32) as the only state equation of the Kalman filter and use (33) as measurement equation. Only the policy loadings on the *realized* components of the output target are relevant for the Kalman filter and neither $\bar{f}_\tau$ nor $\bar{f}_\lambda$.

A key role in the public's signal extraction is played by the ratio of policy loadings f_ε / f_τ. This "mixing ratio" determines how much a policy innovation reveals about τ_t instead of ε_t. Under commitment, this mixing ratio is consciously chosen by the policymaker. It allows the policymaker to change the signal-to-noise ratio in the public's signal extraction problem. A discretion policymaker disregards the systematic effects of his choices on the Kalman filter.

As discussed in Section 3, there are only $N_U \times (N_X - N_Z)$ free coefficients in the innovation part of the commitment problem. In this model, with $N_U = N_Z = 1$ and $N_X = 2$, this leaves only one degree of freedom, corresponding to the choice of the mixing ratio.

There is a neat geometric representation for the role of f_ε / f_τ: Figure 1 depicts the two-dimensional space of innovations in this model,[25] spanned by the orthonormal innovations $\tilde{\tau}_t / \sqrt{\Sigma_\tau}$ and $\varepsilon_t / \sigma_\varepsilon$.[26] Policy innovations span the subspace of what is observable to the public

$$\tilde{x}_t = f_\tau \tilde{\tau}_t + f_\varepsilon \varepsilon_t$$

[24] Discretion policies do not respond to the lagged Lagrange multiplier $\lambda_{t-1|t-1}$, which can be thought of as $\bar{f}_\lambda = 0$ here.

[25] Formally, this is a two-dimensional Hilbert space with inner product defined by the (full information) expectations operator $E(\cdot)$. Hilbert space analysis is for example used by Hansen and Richard (1987) to study the role of information sets for asset pricing restrictions.

[26] Since ε_t is iid we have $\varepsilon_t = \tilde{\varepsilon}_t$.

Figure 1: Innovation Space

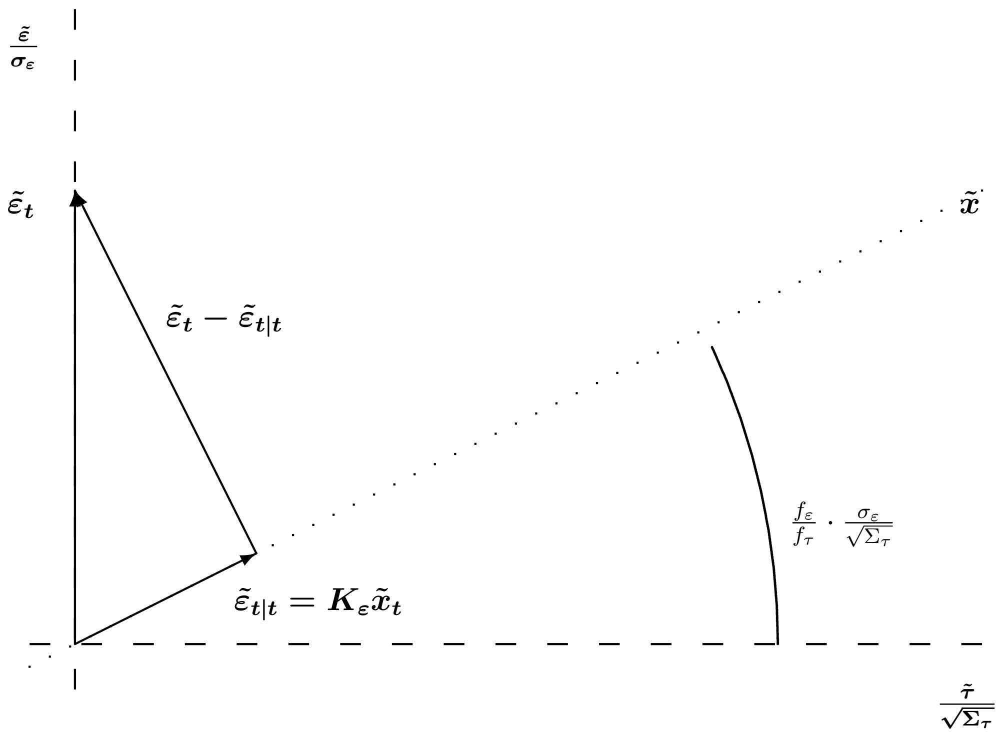

Note: Projections $\varepsilon_{t|t}$ are constrained in the subspace spanned by $\tilde{x}_t$ (dotted line) and projection errors $\tilde{\varepsilon}_t - \tilde{\varepsilon}_{t|t}$ are perpendicular to this line. The same holds for projections $\tilde{\tau}_{t|t}$.

which is just a line in the plane spanned representing the innovation space. Projected innovations $\tilde{v}_{t|t}$ of any variable v_t lie on this line. Projection errors $v_t - v_{t|t}$ are perpendicular to it.

It is easy to see how projections $\tilde{\tau}_{t|t}$ converge to the actual innovations $\tilde{\tau}_t$ as the line flattens, when the mixing ratio approaches zero.[27] (In this case, the innovations τ_t would also converge to the actual shocks η_t.) The opposite happens when the mixing ratio gets very large: Nothing is revealed about the persistent component of the target.

Graphically, Regression R^2 between target and policy innovations are measured by the distance between the $\tilde{x}_t$-space and the axes representing the basis innovations $\tilde{\tau}_t$ and ε_t. The regression R^2 for projecting ε_t is increasing in the mixing ratio f_ε/f_τ and vice versa for $\tilde{\tau}_t$. Analytically, the

[27]The slope of the subspace of observables is equal to $(f_\varepsilon/f_\tau \cdot \sigma_\varepsilon/\sqrt{\Sigma_t})$. For a fixed innovation variance Σ_τ, changes in the mixing ratio carry over 1:1 to changes in the slope. This conjectures also holds when considering the effects of the mixing ratio on the innovation variance Σ_τ. Application of the implicit function theorem to (34) yields $\frac{\partial \Sigma_\tau}{\partial \bar{\sigma}^2} > 0$ The rest follows from the definition of $\bar{\sigma}^2$.

regression R^2 are given by

$$R^2_\varepsilon \equiv \frac{\operatorname{Var}\tilde{\varepsilon}_{t|t}}{\operatorname{Var}\tilde{\varepsilon}_t} = 1 - \frac{\Sigma_\tau}{\Sigma_\tau + \bar{\sigma}^2} = 1 - R^2_\tau$$

$$\text{where} \quad \operatorname{Var}(\tilde{\tau}_t) \equiv \Sigma_\tau = \sigma^2_\eta + \rho^2\,\bar{\sigma}^2\,\frac{\Sigma_\tau}{\Sigma_\tau + \bar{\sigma}^2} = \frac{\sigma^2_\eta}{1 - \rho^2\,\frac{\bar{\sigma}^2}{\Sigma_\tau + \bar{\sigma}^2}} \qquad \text{and} \quad \bar{\sigma}^2 \equiv \frac{f^2_\varepsilon}{f^2_\tau}\,\sigma^2_\varepsilon \tag{34}$$

The solution of the Riccati equation (34) for Σ_τ is bounded between σ^2_η and the AR(1) variance of τ_t.[28] The more persistent τ_t the more exceeds its innovation variance Σ_τ the scale of its structural shocks, σ^2_η.

Expectations are rational. If policy is bold with respect to the persistent target ($f_\tau \to 1$), the public will be more responsive to policy innovations when updating its belief $\tau_{t|t}$. K_τ increases with f_τ (holding f_ε fixed) and vice versa for K_ε. The optimal projection coefficients are

$$K_\tau = \frac{1}{f_\tau}\frac{\Sigma_\tau}{\Sigma_\tau + \bar{\sigma}^2} \quad \text{and} \quad K_\varepsilon = \frac{1}{f_\varepsilon}\frac{\bar{\sigma}^2}{\Sigma_\tau + \bar{\sigma}^2} \tag{35}$$

5.3. *Optimal Policy Outcomes*

This section compares optimal policies under commitment and discretion. 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. Given the limited range of shocks considered, the calibration is not designed to match the level of fluctuations in inflation and activity observed in the data.

5.3.1. *Commitment*

Figure 2 documents impulse responses under commitment. A striking result is that the policy-maker actually overshoots in his policy response to the transitory target, $f_\varepsilon > 1$. Compared to full information, this is surprising, since in this setting the impact coefficients are naturally bounded between zero and one. The output term in the loss function penalizes overshooting the target as much

[28]The fixed point of the Riccati equation is unique and stable. The slope of $\sigma^2_\eta + \rho^2\,\bar{\sigma}^2\,\Sigma_\tau/(\Sigma_\tau + \bar{\sigma}^2)$ with respect to Σ_τ equals $\left(\rho\bar{\sigma}^2/(\Sigma_\tau + \bar{\sigma}^2)\right)'^2$ and is everywhere below one for $|\rho| < 1$.

Table 1: Calibration of New Keynesian Model

		Private Sector Parameters
β	0.99	Time preference
σ	1.00	Risk Aversion / Inverse EIS
θ	0.75	Calvo Probability of not repricing
κ	0.1717	Slope of Phillips Curve: $\kappa = (1-\theta)\cdot(1-\beta\cdot\theta)/\theta\cdot(\sigma+\phi)$
ϕ	1.00	Inverse Frisch Labor Elasticity (used for κ)
		Policy Preferences
α_x	1.00	Relative weight on output stabilization is $\alpha_x/(1+\alpha_x)$
		Driving Processes
σ_ε	1.00	Volatility of iid target component ($\bar{x}_t = \tau_t + \varepsilon_t$ with $\varepsilon_t \sim N(0,\sigma_\varepsilon^2)$)
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. $\sigma_{\bar{x}} \equiv \sigma_\varepsilon^2 + \sigma_\tau^2/(1-\rho^2)$

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.

as undershooting it, and the inflation term unambiguously penalizes the extra inflation resulting from overshooting.

Under imperfect information, an impact coefficient on the iid target larger than one is however rational. Overreacting to iid shocks decreases the public's signal to noise ratio and lowers the reaction of inflation expectation to either shock, since the public places less odds on a given output innovation being persistent.

This effect is fairly persuasive as can be seen from a sensitivity analysis of optimal policy to changes in the policy parameters. Policy trade-offs are particularly affected by two parameters: The relative variance of transitory to persistent target shocks and the preference weight α_x.[29] When varying the latter, the overall variance of the output target will be fixed at a given level $\sigma_{\bar{x}}^2$. Denoting the weight on τ_t by $\omega \in [0;1]$ this translates into

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

[29] Increases in the slope of the Phillips Curve, κ, affect policy trade-offs similarly to decreases in α_x This can be seen from rewriting the loss function in terms of a normalized inflation rate 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 2: Optimal Commitment in the Simple Model

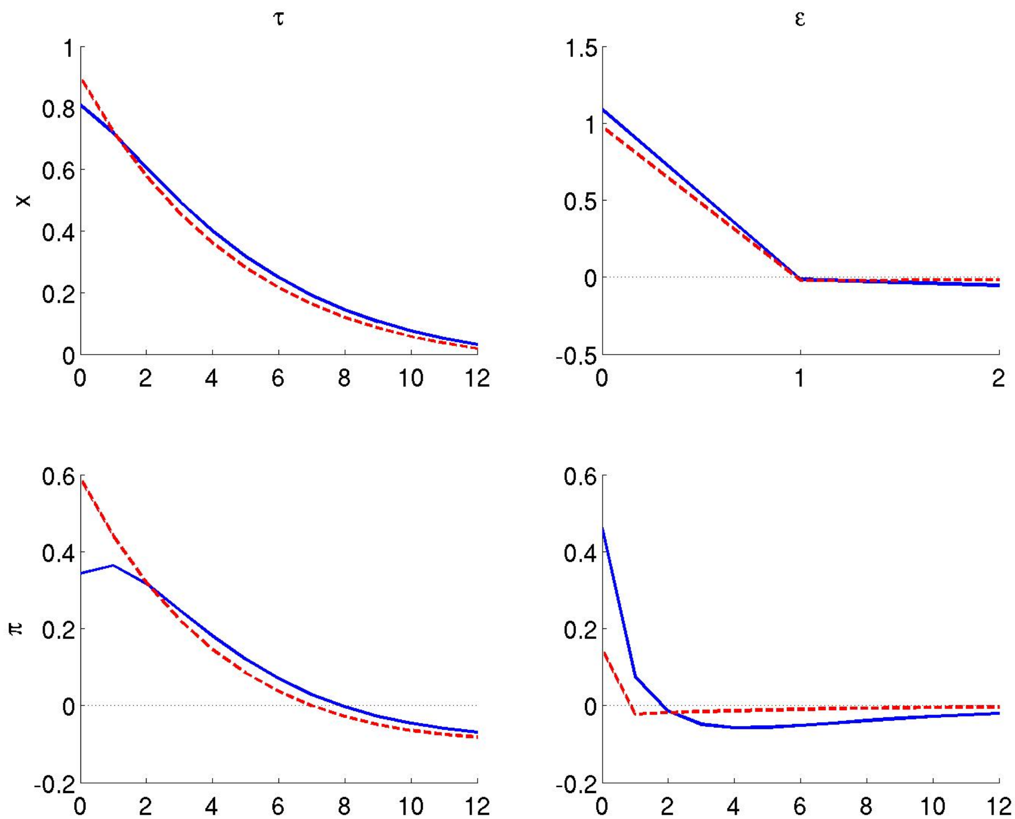

Note: Impulse responses of output gap (x_t) and inflation (π_t) under hidden (straight lines) and full information (dashed) to unit standard deviation shocks τ_t and ε_t.

Figure 3 documents changes in the policy coefficients f_τ and f_ε due to variations in ω and α_x. The corresponding values of $\bar{f}_\tau$ and $\bar{f}_\varepsilon$ under symmetric information are shown as well. Because of certainty equivalence, their surfaces are flat along the ω-axis.

The sensitivity analysis shows how innovation policy seeks to systematically lower the signal-to-noise ratio in the public's signal extraction problem. Panel (a) of Figure 3 shows how f_τ is generally smaller than $\bar{f}_\tau$ over a wide range of calibrations. The policymaker needs to scale back his pursuit of the persistent target, in order to lower the average persistence of a policy shock. In addition, he increases his response to the iid target f_ε. He even commits to overshoot this target for calibrations where variations in the persistent target are large ($\omega \to 1$, see Panel (b)). In doing so, he counteracts the otherwise large weight used by the public in updating $\tau_{t|t}$ in response to observing an innovation in the output gap.

By lowering inflation expectations, these effects lower realized inflation considerably, as can be seen in the case of the baseline calibration shown in Figure 2. When computing the unconditionally expected policy loss EV_t, the benefits of lower inflation significantly outweigh larger target misses. In the absence of analytical solutions, the generality of this result for a wider range of calibrations can only be verified numerically. Panel (a) of Figure 4 shows that policy losses are indeed reduced over a wide range of values for α_x and ω. A similar results will hold under the discretion policies discussed below, see Panel (b).

As an alternative measure one could also evaluate these policies under a loss function which seeks to minimize and inflation and output gaps. $L_t^R = \pi_t^2 + \alpha_x x_t^2$ In a similar setting with unobserved policy targets, Faust and Svensson (2001) use such a measure of "representative" or "welfare" losses.

Strikingly, under this measure the commitment policy fares worse when comparing hidden and full information. The overshooting of policy (f_ε) has benefits for inflation which are worthwhile when weighed against the goal of reaching the output target $\bar{x}_t$, but they are too small compared to a "representative" output target of zero. Under commitment, there is thus a tension in evaluating the benefits of hidden information when the policymaker's loss function should not be representative. This will be different under discretion, where policy losses largely stem from decreasing both f_τ and f_ε under hidden information, such that policy losses carry over to "welfare" losses.

5.3.2. *Discretion*

Under discretion 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:

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

The policy response to $\tau_{t|t-1}$ is synonymous with counterbalancing inflation expectations formed by the public in the past. This response is negative as optimal policy seeks to quell past beliefs. This can be seen from the impulse response shown in Figure 5. 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 to this initial condition 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 and Levin (2003).

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 et al. (2006).

The other lever of policy is the mixing ratio which is higher compared to the full information case.[30] 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.

However, these effects are only a by-product of period-by-period policy choices treating the public's Kalman filter as a given and ignoring the equilibrium effects of policy on the belief system. Compared to the commitment case, the discretionary policy coefficients f_τ and f_ε do not exceed unity. Absent the kind of systematic considerations made by the policymaker under commitment, it is not appealing for the discretionary policymaker to overshoot a target, like ε_t, as it is done under commitment. This underscores also the time-inconsistency of the optimal innovation policy under

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

commitment. An overshooting of policy to the iid target might overall be beneficial as it lowers the sensitivity of beliefs about the persistent target to innovations in policy. But for a given belief system of the public, overshooting merely leads to extra inflation without any direct benefits.

Under discretion, both policy coefficients f_τ and f_ε are mostly smaller when comparing hidden against full information.[31] The main reason is the public's inability to distinguish between realizations in the two target components, leading them to expect some persistence in output gaps and thus inflation in response to any output innovation. (A more nuanced discussion can be found in Mertens (2009).) Scaling back the pursuit of output targets under discretion vastly reduces inflation and leads to better policy losses under hidden information as shown in Panel (b) of Figure 4. Since this result is mostly achieved by producing lower output gaps, this result also carries over to the measure of "welfare" losses discussed above.

Under perfect information, discretion policies are myopic in this model, since the only state variables are exogenous. Managing beliefs adds history dependence to policies and removes a good part of the losses incurred under discretion as opposed to commitment. This is quantified in Panel (d) of Figure 4. Denoting the difference in losses between discretion and commitment under full information by $\left(E(V^{D,F}) - E(V^{C,F})\right)$, the figure plots its ratio against the corresponding difference in losses incurred under hidden information. For the wide range of calibrations considered here, this ration is generally larger than one, implying larger gains from commitment when there is full information. In particular, when the shocks to the output target are mostly persistent ($\omega \to 1$) and when there is a large weight α_x on attaining the output target, the gains from commitment under full information can be up to eight times the corresponding gains under hidden information.

To the extent that imperfect information might be a more plausible condition of actual policymaking, this suggests that the disadvantages from a lack of commitment might be smaller than what would be suggested by an analysis based on the assumption of perfect information.

[31] The discretion solution under full information can be computed here analogously to Clarida et al. (1999) as

$$x_t = \frac{\alpha_x(1-\beta\rho)}{\kappa^2 + \alpha_x(1-\beta\rho)}\,\bar{x}_t \equiv \bar{f}\,\bar{x}_t \quad \text{and} \quad \pi_t = \frac{\kappa}{1-\beta\rho}\,\bar{f}\,\bar{x}_t \tag{36}$$

The general case for linear quadratic models is described by Söderlind (1999) and Svensson (2007).

Figure 3: Sensitivity Analysis of Commitment Policy

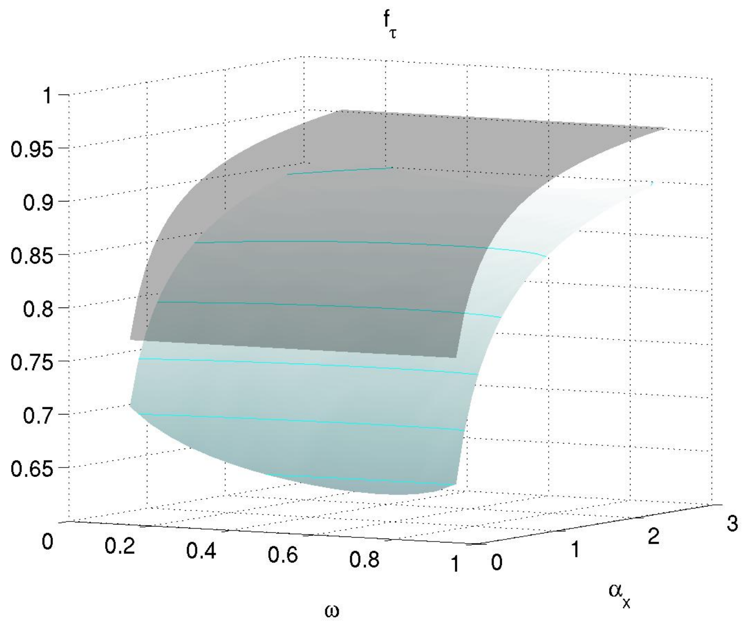

(a) f_τ

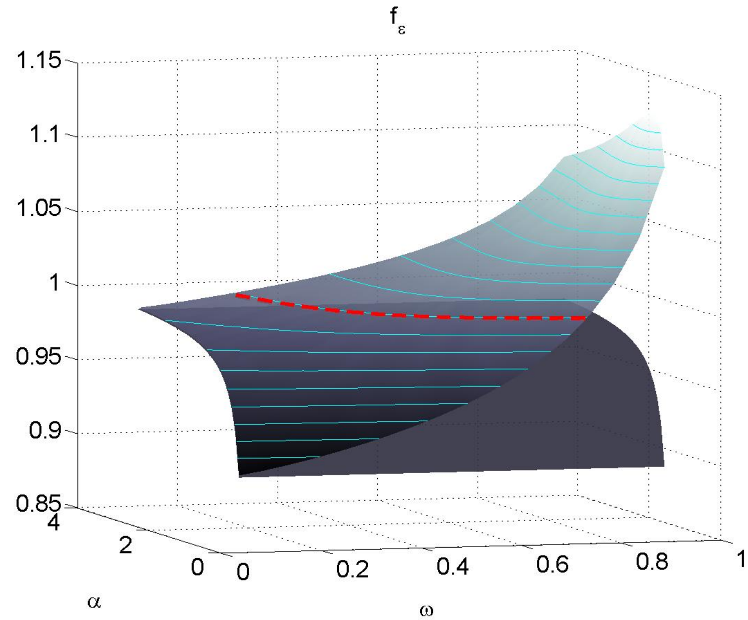

(b) f_ε

Note: Sensitivity of policy function due to variations in preference weight α_x and variance weight σ_ε in the simple New Keynesian model, keeping other parameters at their baseline values from Table 1. In Panel (a), the upper surface, with transparent grey shading, depicts the corresponding variations in $\bar{f}_\tau$ and in Panel (b), the lower surface in dark grey depicts $\bar{f}_\varepsilon$ from the symmetric information model. In Panel (b), the isosurface line for values of f_ε equal to one is depicted with a dashed line (in red). In order to provide the best perspective on each surface, axes are rotated differently in each panel.

Figure 4: Comparison of Policy Losses: Full vs Hidden Info

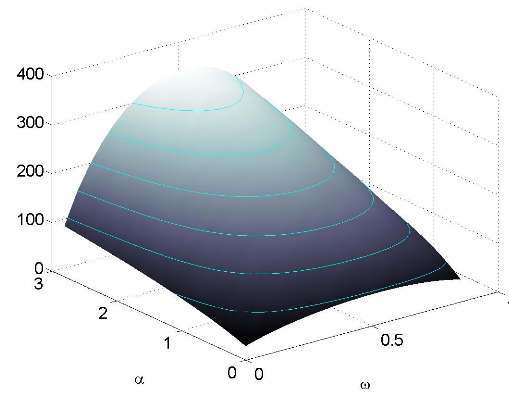

(a) Commitment

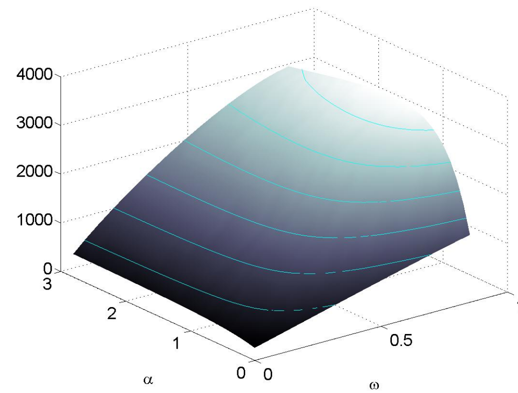

(b) Discretion

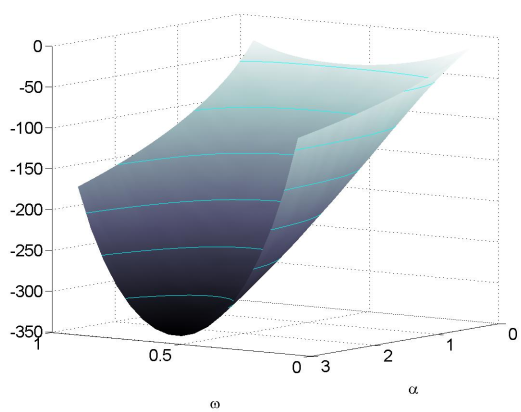

(c) Welfare Losses under Commitment

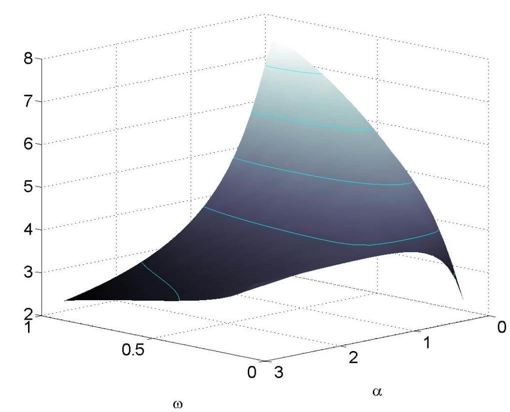

(d) Gains from Commitment

Note: Each panel compares measures of unconditional losses obtained under full versus hidden information. Panel (a) reports the difference between commitment losses under full and hidden information. All numbers are positive, meaning that losses are always larger under full information. Panel (b) reports the corresponding numbers under discretion. Panel (c) compares losses under commitment for an alternative loss function, neglecting the exogenous output target. $L_t^R = \pi_t^2 + \alpha_x x_t^2$. Panel (d) Reports the ratio of differences in losses between discretion and commitment under full and hidden information $\left(E(V^{D,F}) - E(V^{C,F})\right) / \left(E(V^{D,H}) - E(V^{C,H})\right)$ where the superscripts D and C stand for discretion and commitment, and H and F stand for hidden and full information. All numbers are larger than one, meaning that the advantage from commitment is larger under full information.

In order to provide the best perspective on each surface, axes are rotated differently in each panel.

Figure 5: Optimal Discretion Policy in the Simple Model

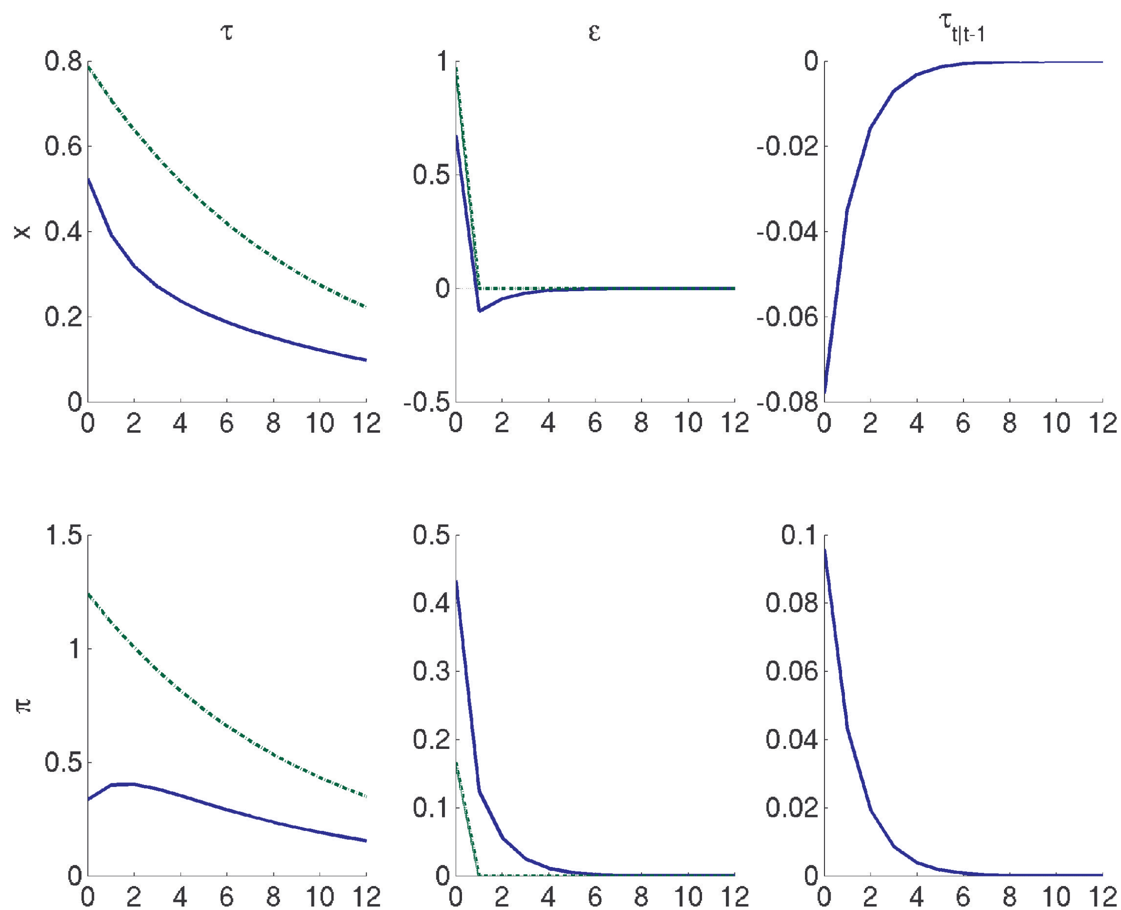

Note: Impulse responses of output gap (x_t) and inflation (π_t) under hidden (straight lines) and full information (dashed) to unit standard deviation shocks τ_t and ε_t. The column labeled $\tau_{t|t-1}$ shows responses to initial conditions $\tau_t = 0, \varepsilon_t = 0, \tau_{t|t-1} = 1$.

6. RELATED LITERATURE

Asymmetric information is such a pertinent issue in policymaking, 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*? Are policy decisions based on explicit optimization problems or merely represented by a fixed rule? In this paper, the public faces hidden information and policy is optimal. The tractability of the solution stems from the unobservable states following smooth, Gaussian processes as opposed to regime switches. A particular novelty is the solution of the commitment problem, which explicitly accounts for the policymaker's *systematic* manipulation of the public's information set.

Discrete regime switches are attractive for modeling central bank "types" like weak/soft or commitment/discretion as in Backus and Driffill (1985), Cukierman and Liviatan (1991), Ball (1995), Walsh (2000) and King et al. (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 2003a; 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 et al. (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.

Papers closest to the simple model illustrated studied in Section 5 are Cukierman and Meltzer (1986) and Faust and Svensson 2001; 2002. However, these paper consider only discretion equilibria.[32] This paper shares with them the linear framework and the Kalman filtering of the public. What is new about the discretionary policies described in Section 4 is the capability of handling various models with endogenous state variables as illustrated in a companion paper (Mertens, 2009).

7. CONCLUSIONS

This paper analyzes optimal policies of a policymaker who is better informed than the public about the realization of structural shocks. General solution methods are presented under discretion and commitment for a general class of linear quadratic models. A key result is the emergence of public beliefs as endogenous state variable. Novel compared to previous work is the systematic manipulation of public information under commitment.

The methods are illustrated with a very simple New Keynesian model, where optimal policy delivers better outcomes under hidden information than under full information. Given the emergence of beliefs as a backward looking state variable, this result might not be surprising for discretion, which would otherwise be fraught by myopic behavior. Particularly interesting is how this result extends to the commitment setting.

[32] Faust and Svensson 2001; 2002 also consider "commitments", but these pertain only to a policymaker's choice of noise variance. These do not account for his systematic influence on public information as it is done here.

Signal extraction dampens the people's response to news towards their unconditional, average expectation about the persistence of a shock. In the simple model analyzed here, this leads to inflation rates which are either too high or too low compared to their full information levels. But for the purpose of reducing price dispersion, it is preferable if inflation reactions do not oscillate between the extremes of observing whether an expansion will be persistent or transitory *for sure*.

This channel of "cross-shock smoothing" appears new and interesting. In the present model, expectational effects are only relevant for price setting, without feedback for example to investment. In a larger model, interesting trade-offs should emerge, leading to more nuanced results about the benefits of transparency in policymaking.

An important extension for future research would be to embed the methods presented here into a setting of mutually imperfect information, where neither policymaker nor public can observe the true state. Realistically, this will require to drop the assumption of a homogeneously informed public. In order to provide better microfoundations to the policy problem, it would also be important to restate the linear quadratic equations used here as approximate equilibrium conditions to a non-linear model.

A. KALMAN FILTER

Both under commitment and discretion, policy innovations have the form

$$\tilde{U}_t = F_x \tilde{X}_t \tag{19}$$

Together with with (1) and (2), this yields the following innovations representation of state and measurement equations

$$\tilde{X}_{t+1} = \underbrace{(A_{xx} + B_x F_x)}_{\equiv A}(X_t - X_{t|t}) + D w_{t+1}$$
$$\tilde{Z}_t = \underbrace{(C_x + C_u F_x)}_{\equiv C} \tilde{X}_t$$

and beliefs evolve as $X_{t|t} = X_{t|t-1} + K_x \tilde{Z}_t$, with Kalman gain $K_x = \text{Cov}\left(X_t, \tilde{Z}_t\right) \ \text{Var}\left(\tilde{Z}_t\right)^{-1}$.

For any variable v_t, innovations are defined as surprises relative to the public information set, $\tilde{v}_t = v_t - v_{t|t-1}$. The Kalman gain is identical to the coefficients of a least squares projection of X_t on $\tilde{Z}_t$. The presence of private sector controls $Y_{t|t}$ and predetermined variables $X_{t|t-1}$ in the transition equation for X_t does not affect the Kalman gain which depends only on the policy coefficients F_x, via which policy reacts to $\tilde{X}_t$.

The optimal Kalman gain is $K_x = \Sigma C' \left(C\, \Sigma C'\right)^{-1}$ where Σ solves the Riccati equation $\Sigma = A\Sigma A' + DD' - A\Sigma C' \left(C\Sigma C'\right)^{-1} C\Sigma A'$. The Kalman filter also implies that the policymaker can forecast future innovations using $E_t \tilde{X}_{t+1} = (A_{xx} + B_x F_x)(I - K_x C)\tilde{X}_t$.[33]

The above assumes that the $N_z \times N_x$ matrix C has full row rank. (Recall that $N_z < N_w \leq N_x$.) In principle (and also in practice) it can happen that C is collinear for some F_x. Numerically it is already critical if C is nearly collinear. This corresponds to situations when there are multiple observables 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 F_x tries to mimic other signals in Z_t. In these

[33] These are closed-loop dynamics for the policy rule given in (19). Equation (48) in Section B.2 below describes the open-loop dynamics of $\tilde{X}_t$.

cases, the Kalman filter can be implemented by pruning the redundancies in the set of observable variables via a singular value decomposition of C.

B. THE SADDLEPOINT PROBLEM UNDER COMMITMENT

This appendix derives the first-order conditions of the saddlepoint problem under commitment, discussed in Section 3. With $L_t^{MM} \equiv L_t + \beta^{-1}\lambda_{t-1|t-1}A_{yy}^1 Y_{t|t-1}$, the Bellman equation is

$$
\begin{aligned}
V_t^{MM} = \min_{U_t, X_{t+1}, Y_{t|t}, Z_t} \; \max_{\mu_{t+1}, \lambda_{t|t}, \zeta_t} \quad & L_t^{MM} + \beta E_t V_{t+1}^{MM} \\
& +\beta\mu'_{t+1}\left[X_{t+1} - A_{xx}X_t - A_{xy}Y_{t|t} - B_x U_t - D w_{t+1}\right] \\
& -\lambda'_{t|t}\left[A_{yy}Y_{t|t} + A_{yx}X_{t|t} + B_y U_{t|t}\right] \\
& +\zeta'_t\left[Z_t - C_x X_t - C_u U_t\right]
\end{aligned}
\tag{12}
$$

Since the problem is linear-quadratic, the value function is linear-quadratic as well

$$
V_t^{MM} = S_t' \boldsymbol{V} S_t + v \tag{37}
$$

with state vector

$$
S_t = \begin{bmatrix} \tilde{X}_t \\ X_{t|t-1} \\ \lambda_{t-1|t-1} \end{bmatrix}
$$

and a conformably partitioned value matrix

$$
\boldsymbol{V} = \begin{bmatrix} \tilde{V}_{xx} & \dots & \dots \\ \vdots & \bar{V}_{xx} & \bar{V}_{x\lambda} \\ \vdots & \bar{V}'_{x\lambda} & \bar{V}_{\lambda\lambda} \end{bmatrix}
$$

Together with the constraints (1)-(3), the first-order conditions for the saddlepoint problem are

$$X_{t+1}: \quad \tilde{V}_{xx}E_t\tilde{X}_{t+1} + \bar{V}_{xx}X_{t+1|t} + \bar{V}_{x\lambda}\lambda_{t|t} + E_t\mu_{t+1} = 0 \tag{38}$$

$$Y_{t|t}: \quad Q_{yy}Y_{t|t} + Q'_{xy}X_{t|t} + N_yU_{t|t} + \beta^{-1}A^1_{yy}\lambda_{t-1|t-1} - \beta A'_{xy}\mu_{t+1|t} - A_{yy}\lambda_{t|t} = 0 \tag{39}$$

$$U_t: \quad N'_yY_{t|t} + N'_xX_t + RU_t - \beta B'_xE_t\mu_{t+1} - B'_y\lambda_{t|t} - C_u\zeta_t = 0 \tag{40}$$

$$Z_t: \quad \Omega_\mu E_t\tilde{\mu}_{t+1} + \Omega_x(X_t - X_{t|t}) + \Omega_u(U_t - U_{t|t}) - \beta\Omega_{xx}E_t\tilde{X}_{t+1} = \zeta_t \tag{41}$$

where the coefficients of the measurement FOC (41) are

$$\Omega_\mu = \beta K'_yA'_{xy} \qquad \Omega_x = K'_\lambda A_{yx} - K'_yQ'_{xy}$$

$$\Omega_u = K'_\lambda B_y - K'_yN'_y \qquad \Omega_{xx} = K'_{xx}(\bar{V}_{xx} - \tilde{V}_{xx}) + K'_\lambda\bar{V}'_{x\lambda}$$

They depend on the yet to be determined Kalman gains

$$K_y = \operatorname{Cov}(Y_{t|t}, \tilde{Z}_t)\operatorname{Var}(\tilde{Z}_t)^{-1} \qquad K_\mu = \operatorname{Cov}(\mu_{t+1}, \tilde{Z}_t)\operatorname{Var}(\tilde{Z}_t)^{-1}$$

$$K_\lambda = \operatorname{Cov}(\lambda_{t|t}, \tilde{Z}_t)\operatorname{Var}(\tilde{Z}_t)^{-1} \qquad K_{xx} = \operatorname{Cov}(X_{t+1}, \tilde{Z}_t)\operatorname{Var}(\tilde{Z}_t)^{-1}$$

Details on the functional derivatives involved in differentiating terms involving the public information set, like $X_{t|t}$, are described below in Section B.3.

Proposition 3 is easily verified by conditioning down (41) onto Z^t and noting that $\tilde{v}_{t+1|t} = 0$ for any variable v_t.

B.1. *Public Information Policy*

The public information policy described in Proposition 2 is obtained by conditioning down the system of first order conditions (38)-(41) onto the public information set Z^t. The measurement FOC (41) drops from the system. As a corollary of Proposition 3, all variables in (41) are orthogonal to Z^t. The remaining system yields

$$0 = Q_{xx}X_{t+1|t} + Q_{xy}Y_{t+1|t} + N_x U_{t+1|t} + \mu_{t+1|t} \tag{42}$$

$$0 = Q_{yy}Y_{t|t} + Q'_{xy}X_{t|t} + N_y U_{t|t} + \beta^{-1}A^1_{yy}\lambda_{t-1|t-1} - \beta A'_{xy}\mu_{t+1|t} - A_{yy}\lambda_{t|t} \tag{43}$$

$$0 = N'_y Y_{t|t} + N'_x X_{t|t} + RU_{t|t} - \beta B'_x \mu_{t+1|t} - B'_y \lambda_{t|t} \tag{44}$$

$$X_{t+1|t} = A_{xx}X_{t|t} + A_{xy}Y_{t|t} + B_x U_{t|t} \tag{45}$$

Equation (42) is obtained from substituting the envelope theorem into (38), and (45) forms public expectations of future backward looking variables according to (1). Since the public sector choices are constrained to lie in the public information set, the associated first-order condition lies already in Z^t and (43) is identical to (39). Apart from the projection of all variables onto Z^t, this system is identical to the FOC of the commitment problem under full and symmetric information described in Section C. The solution has the form shown in Proposition 2, and the coefficients are certainty equivalent and identical to the full information problem.

B.2. *Innovation Policy*

Public information policy solves the projection of first-order conditions (38)-(41) as well as the constraints (1)-(3) onto Z^t. Innovation policy solves the remainder which lies in the part of w^t orthogonal to Z^t:

$$0 = \tilde{V}_{xx}E_t\tilde{X}_{t+1} + E_t\tilde{\mu}_{t+1} \tag{46}$$

$$0 = N'_x(X_t - X_{t|t}) + R(U_t - U_{t|t}) - \beta B'_x E_t\tilde{\mu}_{t+1} - C_u\zeta_t \tag{47}$$

$$\zeta_t = \Omega_\mu E_t\tilde{\mu}_{t+1} + \Omega_x(X_t - X_{t|t}) + \Omega_u(U_t - U_{t|t}) - \beta\Omega_{xx}E_t\tilde{X}_{t+1} \tag{41}$$

$$E_t\tilde{X}_{t+1} = A_{xx}(X_t - X_{t|t}) + B_x(U_t - U_{t|t}) \tag{48}$$

Notice how the forward looking constraint (3) and the associated first-order condition (39) have dropped from the problem since they lie completely in Z^t. Similarly, the third equation reproduces the original first-order condition for Z_t which is completely orthogonal to Z^t (Proposition 3). The last equation projects the transition equation (1) off Z^t.

This system can be transformed into

$$\begin{aligned} 0 &= \Phi_x(X_t - X_{t|t}) + \Phi_u(U_t - U_{t|t}) \\ \text{with} \quad \Phi_x &= N_x' + \beta B_x' \tilde{V}_{xx} A_{xx} - C_u' \Psi_x \\ \Phi_u &= R \;\; + \beta B_x' \tilde{V}_{xx} B_x - C_u' \Psi_u \end{aligned}$$

and where Ψ_x and Ψ_u are the coefficients in

$$\begin{aligned} \zeta_t &= \Psi_x(X_t - X_{t|t}) + \Psi_u(U_t - U_{t|t}) \\ \text{with} \quad \Psi_x &= \Omega_x - \beta \Psi_{xx} A_{xx} = K_\lambda' A_{yx} - K_y' Q_{xy}' - \beta \Psi_{xx} A_{xx} \\ \Psi_u &= \Omega_u - \beta \Psi_{xx} B_x = K_\lambda' B_y - K_y' N_y' - \beta \Psi_{xx} B_x \\ \Psi_{xx} &= \Omega_{xx} + \Omega_\mu = K_{xx}'(\bar{V}_{xx} - \tilde{V}_{xx}) + K_\lambda' \bar{V}_{x\lambda}' + K_y' A_{xy}' \tilde{V}_{xx} \end{aligned}$$

Together with the Kalman Filter $\tilde{X}_{t|t} = K_x \tilde{Z}_t$, $\tilde{U}_{t|t} = \bar{F}_x K_x \tilde{Z}_t$, and measurement innovations $\tilde{Z}_t = C_x \tilde{X}_t + C_u \tilde{U}_t$ this can be solved for

$$\begin{aligned} \Phi_x \tilde{X}_t + \Phi_u \tilde{U}_t = (\Phi_x + \Phi_u \bar{F}_x) K_x \tilde{Z}_t &= \Phi_Z \tilde{Z}_t \\ \tilde{U}_t = -\left(\Phi_u - \Phi_z C_u\right)^{-1} \left(\Phi_x - \Phi_z C_x\right) \tilde{X}_t &= F_x \tilde{X}_t \end{aligned} \tag{21}$$

The coefficients in (19) depend on the Kalman gain K_x and the value matrix $\tilde{V}_{xx}$ which in turn depend on the optimal policy F_x. The solution seeks for a fixed point mapping F_x in itself. A convenient and intuitive algorithm for doing this is the policy improvement algorithm which is described in Section E.

B.3. *Differentiating Conditional Expectations*

This section shows how to differentiate terms involving projections onto the public information set. The usage of functional derivatives follows Aoki (2006) and the related discussion in Svensson and Woodford (2003b). Aoki's presentation is extended here to a multivariate case.

Considering an expression like $X_t' \, Q_{xy} \, Y_{t|t}$, care needs to be taken when perturbing either the variables to be projected, Y_t, or the measurement variables Z_t, which define the space onto which Y_t is projected. This appendix considers both cases and derives the following expressions, which correspond to (10) and (11) of Section 3.

$$\frac{\partial}{\partial Y} \left(X_t' \, Q_{xy} \, Y_{t|t} \right) = Q_{xy}' \, X_{t|t}$$
$$\frac{\partial}{\partial Z} \left(X_t' \, Q_{xy} \, Y_{t|t} \right) = K_y' \, Q_{xy}' \, (X_t - X_{t|t}) + K_x' \, Q_{xy} \, \underbrace{(Y_t - Y_{t|t})}_{=0}$$

where K_x and K_y are the Kalman gains involved in projecting X_t, respectively Y_t, onto the public information set.[34]. Similar expressions hold for other terms in (37) involving expectations conditional on Z^t.[35]

B.3.1. *Probability Space*

Formally, $X_t = X(w^t)$, $Y_{t|t} = Y(Z^t)$ and $Z_t = Z(w^t)$ are functions of the realized history w^t of stochastic events. A functional derivative considers the effects of state-by-state perturbation of these functions on terms like $X_t' \, Q_{xy} Y_{t|t}$. Think of perturbing a simple product of two scalars, say $a(w^t)b(w^t)$ by changing $a(w^t)$ in one particular state denoted $\bar{w}^t$. The marginal effect is $b(\bar{w}^t)$. Summarizing this results over all states, one usually denotes the marginal effect by b_t. A functional derivatives does the same while explicitly accounting for the underlying probability space, which matters, once one differentiates inside a conditional expectations.

First, we the underlying probability space needs to be considered. Let H_t be the set of all possi-

[34]As in (4) this means $K_y = \text{Cov}\,(Y_t, \tilde{Z}_t) \ \text{Var}\,(\tilde{Z}_t)^{-1}$ and $K_x = \text{Cov}\,(X_t, \tilde{Z}_t) \ \text{Var}\,(\tilde{Z}_t)^{-1}$.

[35]A particular property of the two examples given here is that $Y_t = Y_{t|t}$. The derivatives hold however also in the general case where none of the variables lies in Z^t. Such a situation could occur if for example projections $X_{t|t}$ would directly enter the loss function (8) or the transition equation (1).

ble histories w^t up and until time t. There is an exogenously given probability density defined over H_t and denoted $f(w^t)$ with cumulative density $F(w^t)$ and differential $dF(w^t)$.[36] The unconditional expectation of a random function like $X(w^t)$ is then the integral $E\,x_t = \int_{H_t} x(w^t)dF(w^t)$.

A conditional expectations like $X_{t|t}$ can be represented as the unconditional expectation of the product between X_t and a stochastic projection kernel $\kappa(\cdot)$

$$E\left[X_t | Z(w^t) = Z(\bar{w}^t)\right] = E\left[\kappa(Z(w^t), Z(\bar{w}^t))\, X_t\right]$$

where $Z(\bar{w}^t)$ denotes the particular realization of the measurement vector on which $X_{t|t}$ is to be conditioned. An important property of the projection kernel for multivariate Normal variables is its symmetry: $\kappa(Z(w^t), Z(\bar{w}^t)) = Z(w^t)' \operatorname{Var}(Z(w^t))^{-1} Z(\bar{w}^t)$ which will be important below.

B.3.2. *Definition of Functional Derivative*

The functional derivative of a functional $g(X(w^t))$ is defined as

$$\frac{\partial g}{\partial X} \equiv \frac{\partial E(g(X(w^t)))}{\partial X} = \frac{\partial}{\partial X} \int_{H_t} g(X(w^t); w^t) dF(w^t)$$

Defining the derivative as differentiating the *expectation* of the functional, means that the state-by-state derivatives are normalized by their density. For example, let g be the above discussed product, $a_t b_t = g(a_t, b_t)$. As expected, the functional derivative yields

$$\frac{\partial g}{\partial a} = \frac{\partial}{\partial a} \int_{H_t} a(w^t) b(w^t) dF(w^t) = b(w^t) = b_t$$

[36] Since the histories follow Markov processes we also have $f(w^t|w^{t-1}) = f(w^t)/f(w^{t-1})$.

B.3.3. *Derivative with respect to Y_t – Equation (10)*

Applying the functional derivative to (10) works as follows:

$$\begin{aligned}\frac{\partial}{\partial Y}\ \left(X_t' \, Q_{xy} Y_{t|t}\right) &= \frac{\partial}{\partial Y}\ \int_{H_t} X(w^t)' \, Q_{xy}\ \underbrace{\left\{\int_{H_t} \kappa(\bar{Z}^t, Z^t)\, Y(\bar{w}^t)\ \right\}}_{Y_{t|t}} dF(w^t) \\ &= Q_{xy}' \int_{H_t} X(w^t)\kappa(\bar{Z}^t, Z^t) dF(w^t) \\ &= Q_{xy}' \, X_{t|t}\end{aligned}$$

where $\bar{Z}^t$ denotes a particular realization of the history of public information consistent with a particular history of exogenous events, $\bar{w}^t$.

B.3.4. *Derivative with respect to Z_t – Equation (11)*

Differentiating $X_t \, Q_{xy} \, Y_{t|t}$ with respect to the vector of measurement variables involves the expectation

$$E\left(X_t' \, Q_{xy} Y_{t|t}\right) = \operatorname{tr}\left(A_{xx}^2 \, E\ \left(Y_{t|t} X_t'\right)\right) \tag{49}$$

(tr is the trace operator.) Substituting the Kalman filter (4) yields

$$\begin{aligned}E\ \left(Y_{t|t}\, X_t'\right) &= K_y \, E\left((Z_t - Z_{t|t-1})\ X_t'\right) + \text{terms independent of } Z_t \\ &= K_y \Sigma_Z K_x' + \text{terms independent of } Z_t\end{aligned}$$

where the Kalman Gains are defined above and $\Sigma_Z = \operatorname{Var}\left(Z_t - Z_{t|t-1}\right)$.

Differentiating the matrix product $K_y \Sigma_Z K_x'$ with respect to the vector Z_t, is easier conducted using differentials, $d_Z(K_x)$, $d_Z(K_y)$ and $d_Z(\Sigma_Z)$ (all understood with respect to perturbations in Z_t, denoted $d_Z(Z_t)$), than with Jacobian matrices. Since differentials have the same size as the original matrices, matrix manipulations can be carried out in the usual manner.[37] The Identification Theorem of Magnus and Neudecker (1988, Chapters 5 & 9) assures the identification of Jacobians

[37] K_x is an $N_x \times N_Z$ matrix. Its Jacobian with respect to Z_t is a $(N_x \cdot N_Z) \times N_Z$ matrix, see Magnus and Neudecker (1988, p. 170ff).

from differentials. This is particularly convenient when differentiating a scalar function of a vector, so the Jacobian one ultimately cares about is actually a vector.

The conventional product rule of differentiation can be extended to matrix differentiation (Magnus and Neudecker, 1988, Chapter 9):

$$d_Z(K_y \Sigma_Z K_y') = d_Z(K_y)\, \Sigma_Z\, K_x' + K_y\, d_Z(\Sigma_Z)\, K_x' + K_y\, \Sigma_Z\, d_Z(K_x')$$

and the functional differentiation gives

$$\begin{aligned}
d_Z(\Sigma_Z) &= d_Z(Z_t)\tilde{Z}_t' + \tilde{Z}_t d_Z(Z_t)' \\
d_Z(K_x) &= \left[\tilde{X}_t\, d_Z(Z_t)' - K_x d_Z(\Sigma_Z)\right] \Sigma_Z^{-1} \\
d_Z(K_y) &= \left[\tilde{Y}_t\, d_Z(Z_t)' - K_y d_Z(\Sigma_Z)\right] \Sigma_Z^{-1} \\
\Rightarrow d_Z(K_y \Sigma_Z K_x') &= (\tilde{Y}_t - K_y \tilde{Z}_t)\, d_Z(Z_t)'\, K_x' + K_y d_Z(Z_t)\, (\tilde{X}_t - K_x \tilde{Z}_t)' \\
&= (Y_t - Y_{t|t})\, d_Z(Z_t)'\, K_x' + K_y d_Z(Z_t)\, (X_t - X_{t|t})'
\end{aligned}$$

The result in (11) follows from the identification $d_Z(g(Z_t)) = \mathrm{tr}\,(A\ d_Z(Z_t)) \Rightarrow \frac{\partial g}{\partial Z} = A'$, for some matrix A (Magnus and Neudecker, 1988, Chapter 9).

C. COMMITMENT UNDER SYMMETRIC INFORMATION

When public and policymaker can both observe the full information set, w^t, their expectations coincide for every variable v_t, $v_{t+1|t} = E_t v_{t+1}$. The commitment problem solves the following Lagrangian

$$\mathcal{L} = E_{t_0} \sum_{t=0}^{\infty} \Big\{ \beta^t L_t + \beta^{t+1} \mu_{t+1}' \left[X_{t+1} - A_{xx} X_t - A_{xy} Y_t - B_x U_t - D w_{t+1}\right] \tag{50}$$

$$+ \beta^t \lambda_t' \left[A_{yy}^1 E_t Y_{t+1} - A_{yy} Y_t - A_{yx} X_t - B_y U_t\right] \Big\} \tag{51}$$

The FOC to (50) lead to a system of expectational difference equations which is identical to the problem of public information policy (42)-(44), except for the information set used in forming expectations.

$$0 = Q_{xx}E_tX_{t+1} + Q_{xy}E_tY_{t+1} + N_xE_tU_{t+1} + E_t\mu_{t+1} \tag{52}$$

$$0 = Q_{yy}Y_t + Q'_{xy}X_t + N_yU_t + \beta^{-1}A^1_{yy}\lambda_{t-1} - \beta A'_{xy}E_t\mu_{t+1} - A_{yy}\lambda_t \tag{53}$$

$$0 = N'_yY_t + N'_xX_t + RU_t - \beta B'_xE_t\mu_{t+1} - B'_y\lambda_t \tag{54}$$

$$E_tX_{t+1} = A_{xx}X_t + A_{xy}Y_t + B_xU_t \tag{55}$$

The system can be solved using the methods of King and Watson (1998) or Klein (2000). The solution has the form

$$U_t = \bar{F}_xX_t + \bar{F}_\lambda\lambda_{t-1}$$

$$Y_t = \bar{G}_xX_t + \bar{G}_\lambda\lambda_{t-1}$$

$$\lambda_t = \bar{H}_xX_t + \bar{H}_\lambda\lambda_{t-1}$$

In particular, the coefficients $\bar{F}_x$, $\bar{F}_\lambda$, $\bar{H}_x$ and $\bar{H}_\lambda$ are certainty equivalent. They do not depend on the shock loadings D in (1), nor on the information set used in forming expectations. They also solve the public information problem described in Proposition 2.

Equivalently, the symmetric information problem can be solved with a recursive representation analogous to (12). This involves a Riccati equation associated with the linear quadratic value function, see for example Svensson (2007). As a corollary of Proposition 2, it is straightforward to show that the matrices $\bar{V}_{xx}$, $\bar{V}_{x\lambda}$, $\bar{V}_{\lambda\lambda}$ of the value function (37) are identical to the corresponding submatrices of the recursive saddlepoint problem under full information. This will be computationally convenient when iterating over the fixed-point problem associated with the innovation policy under imperfect information described in Appendix E.

D. REGULATOR FOR DISCRETION PROBLEM

To set up the linear regulator problem shown in (25) and (26), the temporary equilibrium (23) and the Kalman filter (4) can be used to substitute $Y_{t|t}$ out of the loss function (8) 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 (23) and the Kalman filter (4) to substitute $Y_{t|t}$ out of the loss function (8) 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 the candidate policy F^0 from (24) (as it will be in equilibrium), such that

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

The Kalman update can be written as

$$X_{t|t} = K_x C_x X_t + (I - K\hat{C})X_{t|t-1} + K_x C_u U_t$$

Together with the temporary equilibrium (23) this yields

$$\begin{aligned} Y_{t|t} &= \Gamma_x^0 X_t + \hat{\Gamma}_x^0 X_{t|t-1} + \Gamma_u^0 U_t \\ \text{with} \quad \Gamma_x^0 &= (G_x^0 + G_u^0 F_x^0) K_x C_x \\ \hat{\Gamma}_x^0 &= (G_x^0 + G_u^0 F_x^0)(I - K\hat{C}) + G_u^0 F_b^0 \\ \Gamma_u^0 &= (G_x^0 + G_u^0 F_x^0) K_x C_u \end{aligned}$$

D.1. *Loss Function*

The loss function (8) can be rewritten in terms of the regulator's states and control using

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

such that

$$\begin{aligned} L_t &= \begin{bmatrix} X_t \\ Y_{t|t} \\ U_t \end{bmatrix}' Q \begin{bmatrix} X_t \\ Y_{t|t} \\ U_t \end{bmatrix} = \begin{bmatrix} S_t \\ U_t \end{bmatrix}' \boldsymbol{H^0}' Q \boldsymbol{H^0} \begin{bmatrix} S_t \\ U_t \end{bmatrix} \\ &= S_t' \boldsymbol{Q^0} S_t + 2 S_t' \boldsymbol{N^0} U_t + U_t' \boldsymbol{R^0} U_t \end{aligned}$$

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

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

D.2. *State Transition*

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

$$\begin{aligned} X_{t+1} &= (A_{xx} + A_{xy}\Gamma_x^0)X_t + A_{xy}\hat{\Gamma}_x^0 X_{t|t-1} + (A_{xy}\Gamma_u^0 + B_x)U_t + Dw_{t+1} \\ X_{t+1|t} &= A_{xx}^0 K_x C_x X_t + \left(A_{xx}^0(I - K\hat{C}) + (A_{xy}G_u^0 + B_x)F_b^0\right) X_{t|t-1} + A_{xx}^0 K_x C_u U_t \end{aligned}$$

where $\quad A_{xx}^0 = A_{xx} + A_{xy}(G_x^0 + G_u^0 F_x^0) + B_x F_x^0$

The matrices $\boldsymbol{A^0}$, $\boldsymbol{B^0}$ and $\boldsymbol{D}$ in (26) are thus given by:

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

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

E. POLICY IMPROVEMENT ALGORITHM

Equilibria under commitment and discretion are fixed points of public beliefs and policy actions. An intuitive and efficient way to compute this fixed point is the policy improvement algorithm. It is efficient, since policy improvement methods converge faster than value function iterations (Whittle, 1996; Bertsekas, 2005).[38] It is intuitive, since the algorithm seeks for deviations from a candidate equilibrium. Non-existence of such a deviation is the defining property of equilibrium.

Formally, the algorithm starts with a candidate policy F^0—and under discretion with beliefs G^0—and computes the Kalman gain K_x^0 and continuation value $\boldsymbol{V^0}$ associated with continuing this policy forever.[39] Section E.2 describes how to compute continuation values for given policy functions in detail.

Under discretion, the solution (27) to the regulator (25) yields the optimal one-period deviation. As long as $F^0 \neq F^*$ and $G^* \neq G^0$ there is no equilibrium. In this case, a new iteration starts using (F^*, G^*) as new candidate policies.

Under commitment, equation (21) provides the optimal innovation policy which must be consistent with the Kalman filter and continuation value assumed when formulating the first order conditions (46)-(48). The certainty equivalence of the public information policy means that the

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

[39] If the conditions for a private sector equilibrium are met (Proposition 1), one can even compute the G^0 consistent with F^0 under discretion.

public information policies ($\bar{F}_x$, $\bar{F}_\lambda$, $\bar{G}_x$, $\bar{G}_\lambda$, $\bar{H}_x$, $\bar{H}_\lambda$ in Proposition 2) and the associated submatrices of the values function ($\bar{V}_{xx}$, $\bar{V}_{x\lambda}$ and $\bar{V}_{\lambda\lambda}$), can be computed before the start of the policy improvement algorithm.

The difference between policy improvement algorithms 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_x^0 and continuation value $\boldsymbol{V^0}$, (and if applicable also G^0), are consistent with the candidate policy F^0.

E.1. *Uniqueness of Equilibrium*

The imperfect information equilibria considered here pose intricate fixed points between optimal policies and public beliefs. Formally, these are fixed points between two Riccati equations, one from the policymaker's regulator problem, the other associated with the public's Kalman Filter.

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. In my practical experience, the algorithm typically converges, and if so always to the same equilibrium from arbitrary starting values.

E.2. *Closed-Loop Value Function*

The policy improvement algorithm uses a continuation value consistent with carrying out a given policy F^0 forever. The remainder of this section describes the necessary computations under commitment and discretion. In both cases, the derivation proceeds by using the given policy function to spell out a "closed-loop" representation of state dynamics and loss function. The value function is then given by the present value of expected future losses. In the scalar case, this would correspond to solving a geometric series, in the matrix case, the computations generalize to the solution of a Lyapunov equation.

E.2.1. *Commitment*

Under commitment, policy is given by $U_t = F^0 S_t = F^0_x \tilde{X}_t + \bar{F}_x X_{t|t-1} + \bar{F}_\lambda \lambda_{t-1|t-1}$ where the coefficients $\bar{F}_x$ and $\bar{F}_\lambda$ are identified by the public information policy (15)-(17) and can be calculated before initializing the improvement algorithm. Only the coefficients F^0_x make up the "candidate" part of policy for which the algorithm has to find an improvement. The Kalman gain and measurement matrix associated with F^0_x will be denoted K^0_x and $C^0 = C_x + C_u F^0_x$ (see also Appendix A).

In closed-loop, the state dynamics are $S_{t+1} = \boldsymbol{A^0_S} S_t + \boldsymbol{D} w_{t+1}$, with

$$\boldsymbol{A^0_S} = \begin{bmatrix} (A_{xx} + B_x F^0_x)(I - K_x C^0) & 0 & 0 \\ \bar{P}_x K^0_x C^0 & \bar{P}_x (I - K^0_x C^0) & \bar{P}_\lambda \\ \bar{H}_x K^0_x C^0 & \bar{H}_x (I - K^0_x C^0) & \bar{H}_\lambda \end{bmatrix} \qquad \boldsymbol{D} = \begin{bmatrix} D \\ 0 \\ 0 \end{bmatrix}$$

Expressed as a function of the states, the modified loss function can be written as $L^{MM}_t = \frac{1}{2}\, S'_t \boldsymbol{Q^0_{SS}} S_t$ where

$$\boldsymbol{Q^0_{SS}} = \boldsymbol{H^0_S}' \begin{bmatrix} Q_{xx} & Q_{xy} & N_x & 0 \\ Q'_{xy} & Q_{yy} & N_y & \beta^{-1} A_{yy} \\ N'_x & N'_y & R & 0 \\ 0 & \beta^{-1} A'_{yy} & 0 & 0 \end{bmatrix} \boldsymbol{H^0_S}$$

$$\text{and} \quad \boldsymbol{H^0_S} = \begin{bmatrix} I & I & 0 \\ \bar{G}_x K^0_x C^0 & \bar{G}_x (I - K^0_x C^0) & \bar{G}_\lambda \\ F^0_x & \bar{F}_x & \bar{F}_\lambda \\ 0 & 0 & I \end{bmatrix}$$

Finally, the value matrix associated with our candidate policy F^0 for the recursive saddlepoint problem (37) solves the Lyapunov equation $\boldsymbol{V^0} = \boldsymbol{Q^0_{SS}} + \beta \boldsymbol{A^0_S}' \boldsymbol{V^0} \boldsymbol{A^0_S}$.

E.2.2. *Discretion*

The policy improvement algorithm uses a continuation value consistent with carrying out the policy F^0 forever. Under discretion, the continuation value is computed from the closed loop representation of the regulator obtained by plugging the policy F^0 into (25) and (26). $\boldsymbol{V^0}$ solves the Lyapunov equation

$$\begin{aligned}\boldsymbol{V^0} &= \left\{\boldsymbol{Q^0} + 2\boldsymbol{N^0}F^0 + F^{0\prime}\boldsymbol{R^0}F^0\right\} + \beta(\boldsymbol{A^0} + \boldsymbol{B^0}F^0)'\boldsymbol{V^0}(\boldsymbol{A^0} + \boldsymbol{B^0}F^0)\\ &= \boldsymbol{Q^0_{SS}} + \beta\boldsymbol{A^0_S}'\boldsymbol{V^0}\boldsymbol{A^0_S}\end{aligned}$$

The equation has a unique solution if $\boldsymbol{Q^0_{SS}}$ is positive definite and if the closed loop transition matrix $\boldsymbol{A^0_S} \equiv (\boldsymbol{A^0} + \boldsymbol{B^0}F^0)$ has all eigenvalues inside the unit circle. The form of the original loss function[40] assures us of the former, and the latter holds if a stationary equilibrium exists.[41]

E.2.3. *Expected Losses*

Both under commitment and discretion *unconditionally* expected losses are computed from the unconditional variance covariance matrix of the states:

$$\begin{aligned}E(V^0_t) &= \mathrm{tr}\,(\boldsymbol{V^0} E S_t S_t') + v^0\\ E S_t S_t' &= (\boldsymbol{A^0} + \boldsymbol{B^0}F^0)\,(E S_t S_t')\,(\boldsymbol{A^0} + \boldsymbol{B^0}F^0)' + DD'\end{aligned}$$

Optimal policies are certainty equivalent (for given F^0) and do not depend on v^0. For the purpose of evaluating expected losses, the scalar v^0 can be computed from:[42] $v^0 = \frac{\beta}{1-\beta}\,\mathrm{tr}\,(\boldsymbol{V^0}\boldsymbol{DD}')$.

[40] Please recall that Q in (8) is assumed to be positive definite.

[41] 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).

[42] tr denotes the trace operator.

REFERENCES

Burton A. Abrams. How richard nixon pressured arthur burns: Evidence from the nixon tapes. *Journal of Economic Perspectives*, 20(4):177–188, Fall 2006.

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

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

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

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

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

Robert J. Barro and David B. Gordon. Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12:101–121, July 1983.

Dimitri P. Bertsekas. *Dynamic Programming and Control*, volume I. Athena Scientific, Belmont, MA, 3rd edition edition, 2005.

Oliver Jean Blanchard and Charles M. Kahn. The solution of linear difference models under rational expectations. *Econometrica*, 48(5):1305–1312, July 1980.

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

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

Richard Clarida, Jordi Gali, and Mark Gertler. The science of monetary policy: A new keynesian perspective. *Journal of Economic Literature*, 37:1661–1707, December 1999.

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

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

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

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

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

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

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

Jordi Gali. New perspectives on monetary policy, inflation, and the business cycle. In Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, editors, *Advances in Economics and Econometrics: Theory and Applications*, Eighth World Congress of the Econometric Society, chapter 6, pages 151–197. Cambridge University Press, Cambridge, UK, 2003.

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

Jean Michel Grandmont. Temporary general equilibrium theory. *Econometrica*, 45(3):535–572, apr 1977.

Lars P. Hansen and Thomas J. Sargent. *Robustness*. Princeton University Press, 2007.

Lars Peter Hansen and Scott F. Richard. The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica*, 55(3):587–613, May 1987.

Robert G. King and Marvin Goodfriend. The new neoclassical synthesis and the role of monetary policy. In Ben S. Bernanke and Julio J. Rotemberg, editors, *NBER Macroeconomics Annual 1997*, pages 231–283. The MIT Press, Cambridge (MA), 1997.

Robert G. King and Mark W. Watson. The solution of singular linear difference systems under rational expectations. *Internatinal Economic Review*, 39(4):1015–1026, November 1998.

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

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

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

Finn E. Kydland and Edward C. Prescott. Dynamic optimal taxation, rational expectations and optimal control. *Journal of Economic Dynamics and Control*, 2(1):79–91, May 1980.

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

Jan R. Magnus and Heinz Neudecker. *Matrix Differential Calculus with Applications in Statistics and Econometrics*. Wiley Series in Probability and Statistics. John Wiley & Sons, 1988.

Albert Marcet and Ramon Marimon. Recursive contracts. Economics Working Papers 337, Department of Economics and Business, Universitat Pompeu Fabra, October 1998.

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

Albert Marcet and Thomas J Sargent. Convergence of least-squares learning in environments with hidden state variables and private information. *Journal of Political Economy*, 97(6):1306–22, December 1989b.

Elmar Mertens. Managing beliefs about monetary policy under discretion. *mimeo*, Federal Reserve Board, February 2009.

Athanasios Orphanides. Monetary policy rules based on real-time data. *American Economic Review*, 91(4):964–985, September 2001.

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

Athanasios Orphanides and John C. Williams. Monetary policy with imperfect knowledge. *Journal of the European Economic Association*, 4(2-3):366–375, 04-05 2006.

Joseph Pearlman, David Currie, and Paul Levine. Rational expectations models with partial information. *Economic Modelling*, 3(2):90–105, April 1986.

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

Kenneth Rogoff. Reputation, coordination and monetary policy. In Robert J. Barro, editor, *Modern Business Cycle Theory*, chapter 6. Harvard University Press, Cambridge, MA, 1989.

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

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

Thomas J. Sargent. *Macroeconomic Theory*. Economic Theory, Econometrics, and Mathematical Economics. Academic Press, New York, 2nd edition, 1987.

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

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

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

Lars E. O. Svensson and Michael Woodford. Optimal policy with partial information in a forward-looking model: Certainty-equivalence redux. NBER Working Papers 9430, National Bureau of Economic Research, Inc, January 2003b.

Lars E.O. Svensson. Optimization under commitment and discretion, the recursive saddlepoint method, and targeting rules and instrument rules: Lecture notes. *mimeo*, June 2007.

Lars E.O. Svensson and Noah Williams. Bayesian and adaptive optimal policy under model uncertainty. *mimeo* Princeton University, November 2006.

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

Carl E. Walsh. *Monetary Theory and Policy*. The MIT Press, Cambridge, MA, 2nd edition edition, 2003.

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

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

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