MANAGING BELIEFS ABOUT MONETARY POLICY UNDER DISCRETION

TECHNICAL APPENDIX

Elmar Mertens*
Federal Reserve Board

October 30, 2014

Abstract

This appendix provides further details on the discretion problem (and its solution) in the generic class of linear-quadratic (LQ) models described in Section 2 of the main paper. For ease of reference, the description of the LQ framework is restated in Appendix A. Appendix B defines the private sector equilibrium and the notion of temporary equilibrium that is necessary for the discretion problem described in Appendix C. Appendix D considers the special case of imperfect information when lagged states are known.

Furthermore, this appendix provides more detailed results for the simple New Keynesian model described in the main paper (Appendix F and G) and illustrates the general framework with an extended model, featuring multiple sources of uncertainty, price and wage stickiness, and several endogenous state variables (Appendix H).

*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.

Contents

Part I
Generic LQ Case

A A Class of Linear Quadratic Models

This appendix restates the description of the class of LQ models given in Section 2 of the main paper. In this class of models, there are four types of variables:

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

These variables will be treated as vectors of dimensions N_x, N_u, N_y, and N_z respectively.

The backward looking variables can capture exogenous forcing variables but also endogenous states like capital, habits, or lagged variables, for example inflation in a model with price indexation. They evolve as

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

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

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

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

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

A sufficient condition to ensure superior information of the policymaker is that $N_z < N_w$ which prevents Z^t from spanning w^t. In principle, the public may not even be able to completely recover lagged values of the structural shocks (w_{t-k}) or lagged backward looking variables (X_{t-k}) from Z^t, which makes the information problem "dynamic" as current signals (Z_t) remain informative about past conditions.

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

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

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

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

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

$$E_t \sum_{k=0}^{\infty} \beta^k L_{t+k} \qquad L_t = \begin{bmatrix} X_t \\ Y_t \\ U_t \end{bmatrix}' Q \begin{bmatrix} X_t \\ Y_t \\ U_t \end{bmatrix} \qquad Q = \begin{bmatrix} Q_{xx} & Q_{xy} & N_x \\ Q'_{xy} & Q_{yy} & N_y \\ N'_x & N'_y & R \end{bmatrix} \tag{4}$$

where the per-period loss function L_t is quadratic in X_t, Y_t and U_t, Q is assumed to be a positive definite matrix, and the expectation operator is conditional on the history of w^t.[3]

Attention is limited here to Markov-perfect equilibria, which exclude reputational mechanisms via the kind of history-dependent strategies considered by Barro and Gordon (1983) or Chari and Kehoe (1990). In the spirit of "bygones are bygones", state variables in a Markov-perfect equilibrium must be relevant for current payoffs.[4] Under perfect information, only the backward-looking variables would qualify here as Markov-perfect state variables. In this case, a current decision-maker can influence a future decision-maker only if some of the backward-looking variables were endogenous, like capital or government debt. Under hidden information, the public's prior beliefs become part of the Markov states since they matter for the public's current decisions and payoffs. The prior beliefs, and not the posterior beliefs, enter the state vector of the policy problem, since the latter will be formed after observing current data which is influenced by current policy.[5]

In principle, the entire distribution of public beliefs needs to be tracked by the policy problem. The framework presented here affords a considerable simplification, which makes the problem very tractable: The model is cast in a Gaussian framework with constant variances, and the equilibrium is required to be linear and time-invariant. In this case, tracking entire distributions collapses to following changes in their means (and knowing their constant variances), which is easily provided by the Kalman filter.

B Private Sector Equilibrium

The policymaker is constraint by the beliefs and the behavior of the private sector. The private sector is atomistic and takes policies as given. Before turning to optimal policy, it is useful to consider notions of private sector equilibrium for a given policy. Private sector decisions are characterized

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

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

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

by the forward-looking constraint (3) above. As shown in this appendix, optimal private-sector beliefs are characterized by the Kalman Filter.

For the time being, the discussion adopts now the perspective of the private sector who takes a policy of the following form as given when forming beliefs and making choices:

$$U_t = F_1 X_t + F_2 X_{t|t-1} = F_1(X_t - X_{t|t-1}) + \hat{F} X_{t|t-1} \qquad (\hat{F} \equiv F_1 + F_2) \tag{5}$$

for some F_1, F_2. This policy anticipates a time-invariant, Markov-perfect equilibria, where policy depends only on current levels of backward-looking variables and prior beliefs about those. In equilibrium, the assumed values of F_1 and F_2 must, of course, be consistent with the optimal discretion policy (to be derived in Appendix C further below).

B.1 Kalman Filter

In a linear time-invariant equilibrium, and given the policy in (5), the solution to the public's optimal inference problem is characterized by the Kalman filter. Combining (5) with (1) and (2), state and measurement equations of the filter are

$$X_{t+1} = \underbrace{(A_{xx} + B_x F_1)}_{\equiv A} X_t + A_{xy} Y_{t|t} + B_x F_2 X_{t|t-1} + D w_{t+1} \tag{6}$$

$$Z_t = \underbrace{(C_x + C_u F_1)}_{\equiv C} X_t + C_u F_2 X_{t|t-1} \tag{7}$$

and beliefs evolve as

$$X_{t|t} = X_{t|t-1} + K \tilde{Z}_t \qquad \tilde{Z}_t \equiv Z_t - Z_{t|t-1} \tag{8}$$

The Kalman gain K is identical to the coefficients of a least squares projection of X_t on "innovations" $\tilde{Z}_t$

$$K \equiv \text{Cov}\left(X_t, \tilde{Z}_t\right) \ \text{Var}\left(\tilde{Z}_t\right)^{-1} = \Sigma C' \left(C\,\Sigma C'\right)^{-1} \tag{9}$$

where $\Sigma \equiv \text{Var}\left(X_t - X_{t|t-1}\right)$ solves the Riccati equation

$$\Sigma = A\Sigma A' + DD' - A\Sigma C' \left(C\Sigma C'\right)^{-1} C\Sigma A' \tag{10}$$

The Kalman filter depends only on the policy coefficients F_1, via which policy reacts to X_t, and is independent of the reaction coefficients F_2 associated with the predetermined state variable $X_{t|t-1}$. The presence of private sector controls $Y_{t|t}$ and predetermined variables $X_{t|t-1}$ in state and measurement equations (6) and (7) does not affect the Kalman gain.

The above assumes that the $N_z \times N_x$ matrix C has full row rank.[6] In principle (and also in practice) it could happen that C is collinear for some F_1. This corresponds to situations when there are multiple observables that are perfectly correlated such that $\text{Var}\,\tilde{Z}_t = C\Sigma C'$ is ill-conditioned. Economically, this means that a candidate policy F_1 tries to mimic other signals in Z_t.[7]

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

[7] I have never observed such mimicking strategies in equilibrium, but depending on initial conditions it can occur along the path of the policy improvement algorithm. In these cases, the Kalman filter can be implemented by pruning the redundancies in the set of observable variables via a singular value decomposition of C. To obtain numerically stable solution, this could be done for singular values of C below some threshold, say 10^{-8}.

For the policy problem, the Kalman filter describes not only the transition dynamics of public beliefs over time, but also their dependence on policy choices U_t via the measurement equation (2). To describe the "open loop" dynamics of beliefs — that is the dynamics of beliefs for an arbitrary sequence of policy choices U_t — it is necessary to specify the public's prior beliefs about current policy $U_{t|t-1}$. As in the derivation of the Kalman Gain, it is assumed that these beliefs are derived from given policy coefficients F_1 and F_2, which will have to coincide with optimal policy choices in equilibrium. It thus follows that $U_{t|t-1} = \hat{F} X_{t|t-1}$ and the Kalman filter's updating equation can be written as follows.

$$X_{t|t} = K\, C_x\, X_t + (I - K\, \hat{C}) X_{t|t-1} + K\, C_u\, U_t \qquad \hat{C} \equiv (C_x + C_u(\hat{F})) \tag{11}$$

B.2 Temporary Private Sector Equilibrium

To constrain the policymaker's discretion problem, a weak form of private sector equilibrium is sufficient. It is a temporary equilibrium in the spirit of Grandmont (1977):

Definition (Temporary Private Sector Equilibrium)**.** *At a given point in time, the private sector has given beliefs about current policy according to (5). They are embodied in a Kalman gain K used to update beliefs about X_t as in (8). Furthermore, people hold beliefs about future policies (which may be different from the public's beliefs about current policy). The beliefs about future policies are embodied in a matrix of coefficients G, mapping public beliefs about future states into public beliefs about future forward-looking variables:*

$$Y_{t+1|t} = G X_{t+1|t} \tag{12}$$

The temporary equilibrium then reduces to optimal choices which satisfy the forward looking constraint (3) given the beliefs in (12).

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

$$\begin{aligned} Y_{t|t} &= G_x X_{t|t} + G_u U_{t|t} && (13)\\ &= \Gamma_1 X_t + \Gamma_2 X_{t|t-1} + \Gamma_u U_t && (14) \end{aligned}$$

where (14) follows from combining (13) with the Kalman filter's open loop equation (11) and where

$$\begin{aligned} G_x &= (A^1_{yy}\, G\, A_{xy} - A_{yy})^{-1} (A_{yx} - A^1_{yy}\, G\, A_{xx})\\ G_u &= (A^1_{yy}\, G\, A_{xy} - A_{yy})^{-1} (B_y - A^1_{yy}\, G\, B_x)\\ \Gamma_1 &= (G_x + G_u\, F_1) K\, C_x\\ \Gamma_2 &= (G_x + G_u F_1)(I - K\hat{C}) + G_u F_2\\ \Gamma_u &= (G_x + G_u F_1) K\, C_u \end{aligned}$$

Given the policy in (5), (12) also implies the following decision rule for the private sector.

$$\begin{aligned} Y_t &= (G_x + G_u\hat{F})X_{t|t-1} + (G_x + G_u F_1)\, K\, C(X_t - X_{t|t-1}) && (15) \\ &= (G_x + G_u F_1)\, K\, C\; X_t + [(G_x + G_u F_1)(I - K\, C) + (G_x + G_u F_2)]\; X_{t|t-1} && (16) \\ &= G_1 X_t + G_2 X_{t|t-1} && (17) \end{aligned}$$

The construction of a temporary equilibrium is not unique to the imperfect information setup considered in this paper. Similar computations are for example employed also by Söderlind (1999) in his derivation of optimal Markov perfect policies under perfect information.

B.3 Existence and Uniqueness of a Private Sector Equilibrium

As a slight digression, it might be useful to consider a stronger notion of private sector equilibrium than the temporary equilibrium defined above. Readers interested in solving the policymaker's discretion problem may however skip this part; only a temporary equilibrium will be assumed in the derivation of the policymaker's decision problem. The stronger notion of private sector equilibrium considered here corresponds to a time-variant rational expectations equilibrium for the private sector, given a time-invariant policy reaction function as in (5). If such an equilibrium exists and can be uniquely determined, it can be used in a policy improvement algorithm for solving the optimal discretion problem as described in Appendix C.

Definition (Private Sector Equilibrium)**.** *Given the policy in (5), the private sector equilibrium is a sequence of observations* $\{Z_t\}$*, perceived states* $\{X_{t|t}\}$*, perceived policies* $\{U_{t|t}\}$ *and private sector choices* $\{Y_t\}$ *such that*

- *Expectations and beliefs are rational. In this linear framework, they are formed using the Kalman filter with measurements* Z_t*.*
- *Choices are optimal, that is they satisfy the forward looking constraint (3).*

Optimal choices of the private sector solve the forward-looking constraint (3) given the policy (5) and private sector beliefs $X_{t|t}$ derived from the Kalman filter (8). The result is a system of expectational difference equations driven by disturbances ($\tilde{Z}_t$) that are serially uncorrelated from the perspective of the public's information set.

$$\begin{aligned} X_{t+1|t} &= (A_{xx} + B_x\hat{F})X_{t|t-1} + A_{xy}Y_{t|t} + (A_{xx} + B_x F_1)\, K\, \tilde{Z}_t && (18) \\ A^1_{yy}Y_{t+1|t} &= (A_{yx} + B_y\hat{F})X_{t|t-1} + A_{yy}Y_{t|t} + (A_{yx} + B_y F_1)\, K\, \tilde{Z}_t && (19) \end{aligned}$$

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

$$\bar{\boldsymbol{A}} = \begin{bmatrix} I & 0 \\ 0 & A^1_{yy} \end{bmatrix} \quad \text{and} \quad \bar{\boldsymbol{B}} = \begin{bmatrix} (A_{xx} + B_x\hat{F}) & A_{xy} \\ (A_{yx} + B_y\hat{F}) & A_{yy} \end{bmatrix} \qquad (20)$$

collect the coefficients on the endogenous variables in the system described by (18) and (19). This is the kind of linear systems studied by King and Watson (1998) and Klein (2000), and their "counting rules" for stable and unstable roots can be applied to derive conditions for existence and uniqueness of the private sector equilibrium.

Proposition 1 (Existence and Uniqueness of Private Sector Equilibrium)**.** *Existence and uniqueness of a private sector equilibrium depend on the roots z of $|\bar{\boldsymbol{A}}z - \bar{\boldsymbol{B}}| = 0$ for matrices $\bar{\boldsymbol{A}}$ and $\bar{\boldsymbol{B}}$ defined above. A unique equilibrium exists only if there are N_x roots inside the unit circle and N_y outside. The matrices $\bar{\boldsymbol{A}}$ and $\bar{\boldsymbol{B}}$, and thus also the condition for existence and uniqueness, depend on the policy rule (5) but not on the Kalman gain K. 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

$$\begin{aligned} Y_{t|t} &= \bar{G}X_{t|t-1} + H_y\tilde{Z}_t \\ X_{t+1|t} &= \bar{A}X_{t|t-1} + H_x\tilde{Z}_t \end{aligned}$$

where $\bar{G}$ and $\bar{A}$ depend only on $\bar{\boldsymbol{A}}$ and $\bar{\boldsymbol{B}}$ but not on K (for given policy coefficients F_1 and F_2.) □

In the simple New Keynesian model discussed in the main paper, the condition is trivially met but in general this needs not be the case.[8] A pertinent example is the nominal indeterminacy of Sargent and Wallace (1975), which holds for any exogenous policy like (5) when the interest rate is the control variable and applies also to the New Keynesian model as discussed for example by Galì (2008).

C Discretion Policy and Equilibrium

Discretionary policy is time-consistent. At each point in time the policymaker can reoptimize while taking his future optimizations as given. A recursive representation of the optimization problem as a dynamic program is thus an inherent feature of the discretionary policy problem. This appendix describes the discretionary policy problem, its representation as a linear regulator problem and a solution algorithm. The state variables of the markov-perfect discretion problem are the backward looking variables (X_t) and the public's prior beliefs ($X_{t|t-1}$), there is no further history dependence. Furthermore, the policymaker must account for the rational expectations and optimal choices of the private sector as captured by the Kalman filter and temporary private sector equilibrium described in Appendix B. This is summarized in the following definition.

Definition (Discretionary Policy)**.** *At each point in time, for given private beliefs embodied in F_1, F_2, K and and G, the policymaker chooses U_t to minimize*

$$\begin{aligned} V_t(S_t) &= \min_{U_t, Y_t, X_{t+1}, X_{t|t}} \{L_t + \beta E_t V_{t+1}(S_{t+1})\} \\ \textit{s.t.} \quad X_{t+1} &= A_{xx}X_t + A_{xy}Y_t + B_xU_t + Dw_{t+1} \\ X_{t|t} &= X_{t|t-1} + K(Z_t - Z_{t|t-1}) \\ Y_{t|t} &= \Gamma_1 X_t + \hat{\Gamma}_2 X_{t|t-1} + \Gamma_u U_t \end{aligned}$$

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

where beliefs $X_{t|t}$, $Z_t tm1$ etc. are characterized by the Kalman filter (8), the period-loss function L_t is defined in (4) and Γ_1, Γ_2, Γ_u are derived from the temporary equilibrium in (14) above. The constraints correspond to the transition equation for X_t (1), the private sector's Kalman Filter (8) and temporary equilibrium. The continuation value of this dynamic optimization problem, $V_{t+1}(S_{t+1})$, is a given function of future, Markov perfect state variables

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

$V_t(\cdot)$ and $V_{t+1}(\cdot)$ can be different functions of the state vector, but in a time-invariant equilibrium, both must, of course, collapse to a time-invariant value function $V(\cdot)$.

Equilibrium requires that the policymaker's policy choices, public beliefs and and the private sector's temporary equilibrium are mutually consistent.

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.

This equilibrium concept is similar to the self-confirming equilibria of Fudenberg and Levine (1993) and Sargent (1999) in that both are a fixed point of mutual beliefs and actions in multi-player games. However, in a self-confirming equilibrium, players can hold erroneous beliefs about the structure of the economy, as long as they are justified by observable outcomes.[9] However, in the setup considered here, the public knows and understands the true structure of the economy.

In light of the linear-quadratic structure of the model, the value function can be taken to be linear quadratic as well (Bertsekas, 2005) and for given coefficients of the private sector's temporary equilibrium and Kalman filter discretionary policy solves a linear regulator problem, which will be described shortly below. The resulting optimal policy has then exactly the form anticipated in (5).

$$U_t = F_1 X_t + F_2 X_{t|t-1} = \boldsymbol{F} S_t \tag{21}$$

Formally, the discretionary equilibrium then requires that $\boldsymbol{F}$ supports the Kalman Gain and coefficients Γ_1, Γ_2 and Γ_u of the private sector's temporary equilibrium (14).

[9] A similar fixed point of beliefs and outcomes is used in the limited-information rational expectations equilibria of Marcet and Sargent 1989b; 1989a and Sargent (1991).

C.1 Regulator for Discretion Problem

A linear regulator is a dynamic program choosing policy controls U_t to minimize a present value of expected discounted losses L_t subject to the evolution of states S_t. Critically, a linear regulator is a backward-looking dynamic program, whereas the LQ class of models laid out in Appendix A is forward looking. However, the discretionary policymaker needs to take the private sector's temporary equilibrium as given and the forward-looking variables Y_t can be substituted out using (14). All that is required to set up the linear regulator problem is thus to rewrite the loss function and state transitions in terms of U_t and S_t. Public beliefs $X_{t|t-1}$ are part of the state vector and — for a given Kalman gain K — their transition equation is given by the Kalman filter. The coefficients involved in the Kalman filter and the temporary equilibrium — in particular Γ_1, Γ_2, Γ_u, K, C and $\hat{C}$ — are for now taken as given, typically derived from some initial guess about the policymaker's reaction coefficients F_1 and F_2. The linear regulator will then yield the policymaker's optimal choice of these reaction coefficients which can be used to update the Kalman gain and the coefficients of the private sector's temporary equilibrium until convergence.

Loss Function

The loss function (4) can be rewritten in terms of the regulator's states and control using

$$\begin{bmatrix} X_t \\ Y_{t|t} \\ U_t \end{bmatrix} = \underbrace{\begin{bmatrix} I & 0 & 0 \\ \Gamma_1 & \Gamma_2 & \Gamma_u \\ 0 & 0 & I \end{bmatrix}}_{\boldsymbol{H}} \begin{bmatrix} X_t \\ X_{t|t-1} \\ U_t \end{bmatrix}$$

such that

$$L_t = \begin{bmatrix} X_t \\ Y_t \\ U_t \end{bmatrix}' Q \begin{bmatrix} X_t \\ Y_t \\ U_t \end{bmatrix} = \begin{bmatrix} S_t \\ U_t \end{bmatrix}' \boldsymbol{H}' Q \boldsymbol{H} \begin{bmatrix} S_t \\ U_t \end{bmatrix} = S_t' \boldsymbol{Q} S_t + 2 S_t' \boldsymbol{N} U_t + U_t' \boldsymbol{R} U_t \tag{22}$$

where $\boldsymbol{Q}$, $\boldsymbol{N}$ and $\boldsymbol{R}$ conformably partition the quadratic form $\boldsymbol{H}' Q \boldsymbol{H}$.

State Transition

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

$$\begin{aligned} X_{t+1} &= (A_{xx} + A_{xy}\Gamma_1) X_t + A_{xy}\Gamma_2 X_{t|t-1} + (A_{xy}\Gamma_u + B_x) U_t + D w_{t+1} \\ X_{t+1|t} &= \boldsymbol{A_{xx}} K C_x X_t + \left(\boldsymbol{A_{xx}} (I - K\hat{C}) + (A_{xy} G_u + B_x) F_2 \right) X_{t|t-1} + \boldsymbol{A_{xx}} K C_u U_t \\ \text{where} \quad \boldsymbol{A_{xx}} &= A_{xx} + A_{xy}(G_x + G_u F_1) + B_x F_1 \end{aligned}$$

The state transition can thus compactly be written as

$$S_{t+1} = \boldsymbol{A} S_t + \boldsymbol{B} U_t + \boldsymbol{D} w_t \tag{23}$$

where $\boldsymbol{A}$, $\boldsymbol{B}$ and $\boldsymbol{D}$ given by

$$\boldsymbol{A} = \begin{bmatrix} (A_{xx} + A_{xy}\Gamma_1) & A_{xy}\Gamma_2 \\ \boldsymbol{A_{xx}}KC_x & \left(\boldsymbol{A_{xx}}(I - K\hat{C}) + (A_{xy}G_u + B_y)F_2\right) \end{bmatrix},$$

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

Linear Regulator

What remains to be specified is the the value function, which is quadratic in the states in a linear-quadratic regulator problem (Bertsekas, 2005):

$$V_{t+1}(S_{t+1}) = S'_{t+1}\boldsymbol{V}S_{t+1} + v \tag{24}$$

where $\boldsymbol{V}$ is taken as given.[10] The regulator problem thus has the following form:

$$S'_t\boldsymbol{V}^*S_t + v^* = \min_{U_t, S_{t+1}} \left\{ S'_t\boldsymbol{Q}S_t + 2S'_t\boldsymbol{N}U_t + U'_t\boldsymbol{R}U_t + \beta E_t S'_{t+1}\boldsymbol{V}S_{t+1} + v \right\} \tag{25}$$

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

and the first-order conditions yield

$$\begin{aligned} U_t &= -(\boldsymbol{R} + \beta\boldsymbol{B}'\boldsymbol{V}\boldsymbol{B})^{-1}(\boldsymbol{N} + \beta\boldsymbol{B}'\boldsymbol{V}\boldsymbol{A})\, S_t \\ &\equiv \boldsymbol{F}^* S_t. \end{aligned} \tag{27}$$

The optimal policy is linear in S_t as has been anticipated in (5). The policy appears certainty equivalent since it is independent of the shock loadings $\boldsymbol{D}$.[11] However, the setup of the regulator itself is not certainty equivalent since it depends on the private sector's Kalman filter and the optimal policy coefficients $\boldsymbol{F}^*$ are *not* certainty equivalent.

Value Function consistent with a given $\boldsymbol{F}$

The policy improvement algorithm described further below uses a continuation value consistent with carrying out the policy $\boldsymbol{F}$ forever. This continuation value ca be computed from the closed loop representation of the regulator — for given $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{D}$, $\boldsymbol{Q}$, $\boldsymbol{N}$, $\boldsymbol{R}$ — by plugging the policy $\boldsymbol{F}$ into (25) and (26). $\boldsymbol{V}$ then solves the Lyapunov equation

$$\boldsymbol{V} = \{\boldsymbol{Q} + \boldsymbol{N}\boldsymbol{F} + \boldsymbol{F}'\boldsymbol{R}\boldsymbol{F}\} + \beta(\boldsymbol{A} + \boldsymbol{B}\boldsymbol{F})'\boldsymbol{V}(\boldsymbol{A} + \boldsymbol{B}\boldsymbol{F})$$

The equation has a unique solution if the matrix in curly braces is positive definite and if the closed loop transition matrix $(\boldsymbol{A} + \boldsymbol{B}\boldsymbol{F})$ has all eigenvalues inside the unit circle. The former is assured by the form of the original loss function and the latter holds if a stationary equilibrium exists.[12]

[10] The policy improvement algorithm described below will derive $\boldsymbol{V}$ to be consistent with continuing a given candidate policy forever, using the same candidate policy underlying the private sector's Kalman filter and temporary equilibrium.

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

[12] 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).

The constant v in (24) does not affect the optimal policy coefficients. Still, supposing that the policy $\boldsymbol{F}$ is applied forever, the scalar v can be computed from

$$v = \frac{\beta}{1-\beta} \operatorname{tr}(\boldsymbol{V}\boldsymbol{D}\boldsymbol{D}').$$

Similarly, unconditionally expected losses can be computed from the unconditional variance covariance matrix of the states:

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

where the last equation is another Lyapunov equation.

C.2 Policy Improvement Algorithm

The equilibrium is a fixed point of public beliefs and policy actions and maps $(\boldsymbol{F}, G, \boldsymbol{V})$ into itself. An intuitive and efficient way to compute this fixed point is a policy improvement algorithm. It is efficient, since policy improvement methods converge faster than value function iterations (Whittle, 1996; Bertsekas, 2005).[13] It is intuitive, since the algorithm uses the regulator (25) to seek for a one-period deviation from a candidate equilibrium. Non-existence of such a deviation is the defining property of equilibrium.

Formally, the algorithm starts with a candidate policy $\boldsymbol{F^0}$ and beliefs G^0 and computes the Kalman gain K and continuation value $\boldsymbol{V}$ associated with continuing this policy forever. If the conditions for a private sector equilibrium are met (Proposition 1), one can even compute the G^0 consistent with $\boldsymbol{F^0}$. The solution (27) to the above regulator problem then yields the optimal one-period deviation. As long as $\boldsymbol{F^0} \neq \boldsymbol{F}$ and $G^0 \neq G$ there is no equilibrium. In this case, a new iteration starts using (F, G) as new candidate policies.

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

The above equilibrium is an intricate fixed point between optimal one-period policies ($\boldsymbol{F}$), and public beliefs ($\boldsymbol{F^0}$, G^0). Formally, it is a fixed point between two Riccati equations, one from the policymaker's regulator problem, the other from 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. However, in my practical experience, the algorithm typically converges, and if so always to the same equilibrium from arbitrary starting values for $(\boldsymbol{F^0}, G^0)$.

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

D Imperfect Information when Lagged States are Known

This appendix considers a special case of imperfect information, discussed in Section 2.2 of the main paper, where it is assumed that lagged variables are always known to the public. Specifically, it is assumed that lagged expectations (conditioned on the full information set at $t-1$ spanned by w^{t-1}) of the backward-looking variables, $E_{t-1}X_t$, are known to the public, since these expectations are sufficient statistics for the information content in w^{t-1} that is relevant for decision making at t. As discussed in the main paper, imperfect information is then purely static in nature, and the lagged public beliefs $X_{t|t-1}$ are replaced by lagged expectations $E_{t-1}X_t$ (in time t decision problems). In principle, this can be handled within the framework described in the preceding appendices by enlarging the set of backward-looking variables to also track $E_{t-1}X_t$ and by including these extra variables in the public's measurement vector. However, this approach would also track lagged beliefs $X_{t|t-1}$ and compute policy reactions to these beliefs even though they are not relevant in Markov-perfect policy problem. The resulting dynamic system would however be inefficiently large, and unneccessarily involve the solution of a Riccati equation.

This appendix briefly describes how the general framework described above can easily be adapted to the case when lagged states are known. Only two changes are necessary: First, the Kalman filter is replaced by a static signal extraction problem. Second, the transition equation for lagged beliefs is replaced the definition of lagged expectations $E_{t-1}X_t$. Importantly, as the setup of the temporary equilibrium is not germane to the imperfect information structure of the policy problem — see for example Söderlind (1999) — its structure remains unaffected in this special case. Despite some differences, there is thus quite some similarity between the general framework and the special case described in this appendix; accordingly I will deliberately recycle some notation.

As before, beliefs and temporary equilibrium of the private sector are constructed while taking policy as a given linear function of the states. In this particular equilibrium, the markov-perfect states are the backward-looking variables and their lagged expectations:

$$U_t = \begin{bmatrix} F_1 & F_2 \end{bmatrix} \begin{bmatrix} X_t \\ E_{t-1}X_t \end{bmatrix} = \boldsymbol{F} S_t \tag{28}$$

Signal Extraction Problem

Given (28), the public forms optimal beliefs about X_t with a static signal extraction problem.

$$\begin{aligned} X_{t|t} &\equiv E(X_t|Z^t, E_{t-1}X_t) && (29)\\ &= E_{t-1}X_t + K(Z_t - E_{t-1}Z_t) && (30)\\ &= (I - K\hat{C})E_{t-1}X_t + K\,C_x X_t + K\,C_U U_t && (31)\\ \text{with} \quad K &= (DD')C'(CDD'C')^{-1} && (32) \end{aligned}$$

and where Z_t is still defined by (2), $\hat{C} = C_x + C_U\hat{F}$ and $\hat{F} = F_1 + F_2$ as before.

Temporary Equilibrium

The construction of the temporary equilibrium is the same as in Appendix B, except for replacing $X_{t|t-1}$ by $E_{t-1}X_t$ and forward-looking variables are restricted to be

$$Y_{t|t} = \Gamma_1 X_t + \Gamma_2 E_{t-1}X_t + \Gamma_u U_t \tag{33}$$

State Transition

The state transition equation for S_t follows from the definitions of X_t and $E_{t-1}X_t$ and the temporary equilibrium's restriction (33) on the forward-looking variables.

$$S_{t+1} = \begin{bmatrix} A_{xx} + A_{xy}\Gamma_1 & A_{xy}\Gamma_2 \\ A_{xx} + A_{xy}\Gamma_1 & A_{xy}\Gamma_2 \end{bmatrix} S_t + \begin{bmatrix} A_{xy}\Gamma_u \\ B_x A_{xy}\Gamma_u \end{bmatrix} U_t + \begin{bmatrix} D \\ 0 \end{bmatrix} w_{t+1} \tag{34}$$

$$= \boldsymbol{A} S_t + \boldsymbol{B} U_t + \boldsymbol{D} w_{t+1} \tag{35}$$

When these modifications are taken into account, the regulator problem and improvement algorithm described in Appendix C can be applied to solving the discretion equilibrium under imperfect information when lagged states are known.

E Mapping the simple model into the LQ framework

The simple New Keynesian model described in the main paper can easily be represented within the general framework described above. The output gap equals the policy control, $U_t = x_t$. The backward looking variables are the two exogenous components of the output gap target, and their transition equation is

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

with $A_{xy} = 0$ and $B_x = 0$. Inflation is the only forward looking variable of the private sector, $Y_t = \pi_t$, and the Phillips curve corresponds to the associated forward looking constraint with $A^1_{yy} = \beta$, $A_{yy} = 1$, $A_{yx} = \begin{bmatrix} 0 & 0 \end{bmatrix}$, and $B_y = -\kappa$. In the case of imperfect information (when lagged states are not known), the signal is x_t and thus $C_x = \begin{bmatrix} 0 & 0 \end{bmatrix}$, $C_u = 1$.[14]

[14]The special case of imperfect information with known lagged states can be replicated by augmenting the system of backward-looking variables to track $\begin{bmatrix} \gamma_t & \varepsilon_t & E_{t-1}\gamma_t \end{bmatrix}'$ and including $E_{t-1}\gamma_t$ in the measurement vector.

Part II
Simple New Keynesian Model

This part of the appendix presents further details about the simple New Keynesian Model discussed in Sections 3 and 4 of the main paper.

F Equilibrium Determination When Lagged States are Known

This appendix provides a graphical illustration of the roots to the second-order polynomial that determines the optimal policy coefficients in the static imperfect-information case of the simple New Keynesian model of Section 3.3 in the main paper. In this case, optimal policy has the form

$$x_t = \tilde{f}\eta_t + \tilde{f}\varepsilon_t + \hat{f}\ E_{t-1}\gamma_t \tag{36}$$

and as stated in Proposition 1 of the main paper, the optimal policy coefficients are given by

$$\tilde{f} = \frac{\alpha_x}{\alpha_x + \alpha_\pi \tilde{\kappa}^2} \qquad \hat{f} = \frac{(1-\beta\rho(1-\hat{R}^2))\alpha_x}{(1-\beta\rho(1-\hat{R}^2))\alpha_x + \alpha_\pi\tilde{\kappa}^2} \tag{37}$$

$$\text{with} \quad \tilde{\kappa} = \kappa \cdot \left(1 + \frac{\beta\rho}{1-\beta\rho}\hat{R}^2\right) \qquad \hat{R}^2 = K \cdot \hat{f} = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_\varepsilon^2} \cdot \frac{\hat{f}}{\tilde{f}} \tag{38}$$

Inspection of (37) shows that a given ratio $\phi \equiv \hat{f}/\tilde{f}$, uniquely determines the individual policy coefficients $\tilde{f}$ and $\hat{f}$. The ratio ϕ itself solves a second-order polynomial:

$$\phi = \frac{\hat{f}}{\tilde{f}} = \frac{(1-\beta\rho)^2\alpha_x + (1-\beta\rho(1-R^2\phi))^2\alpha_\pi\kappa^2)}{(1-\beta\rho)^2\alpha_x + (1-\beta\rho(1-R^2\phi))\alpha_\pi\kappa^2)} \qquad \Leftrightarrow \qquad q\phi^2 = c - l\phi \tag{39}$$

With

$$\begin{aligned} q &= \alpha_\pi\kappa^2\beta\rho R^2(1-\beta\rho R^2) \\ l &= (1-\beta\rho)^2\alpha_x + \alpha_\pi\kappa^2(1-\beta\rho)(1-2\beta\rho R^2) \\ c &= (1-\beta\rho)^2\alpha_x + \alpha_\pi\kappa^2(1-\beta\rho) \end{aligned}$$

The roots of (39) are given by the intersection of an upwardly open parabola and a downward sloping straight line. The parabola has its minimum at the origin of the coordinate system whereas the line has a positive intercept; there are thus two real roots, one positive and one negative. Since $q > c - l$ the positive root must lie between zero and one. (There is no such constraint on the negative root.) A graphical illustration for two different parameter choices is provided in Figure 1. In both cases, the positive root lies between zero and one and the second root is negative. However, depending on the choice of α_x the negative root can be inside or outside the unit circle.

Figure 1: Roots of second-order polynomial in ϕ

(a) Baseline calibration with $\alpha_x = \alpha_\pi = 0.5$ (b) Alternative case with $\alpha_x = 0.1$ and $\alpha_\pi = 0.9$

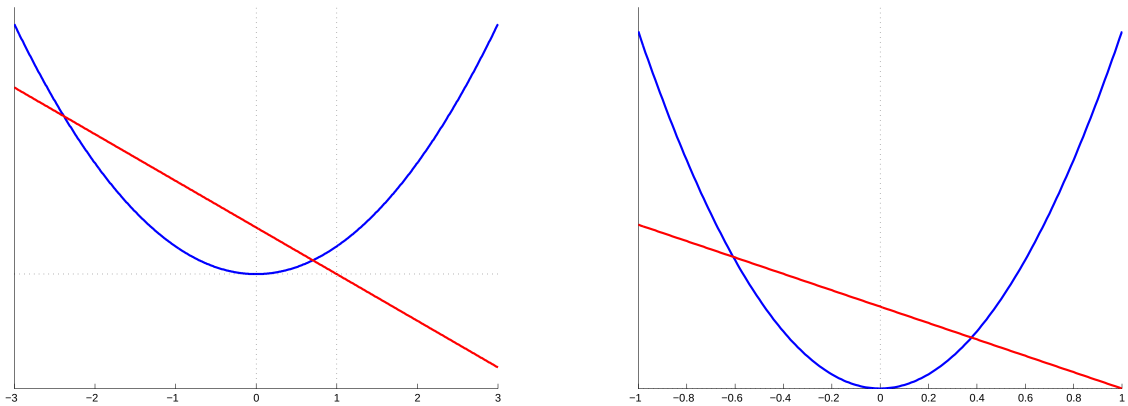

Note: Roots of the polynomial (39) for two different parameter choices for $\alpha_x = 1 - \alpha_\pi$. Other parameters are held at their baseline values listed in Table 1.

G Alternative Calibrations of the New Keynesian Example

G.1 Response Coefficients

This appendix supplements the comparison of full- and hidden-information equilibria in Section 3 of the main paper. Propositions 2 and 4 of the paper compare the response coefficents for optimal output gaps and inflation under full information as well as the static and dynamic cases of hidden information. Figures 2 and 3 provide surface plots of these coefficients for different model calibrations. Each figures considers variations in $\alpha_x = 1 - \alpha_\pi$. the weight placed in the policymaker's loss function on minimizing the shortfall of the output gap from its exogenous target value, and $R^2 = \sigma^2_\eta / (\sigma^2_\eta + \sigma^2_\varepsilon)$, the share of the target shock variance accounted for by shocks to the persistent target component. All other model parameters are held at their baseline values, listed in Table 1.[15]

The left-hand panels of Figures 2 and 3 compare response coefficients for the static hidden-information case against those obtained under full information; the right-hand panels compare the dynamic case of imperfect information against full information. For added contrast, surfaces are either colored black (full information) or white (hidden information). In the three information cases considered, output gap and inflation are characterized by:

$$\text{Full information:} \quad x_t = \bar{f}_\gamma \gamma_t + \bar{f}_\varepsilon \varepsilon_t \qquad \pi_t = \bar{g}_\gamma \gamma_t + \bar{g}_\varepsilon \varepsilon_t \tag{40}$$

$$\text{Static hidden information:} \quad x_t = \tilde{f}(\eta_t + \varepsilon_t) + \hat{f} E_{t-1}\gamma_t \qquad \pi_t = \tilde{g}(\eta_t + \varepsilon_t) + \hat{g} E_{t-1}\gamma_t \tag{41}$$

$$\text{Dynamic hidden information:} \quad x_t = \tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}} \tilde{\gamma}_t + \tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}} \tilde{\varepsilon}_t + \hat{\boldsymbol{f}} \gamma_{t|t-1} \qquad \pi_t = \tilde{\boldsymbol{g}}_{\boldsymbol{\gamma}} \tilde{\gamma}_t + \tilde{\boldsymbol{g}}_{\boldsymbol{\varepsilon}} \tilde{\varepsilon}_t + \hat{\boldsymbol{g}} \gamma_{t|t-1} \tag{42}$$

[15] Qualitatively similar pictures, not shown here, can also be obtained when using different values for the remaining model parameters; notably β, ρ and κ.

As predicted by Propositions 2 and 4, output gap responses to ε_t and $E_{t-1}\gamma_t$ (or $\gamma_{t|t-1}$) are always smaller than their full-information counterparts; $\tilde{f} < \bar{f}_\varepsilon$ and $\tilde{\boldsymbol{f}}_{\boldsymbol{\varepsilon}} < \bar{f}_\varepsilon$ as well as $\hat{f} < \bar{f}_\gamma$ and $\hat{\boldsymbol{f}} < \bar{f}_\varepsilon$. In consequence, the black surfaces (full information) lie everywhere above the white surfaces (hidden information) in Panels (c) through (f) of Figure 2. As stated in the propositions, $\bar{f}_\gamma$ can however be larger or smaller than $\tilde{f}$ and $\tilde{\boldsymbol{f}}_{\boldsymbol{\gamma}}$, respectively, which can be verified in Panels (a) and (b). In particular, notice that hidden information responses to shocks η_t (or innovations $\tilde{\gamma}_t$ in the dynamic case) are larger than in the full-information case, when persistent shocks are rare (low R^2) or unimportant for the policymaker (low α_x). In those cases, the public can (correctly) anticipate that persistent variations in the output gap are unlikely to occur or small (or both), see Panels (a) and (b) of Figure 3. In consequence, the sensitivity of inflation in either case will be relatively small, such that the policymaker "can get away" with pursuing the persistent output target more aggressively than in the full-information case.

Similarly, Panels (a), (b), (e) and (f) of Figure 3 confirm the predictions of Propositions 2 and 4 for inflation responses, in particular that $\tilde{g} < \bar{g}_\gamma$ and $\tilde{\boldsymbol{g}}_{\boldsymbol{\gamma}} < \bar{g}_\gamma$ as well as $\hat{g} < \bar{g}_\gamma$ and $\hat{\boldsymbol{g}} < \bar{g}_\gamma$, whereas Panels (c) and (d) show that $\bar{g}_\varepsilon$ can be larger or smaller than $\tilde{g}$ and $\tilde{\boldsymbol{g}}_{\boldsymbol{\varepsilon}}$, respectively, for different parameter values. Notice that $\bar{g}_\varepsilon$ lies above $\tilde{g}$ and $\tilde{\boldsymbol{g}}_{\boldsymbol{\varepsilon}}$, respectively, only for a small part of the surface areas shown in both panels, representing mostly combinations of particularly low values of α_x and high R^2. Elsewhere, the hidden-information responses of inflation to ε_t are larger than $\bar{g}_\varepsilon$ since the potential confusion with persistent changes in the output target causes a stronger inflation responds, which tends to dominate the outcomes under optimal policy. However, for small α_x and high R^2, optimal policy scales its output gap responses to both target shocks so far below their full-information counterparts that, on net, the hidden-information response of inflation to ε_t is below $\bar{g}_\varepsilon$. (Of course, for small values of R^2 the likelihood of a transitory shock to the output gap target is fairly small.)

Figure 2: Response Coefficients for the Output Gap

Responses to η_t and $\tilde{\gamma}_t$, respectively

(a) Static: $\tilde{f}$ vs. $\bar{f}_\gamma$ (b) Dynamic: $\tilde{\boldsymbol{f}}_\gamma$ vs. $\bar{f}_\gamma$

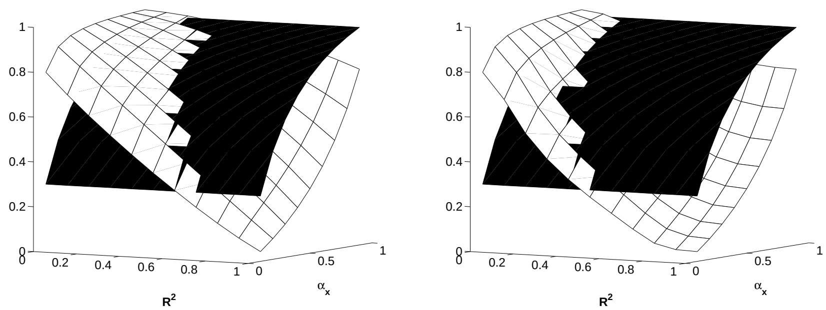

Response to ε_t

(c) Static: $\tilde{f}$ vs. $\bar{f}_\varepsilon$ (d) Dynamic: $\tilde{\boldsymbol{f}}_\varepsilon$ vs. $\bar{f}_\varepsilon$

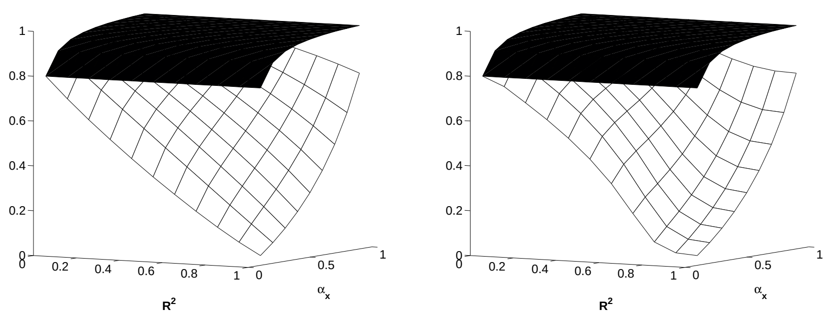

Response to $E_{t-1}\gamma_t$ and $\gamma_{t|t-1}$, respectively

(e) Static: $\hat{f}$ vs. $\bar{f}_\gamma$ (f) Dynamic: $\hat{\boldsymbol{f}}$ vs. $\bar{f}_\gamma$

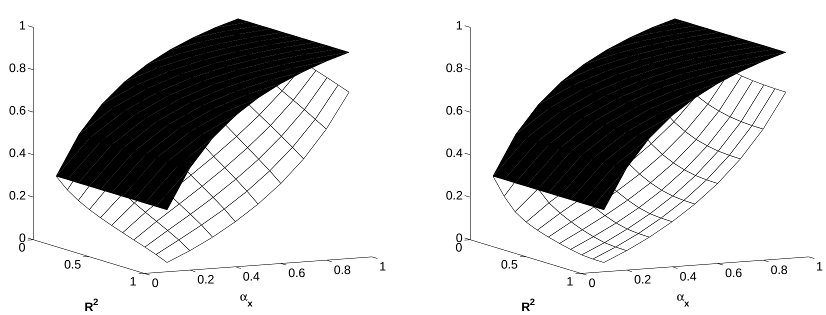

Note: Comparison of response coefficients of the output gap as given by equations (40), (41), and (42) for the cases of full information, as well as static and dynamic hidden information, respectively. In each panel, different values for the full-information coefficients are depicted as a black surface, whereas the white surfaces depict values of the corresponding hidden-information coefficients. Each panel considers variations in $\alpha_x = 1 - \alpha_\pi$ and R^2. All other model parameters are held at their baseline values, listed in Table 1.

Figure 3: Response Coefficients for Inflation

Responses to η_t and $\tilde{\gamma}_t$, respectively

(a) Static: $\tilde{g}$ vs. $\bar{g}_\gamma$

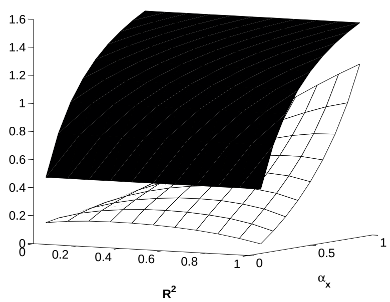

(b) Dynamic: $\tilde{\boldsymbol{g}}_{\boldsymbol{\gamma}}$ vs. $\bar{g}_\gamma$

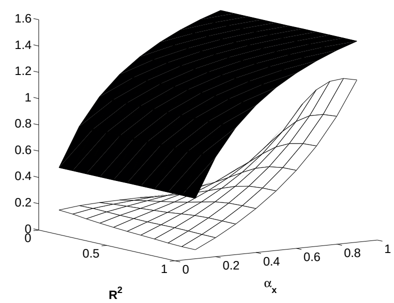

Response to ε_t

(c) Static: $\tilde{g}$ vs. $\bar{g}_\varepsilon$

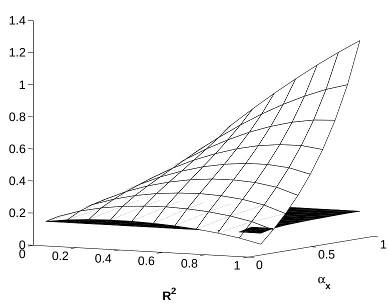

(d) Dynamic: $\tilde{\boldsymbol{g}}_{\boldsymbol{\varepsilon}}$ vs. $\bar{g}_\varepsilon$

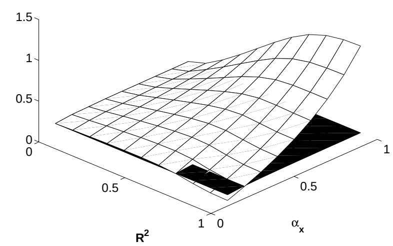

Response to $E_{t-1}\gamma_t$ and $\gamma_{t|t-1}$, respectively

(e) Static: $\hat{g}$ vs. $\bar{g}_\gamma$

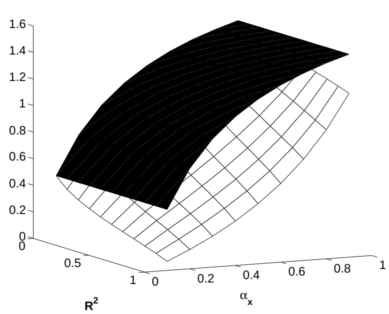

(f) Dynamic: $\hat{\boldsymbol{g}}$ vs. $\bar{g}_\gamma$

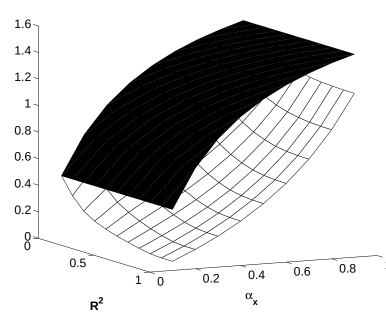

Note: Note: Comparison of response coefficients for inflation as given by equations (40), (41), and (42) for the cases of full information, as well as static and dynamic hidden information, respectively. In each panel, different values for the full-information coefficients are depicted as a black surface, whereas the white surfaces depict values of the corresponding hidden-information coefficients. Each panel considers variations in $\alpha_x = 1 - \alpha_\pi$ and R^2. All other model parameters are held at their baseline values, listed in Table 1.

Figure 4: Gains from Hidden Info

$$\boldsymbol{E(L^{FI}) - E(L^{HI})}$$

(a) Static hidden info (b) Dynamic hidden info

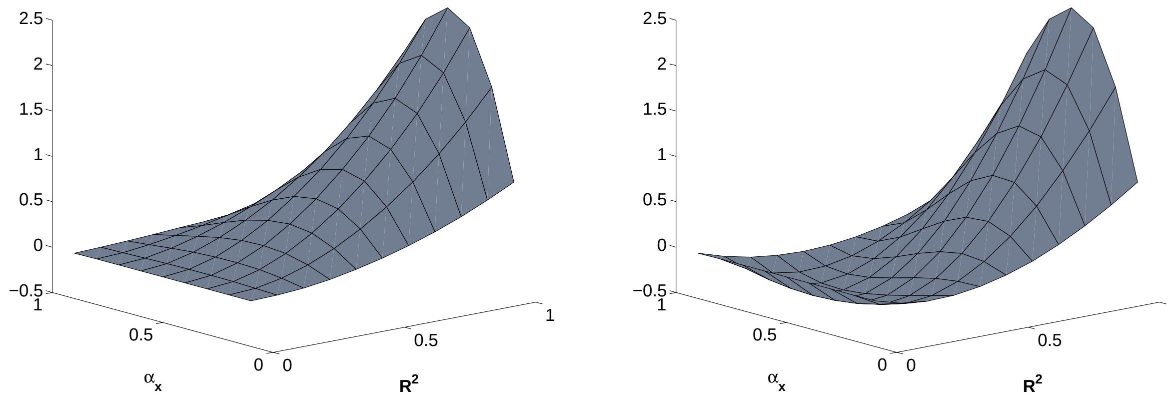

Note: Differences in expected losses generated by optimal policy in the simple New Keynesian model of Section 3 in the main paper under full information and the two hidden information cases. Positive values indicate that policy losses are larger in the full information case, Each panel considers variations in $\alpha_x = 1 - \alpha_\pi$ and R^2. All other model parameters are held at their baseline values, listed in Table 1.

G.2 Expected Losses

Figure 4 displays the difference in expected losses generated by optimal policy in the simple model of Section 3 in the main paper under full information and the two hidden information cases for the same set of alternative calibrations described above. Comparing the full-information case against the hidden information with known lagged state values ("the static case"), losses are always larger under full information as predicted by Proposition 3 and shown in Panel (a) of the figure. For the set of calibrations considered here, the same holds also for the comparison of full information against the dynamic case of imperfect information, as shown in Panel (b). Also, similar results (not shown) can be generated for alternative calibrations of the remaining model parameters that are held fixed in the figure.

Finally, to complement the discussion of the dynamic case of hidden information in Appendix A.3 of the main paper, Figure 5 displays the non-monotone relationship between the exogenous variance shares $R^2 = \sigma_\eta^2/(\sigma_\eta^2 + \sigma_\varepsilon^2)$ and $\boldsymbol{R^2} = \text{Var}\,(\tilde{\gamma}_t)/(\text{Var}\,(\tilde{\gamma}_t) + \sigma_\varepsilon^2)$. As discussed in the context of equation (36) of the main paper, due to the influence of endogenous policy choices on the output gap's signal-to-noise ratio for the public's inference about the persistent target component, the endogenous variance share $\boldsymbol{R^2}$ can be larger or smaller than the exogenous ratio of shock variances R^2. For the calculation of expected losses, the latter is the appropriate probability weight in the static case but the former needs to be used in the dynamic case.

Figure 5: Comparison of R^2 and $\boldsymbol{R^2}$ in the dynamic case of hidden info

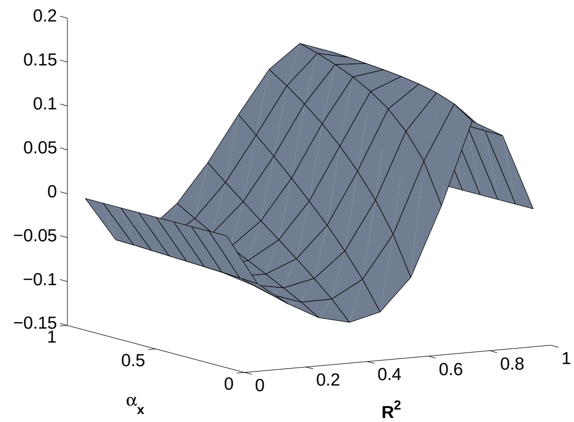

Note: Difference between $R^2 = \sigma_\eta^2/(\sigma_\eta^2+\sigma_\varepsilon^2)$ — the share of persistent shocks — and $\boldsymbol{R^2} = \text{Var}\,(\tilde{\gamma}_t)/(\text{Var}\,(\tilde{\gamma}_t)+\sigma_\varepsilon^2)$ — the share of persistent innovations in the Kalman filter— in the dynamic case of hidden information for the simple model of Section 3 in the main paper.

Part III
An Extended Model

This part applies the general LQ framework described in Part I of this teachnical appendix to a richer model, featuring multiple sources of hidden information and more endogenous state variables than the simple New Keynesian model studied in the main paper. Key elements of this extended model are the use of the nominal interest rate — not the output gap — as the control instrument for monetary policy. Furthermore, the models embeds the assumption that both wages and prices are sticky (EHL) and subject to indexation. The latter assumption introduces the lagged values of wage and price inflation as endogenous state variables. Finally, whereas the simple New Keynesian example discussed in the main paper considered only hidden information about policy preferences, this extended model also illustrates the interplay of hidden information about economic fundamentals and policy targets. In particular, the policymaker is assumed to observe the expected growth rate of potential output — and thus natural rate of interest — whilst the public can only observe the current level of potential output. In this environment, an optimizing policymaker needs to consider how policy rate changes — that may be triggered either by a shocks to the natural rate or policy targets — will cause the public to update its expectations about both economic fundamentals and perceived policy targets. Compared to the simple example studied in the main text, this is a substantially richer setting for studying the benefits of transparency and the design of optimal policy when the policymaker is (potentially) better informed than the public. Nevertheless, this is still a highly stylized model, mostly chosen to illustrate the versatility of the general LQ framework, described in Part I above.

H Optimal Disinflation when Inflation is Sticky

The framework of Part I is applied here to a long-standing question in monetary economics: How to disinflate to a permanently lower inflation rate? Based on historical disinflations Sargent (1982) and Bordo et al. (2007) have argued in favor of disinflating more quickly, when credibility is at stake, while other economists, for example Gordon (1982), have rather argued for prolonged and modest disinflation paths.

If inflation inertia were entirely due to a lack of transparency and the public's slow learning about the new target, a rapid disinflation — even a "cold turkey" strategy which forces inflation to the new target at all cost — might seem highly advisable, since it should anchor public expectations much more swiftly around the new target. Erceg and Levin (2003) and Goodfriend and King (2005) have for example characterized the Volcker disinflation of the early 1980s, as having been hampered — at least initially — by public concerns about policymaker's intentions to permanently lower the inflation rate in the U.S.[16] However, if inflation persistence should be due to some innate features of price- and wage setters' contracting arrangements — as argued by Gordon (1982) — a rapid disinflation might be too costly in terms of lost output.

Without taking a stand on the true nature of inflation persistence, this section derives the optimal time-consistent policy under different levels of (exogenous) inflation inertia.[17] In contrast to the previous literature, my approach takes explicitly into account the signaling role of monetary policy.[18] As elsewhere in this paper, optimal policy is required to be time-consistent, which seems a particularly sensible in the context of analyzing disinflation policies when credibility and a lack transparency are at stake. As will be seen, the required time-consistency limits the extent of credible signaling in equilibrium.

H.1 An Extended Model

The disinflation experiment in this section is built around the notion of an exogenous drop in the policymaker's inflation target. Specifically, the policymaker's loss function is now:

$$L_t = \alpha_p \left(\pi_t - \bar{\pi}_t\right)^2 + \alpha_x \left(x_t - \bar{x}_t\right)^2 \tag{43}$$

Changes in the inflation target are assumed to be permanent, and $\bar{\pi}_t$ follows a random walk.[19] An exogenously shifting inflation target can, of course, not yield a satisfactory account of *why* inflation

[16]Similarly, Ireland (1995) advocates rapid disinflations when credibility is at stake. Compared to my framework, private-sector beliefs are less sophisticated — without imposing rational expectations, or some other form of optimality, public beliefs of future policies are formed by extrapolating from past policies with a fixed rule — which limits the signaling role for a policymaker seeking to build credibility.

[17]Assuming perfect credibility and transparency, Ireland (1997) has addressed the question of optimal disinflation when there is state-dependent price setting. In his case, prices adjust more flexibly when inflation is larger, thus facilitating the conduct of rapid disinflations when target changes are large.

[18]To my knowledge, the sole exception is the recent work by Cogley et al. (2010) who analyze disinflation strategies, when the public employs least-squares learning to uncover the (new) rule chosen by the policymaker.

[19]The model is illustrated below, with simulations where the shock variance of this random walk has been calibrated to a relatively small value, such that target shifts are not a dominant source of business cycle variations in inflation and output, accounting for 15 and two percent in the bandpass-filtered fluctuations of output and inflation, respectively. While the calibration is purely illustrative, it might be worthwhile to notice that, in line with Ireland's (2007) estimates, the target shock explains about 70 and 60 percent of the forecast error variance of inflation and the nominal rate, respectively, at a horizon of 40 quarters, while the target shock accounts for barely none of the forecast error variance

rose and fell during the postwar history of the U.S. However, for the purpose of analyzing the design of optimal disinflation the abstraction seems reasonable.[20] As in the simple example described in the main paper, the stochastic output gap target consists of its two stationary components, with different levels of persistence.[21]

To illustrate the versatility of the general framework, there is not only staggered price setting but also staggered wage setting as in Erceg et al. (2000). Price inflation (π_t^p) and wage inflation (π_t^w) are characterized by two Phillips Curves, which relate current and anticipated future inflation rates to the deviations in output and the real wage from their values in the absence of sticky prices and wages. Erceg et al. (2000) show also that optimization of social welfare adds the stabilization of wage inflation to the social loss function. For simplicity, I have chosen to follow a dual mandate, which places equal weight on the stabilization of output and price inflation.[22] With sticky wages, there is also a trade-off between stabilizing inflation and the output gap in response to productivity shocks — as discussed by Blanchard and Galí (2007) — which will be introduced into the model further below.

Sticky prices and wages do not necessarily translate into inertial inflation rates. In order to generate additional inflation persistence, prices (and wages) are assumed to be indexed to their lagged values at the rate γ, when they cannot be re-optimized in the current period. If $\gamma = 1$, contracts are fully indexed to lagged prices (and wages) as in Christiano et al. (2005). For $\gamma = 0$, contracts are fully indexed to the public's current belief about the trend rate of inflation ($\bar{\pi}_{t|t}$).[23] This indexation scheme also avoids the problems arising from non-zero trend inflation studied by Ascari (2004) and Cogley and Sbordone (2008).[24] In sum, the Phillips Curves and indexation schemes for prices and wages are:[25]

$$\hat{\pi}_t^j = \beta\, \hat{\pi}_{t+1|t}^j + \kappa_j\, x_t + \lambda_j\, \omega_t \qquad j \in \{p, w\} \tag{44}$$

$$\hat{\pi}_t^j = \pi_t^j - \gamma \pi_{t-1}^j - (1-\gamma)\bar{\pi}_{t|t} \tag{45}$$

where x_t and ω_t denote the gaps in output and the real wage.[26]

Crucially, when only a small indexation weight is placed on lagged prices (and wages), inflation can adjust almost without cost to changes in the public's beliefs about the true target rate. If the

of output at this long horizon. (Not surprisingly, the short run shares explained by the target shock are much smaller than in Ireland's full information model.

[20]In a similar vein, Erceg and Levin (2003) characterize elements of the Volcker disinflation by an unobserved intercept shift in the policymaker's interest rate rule. Empirically, Ireland (2007) finds that an estimated New Keynesian model, solved under perfect information, fits the postwar U.S. data, including the high inflation periods of the late 1970s.

[21]The output gap target thus follows an exogenous law of motion described by: $\bar{x}_t = \gamma_t + \varepsilon_t$, with $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$, $\gamma_{t+1} = \rho\, \gamma_t + \eta_{t+1}$, $\eta_t \sim N(0, \sigma_\eta^2)$, and $|\rho| < 1$.

[22]Qualitatively, this choice is largely inconsequential for the results presented below, when compared against a loss function which considers also the squared difference between wage inflation and $\bar{\pi}_t$ in addition to the output gap and the differences between price inflation and $\bar{\pi}_t$; both for the case when all three terms are equally weighted or in the case of using the micro-founded weights, derived by Erceg et al. (2000) under the assumption of perfect information.

[23]When $\bar{\pi}_t$ were a constant, the case of $\gamma = 0$ would also correspond to the indexation used by Yun (1996).

[24]Assuming a perfectly observable inflation target, Ireland (2007) uses the same indexation scheme, where the price adjustment depends on a weighted average of past inflation and the current target level.

[25]For simplicity, the same indexation rate is assumed to apply to both prices and wages.

[26]Wage gap dynamics depend on he difference in price and wage inflation as well as the exogenous growth rate in productivity, denoted Δa_t: $\omega_t = \omega_{t-1} + \pi_t^w - \pi_t^p - \psi \Delta a_t$.

Table 1: Model Calibration

Parameters for Model with Sticky Prices and Wages (Galì, 2008, Chapter 6)		
β	0.9900	Time preference
σ	1.0000	Risk Aversion / Inverse EIS
ϕ	1.0000	Inverse Frisch Labor Elasticity
κ_p	0.0262	Output gap coefficient in Phillips Curve for price inflation
κ_w	0.0490	Output gap coefficient in Phillips Curve for wage inflation
λ_p	0.0132	Wage gap coefficient in Phillips Curve for price inflation
λ_w	-0.0123	Wage gap coefficient in Phillips Curve for wage inflation
ψ_{ya}	1.0000	Sensitivity of potential output to productivity (consistent with $\sigma = 1$ above).
ψ_{wa}	1.0000	Sensitivity of natural wage to productivity (consistent with $\sigma = 1$ above).
Policy Preferences, $L_t = \pi_t^2 + \alpha_x(x_t + \bar{x}_t)^2$		
α_x	0.5000	Weight on output stabilization
α_p	0.5000	Weight on inflation stabilization
Process for Output Gap Target $\bar{x}_t = \gamma_t + \varepsilon_t$		
σ_ε	1.0000	Volatility of iid target component, $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$
ρ	0.9000	Persistence of target component, $\gamma_{t+1} = \rho\gamma_t + \eta_{t+1}$
σ_η	1.0000	Volatility of persistent target shocks, $\eta_t \sim N(0, \sigma_\eta^2)$
Process for Inflation Target $\bar{\pi}_t = \bar{\pi}_{t-1} + \varepsilon_t^\pi$		
σ_π	0.200	Volatility of target shock $\varepsilon^p i_t \sim N(0, \sigma_\pi^2)$
Productivity Process $\Delta a_t = \mu_{t-1} + \varepsilon_t^a$		
Var (Δa_t)	2.0000	Variance of technology growth Var (Δa_t) = Var $(\mu_t) + \sigma_a^2$
σ_a	1.0000	Volatility of iid component, $\varepsilon_t^a \sim N(0, \sigma_a^2)$
ρ_μ	0.2500	Persistence of conditional mean, $\mu_{t+1} = \rho_\mu \mu_t + \varepsilon_{t+1}^\mu$
σ_μ	0.9682	Volatility of shocks to conditional mean, $\varepsilon_t^\mu \sim N(0, \sigma_\mu^2)$

Notes: Private-sector parameters taken from Galì (2003)'s calibration to quarterly U.S. data. Shock variances are each normalized to unity and not intended to match the scale of any second moments.

inflation target were fully transparent, $\bar{\pi}_t = \bar{\pi}_{t|t}$, target changes would thus not incur any output costs. However, when the indexation weight on lagged prices is large, there is a strong backward-looking component in price and wage inflation, independently of the policy regime, and Gordon's arguments in favor of a more cautious disinflation might apply.

Monetary policy controls the nominal short-term interest rate (i_t), instead of the output gap, and the model is closed with the New Keynesian IS curve:

$$i_t = \pi_{t+1|t} + \sigma \left(\Delta x_{t+1|t} + \psi_{ya}\, \Delta a_{t+1|t} \right) \tag{46}$$

where σ is the intertemporal rate of substitution. As in Galì (2008, Chapter 6), potential output ($\bar{y}_t$) is exogenously given as a multiple of the level of technology ($\bar{y}_t = \psi_{ya}\, a_t$).

There are multiple sources of uncertainty in this model. As before, the output gap target ($\bar{x}_t$) has two components, and none of the policy targets shall be directly observable. In addition, the model also allows for disagreement between policymaker and public about the outlook for future productivity. Specifically, the growth rate of technology is assumed to follow a mixture process

$$\Delta a_t = \mu_{t-1} + \varepsilon_t^a \qquad \text{with} \quad \mu_t = \rho^\mu \mu_{t-1} + \varepsilon_t^\mu \qquad \varepsilon_t^a \sim N(0, \sigma^a) \qquad \varepsilon_t^\mu \sim N(0, \sigma^\mu)$$

The public can only directly observe the current and past levels of technology, a^t, while the policymaker sees μ_t and ε_t^a separately. As before, the public observes also the policy instrument, so that its measurement vector consists of one exogenous variable (a_t) as well as one endogenous variable (i_t).[27] In this simple model, there is a one-to-one relationship between the level of productivity and the level of potential output — that is the output level absent any nominal frictions. Having observed an unexpected change in the interest rate, the public has to filter out, whether this policy surprise was due to a change in the expected growth rate of technology — and thus the "natural rate of interest" (Woodford, 2003)[28] — or whether the surprise might be motivated by a shift in the policy targets. The model calibration is described in Table 1.

The remainder of this section considers the experiment of negative shock to $\bar{\pi}_t$, which permanently lowers the inflation target. First, it will be assumed that the public can only observe the histories a^t and i^t. The discussion will then turn to the implications of improved transparency about the inflation target, which will be modeled in the form of an exogenous, noisy signal about the true target level, similar to what has been studied in Section 4 of the main paper.

H.2 Disinflation at Different Degrees of Inflation Inertia

Figure 6 displays the economy's dynamic responses to a unit-sized shock in the inflation target at different levels of exogenous inertia, parameterized by γ.[29] Significantly, when γ is large the optimal policy contracts the economy more vigorously, as can be seen from the higher real rates in Panel (a) and the more negative output gap responses in Panel (b).[30] In this sense, the optimal policy bears resemblance to what has been advocated by Sargent (1982).

However, as can be seen from the two bottom panels, the optimal policy does not disinflate in a "cold-turkey" manner, by forcing inflation to drop down rapidly to the new target level. For any level of γ, learning is slow and public beliefs about the target ($\bar{\pi}_{t|t}$, Panel c) adjust only belatedly to the new level — taking about 24 quarters. Consequently, price inflation (π_t^p, Panel d) and wage inflation (not shown) fall only in a similarly protracted manner.

The slow speed of the disinflation is instructive about the limits imposed on optimal policy by the required time-consistency. The policymaker *could* disinflate swiftly and at low cost, when γ is low *and* when the interest rate were a perfect signal for the inflation target, for example by setting

[27] Having the public observe the level of technology is important here, to preserve the assumption of a homogeneously informed household, who should be aware of the level of his productivity level when working.

[28] The natural rate of interest corresponds to the real interest rate that would prevail in the absence of nominal frictions, i.e. when $\pi_t = 0$ and $x_t = 0$ for all t, and is thus proportional to the expected growth rate of productivity $\bar{r}_t = \sigma \Delta E_t \bar{y}_{t+1} = \sigma \psi_{ya} \mu_t$.

[29] These responses are different over a grid of values for γ ranging from zero to one.

[30] Not shown in the figure are the responses of the nominal rate — which is however displayed in Figure 7 below. As the level of inflation is ultimately lowered in these simulations. the nominal rate falls during the disinflation (after an initial rise) because of the Fisher equation. For brevity, Figure 6 displays thus only the real rate, since it provides a cleaner picture about the stance of policy. (Please note that the inflation trajectories shown in the figure correspond to expectations $E_0 \pi_{t+k}$ whereas the real rate is constructed from the public's beliefs $r_t = i_t - \pi_{t+1|t}$.

$i_t = \bar{\pi}_t + \sigma\, \psi_{ya}\, \Delta a_{t+1|t}$.[31] Obviously such a policy would come at the expense of ignoring — and thus missing – the output targets and eschewing any policy signals about the natural rate. Even if such a policy rule might overall be beneficial, it is not necessarily time consistent, and — absent any direct communication about the inflation target as analyzed further below — the policymaker has to trade off the benefits of a clear inflation signal against his other objectives.[32] Despite the differences in real contractions across different values of γ, shown in Panels (a) and (b), the optimal policy achieves fairly similar outcomes in terms of inflation and target beliefs, shown in the bottom panels of the figure. While it is not time-consistent for the optimal policy to go "cold turkey", even when $\gamma = 0$, the optimal policy raises the real rate more strongly, when γ is large,

[31]Notice that this policy conditions only on public beliefs about the growth rate of productivity, ($\Delta a_{t+1|t}$ instead of $E_t \Delta\, a_{t+1}$), thus providing a signal about no other variable than the inflation target.

[32]Furthermore, even if such a policy should be desirable, its implementation as a rule of the form $i_t = \bar{\pi}_t + \sigma\, \psi_{ya}\, \Delta a_{t+1|t}$ would not uniquely determine the equilibrium and would require to be augmented as discussed in Galì (2011).

Figure 6: Disinflations and At Different Levels of Inflation Inertia

(a) Real Rate

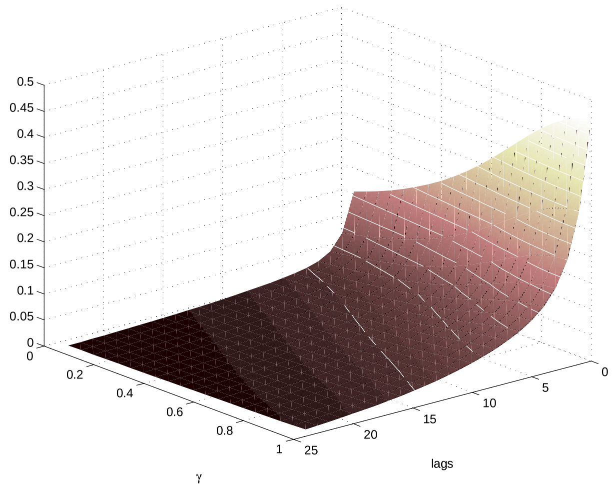

(b) Output Gap

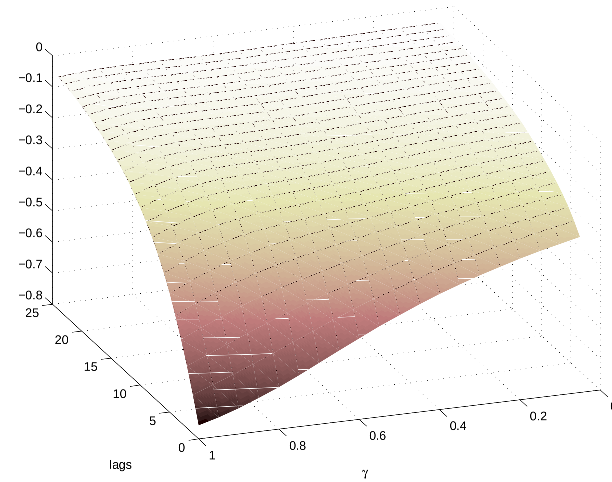

(c) Target Belief $\bar{\pi}_{t|t}$

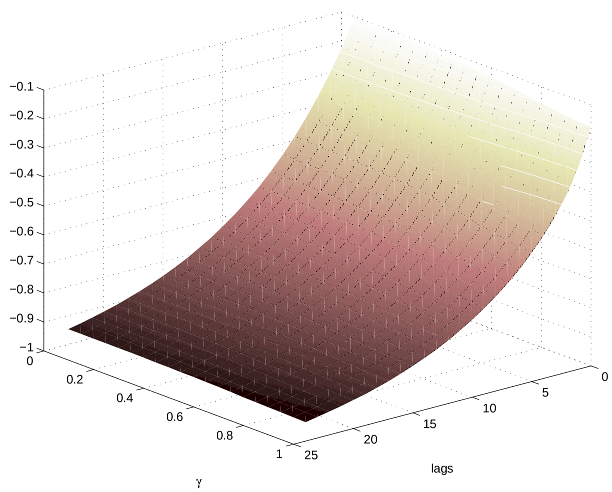

(d) Price Inflation π_t^p

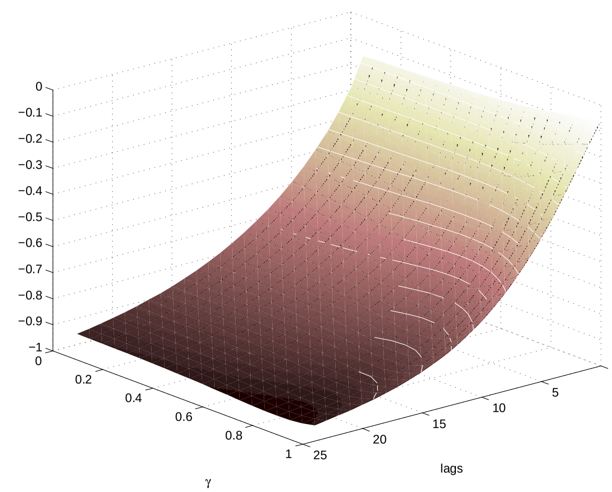

Note: Impulse response of selected variables — derived from the extended model of Section H — to a unit decrease in the inflation target at time $t = 0$. The parameter γ represents the weight on lagged inflation in the indexation of prices and wages. For $\gamma = 1$ prices and wages are indexed completely to lagged prices and wage, respectively. For $\gamma = 0$, prices and wages are indexed completely to the public's belief in the current inflation target, $\bar{\pi}_{t|t}$. For intermediary values, the indexation places a weight of γ on lagged prices (wages) and the remainder on target beliefs. Please note that the axes on Panel (b) are rotated differently than in the other panels, in order to provide a better perspective on the responses of the output gap. Since these simulation are holding fixed the level of potential output, changes in the output gap are identical to changes in output.

H.3 Transparency about the Inflation Target

In the previous simulations, the public could learn about the level of the inflation target only by observing the nominal interest rate (the policy instrument), and there were no other means of communicating this central state variable. The remainder of this section considers now variations in the degree of transparency about the inflation target, by adding an exogenous, noisy signal to the public's information set, which continues to include productivity (a^t) and the nominal rate (i^t):

$$s_t = \bar{\pi}_t + n_t \qquad \text{with} \quad n_t \sim N(0, \sigma_n^2) \tag{47}$$

As discussed in Section 4 of the main paper, variations in the noise variance σ_n^2, translate into varying degrees of target transparency in an "ex-ante" sense; it is "ex-ante" since the signal is observed independently of the policy instrument and the policymaker may or may not want to increase this transparency via the design of his interest rate policy.

Figure 7 displays the economy's impulse responses for three cases: Full transparency ($\sigma_n^2 = 0 \Leftrightarrow s_t = \bar{\pi}_t = \bar{\pi}_{t|t}$), low transparency ($\sigma_n^2 \to \infty$), and the intermediate case of an imperfect, but informative signal about the inflation target. Throughout, these simulations assume a high level of inflation inertia ($\gamma = 1$), in order to assess the effects the transparency regime, in the case that Gordon (1982) has been concerned with.[33] Also, even when s_t is a perfect signal of the inflation target, the public remains uncertain about the natural rate and the state of the policymaker's output target.

As shown in the top panels of the figure, nominal and real rate are raised more strongly, when the inflation target lacks transparency. In fact, when the target is directly observed by the public, the optimal policy is less contractionary (and hence there is less growth back to the level of potential output) and the Fisher-effect of lower inflation expectations dominates, such that the optimal nominal rate falls when the inflation target decreases; in this case the real rate rises only barely. Only under imperfect information, nominal and real rate increase at the onset of the disinflation, and the real rate stays elevated until public beliefs $\bar{\pi}_{t|t}$ have converged to the new target level. The consequences of strong backward-looking component ($\gamma = 1$) are evident in the simulated paths for price inflation, which decrease only slowly under all three transparency regimes.[34]

Finally, it should be stressed, that policy losses are strictly higher, the lower the transparency of the inflation target, since inflation reaches the target later and output needs to contract stronger.

This results is particularly interesting, when compared against the simple model of the main paper, where intransparent *output* targets were actually beneficial for lower policy losses — resembling the findings of Faust and Svensson (2001, 2002), Gaspar et al. (2006, 2010), Westelius (2009) in similar settings.[35] In contrast, the results for the case of an intransparent inflation target — sensibly — document that intransparency is not beneficial per se, and that results derived for a certain setting and a certain class of shocks, need not carry over to other cases.

[33] The results shown in Figure 6 for $\gamma = 1$ correspond to what is shown on Figure 7 for the case of an infinite noise variance, $\sigma_n^2 \to \infty$.

[34] Wage inflation (not shown) decreases to the new target level at a similarly slow pace.

[35] In contrast to the other references, the model of Gaspar et al. (2006, 2010) is driven by cost-push shocks with unknown persistence, instead of output targets, but their mechanism is very similar.

Figure 7: Optimal Disinflations when Inflation is Highly Inertial ($\gamma = 1$)

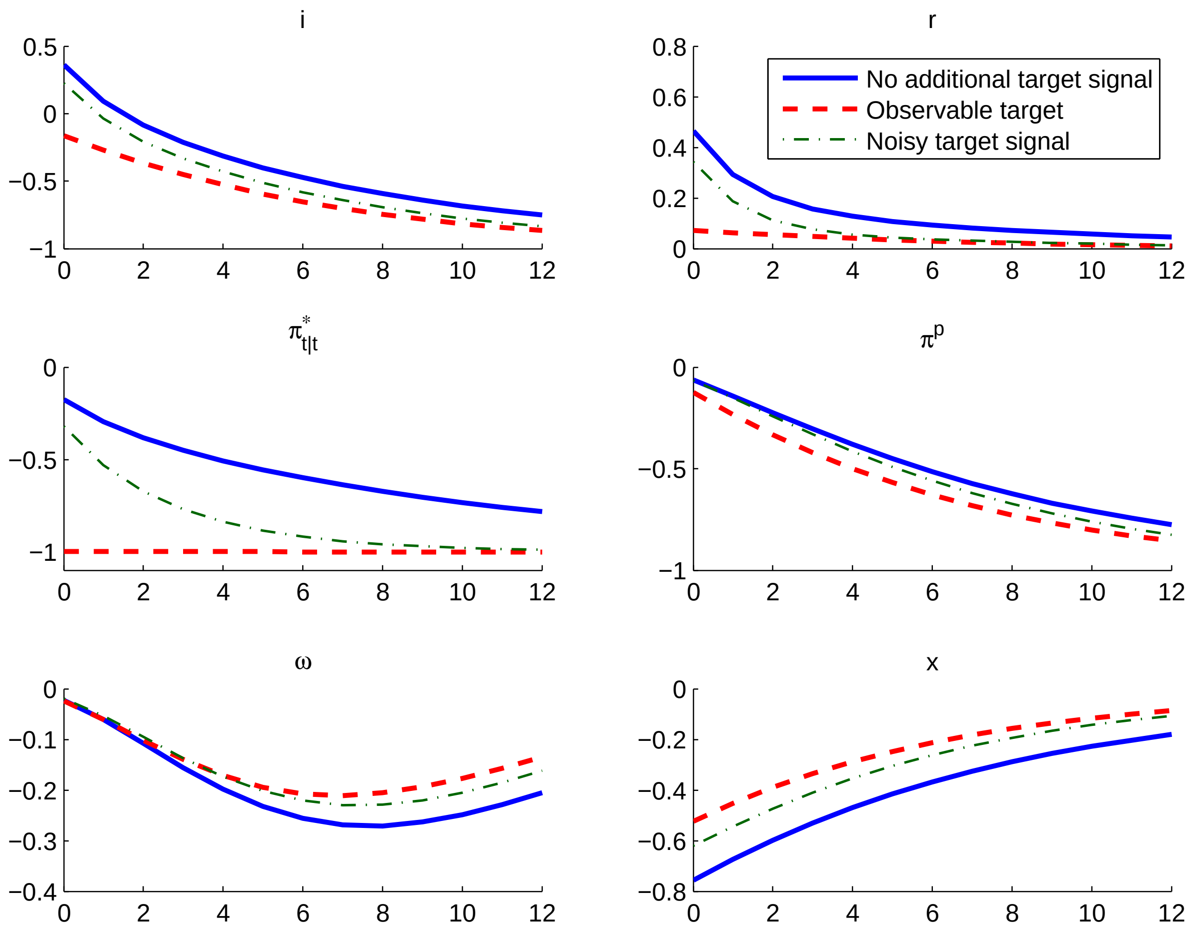

Note: Impulse response to a unit decrease in the inflation target for three different calibrations of the noise variance in the signal $s_t = \bar{\pi}_t + n_t$: all derived under the assumption of perfect indexation to lagged prices ($\gamma = 1$). Thick solid lines (blue) correspond to the case of a useless signal ($\sigma_n \to \infty$). These response are identical to what is shown in Figure 6 for $\gamma = 1$. Dashed lines (red) denote the model's dynamic responses when the target is publicly observable ($\sigma_n^2 = 0 \Rightarrow \bar{\pi}_{t|t} = \bar{\pi}_t$), while the intermediate case is displayed by the dash-dotted lines (green). For the intermediate case, the noise variance has been chosen such that the signal s_t explains about 10% of variations in the target innovations ($\bar{\pi}_t - \bar{\pi}_{t|t-1}$), which induces the optimal policy to implement outcomes, that are already fairly close to the case of an observable target.

References

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

Ascari, G. (2004, July). Staggered prices and trend inflation: Some nuisances. *Review of Economic Dynamics 7*(3), 642–667.

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

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

Blanchard, O. and J. Galí (2007, 02). Real wage rigidities and the New Keynesian model. *Journal of Money, Credit and Banking 39*(s1), 35–65.

Bordo, M. D., C. Erceg, A. Levin, and R. Michaels (2007, March). Three great American disinflations. NBER Working Papers 12982, National Bureau of Economic Research, Inc.

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

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005, February). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy 113*(1), 1–45.

Cogley, T., C. Matthes, and A. M. Sbordone (2010, November). Optimal disinflation under learning. *mimeo* New York University.

Cogley, T. and A. M. Sbordone (2008, December). Trend inflation, indexation, and inflation persistence in the New Keynesian Phillips curve. *The American Economic Review 98*(5), 2101–26.

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics 46*, 281–313.

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

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

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

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

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

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

Galì, J. (2011). Are central banks' projections meaningful? *Journal of Monetary Economics 58*(6 - 8), 537 – 550.

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

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

Goodfriend, M. and R. G. King (2005, July). The incredible Volcker disinflation. *Journal of Monetary Economics 52*(5), 981–1015.

Gordon, R. J. (1982). Why stopping inflation may be costly: Evidence from fourteen historical episodes. See Hall (1982), Chapter 2, pp. 11 – 40.

Grandmont, J. M. (1977, April). Temporary general equilibrium theory. *Econometrica 45*(3), 535–572.

Hall, R. E. (Ed.) (1982). *Inflation: Causes and Effects*. Chicago: The University of Chicago Press.

Hansen, L. P. and T. J. Sargent (2007). *Robustness*. Princeton University Press.

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

Ireland, P. N. (1997, November). Stopping inflations, big and small. *Journal of Money, Credit and Banking 29*(4), 759–775.

Ireland, P. N. (2007, December). Changes in the Federal Reserve's inflation target: Causes and consequences. *Journal of Money, Credit and Banking 39*(8), 1851–1882.

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

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

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

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

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

Sargent, T. J. (1982). The end of four big inflations. See Hall (1982), Chapter 3, pp. 41 – 97.

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

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

Sargent, T. J. and N. Wallace (1975, April). "Rational" expectations, the optimal monetary instrument, and the optimal money supply rule. *The Journal of Political Economy 83*(2), 241–254.

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

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

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

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

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