Available online at www.sciencedirect.com
SciVerse ScienceDirect
Journal of Economic Theory 147 (2012) 1–42
JOURNAL OF Economic Theory
www.elsevier.com/locate/jet

Linear-quadratic approximation of optimal policy problems ☆

Pierpaolo Benigno [a,b], Michael Woodford [c,*]

[a] Dipartimento di Economia e Finanza, LUISS Guido Carli, viale Romania 32, 00197 Roma, Italy
[b] EIEF, Italy
[c] Department of Economics, Columbia University, 420 W. 118th Street, New York, NY 10027, USA

Received 6 August 2008; final version received 9 September 2011; accepted 11 September 2011

Available online 18 October 2011

Abstract

We consider a general class of nonlinear optimal policy problems with forward-looking constraints, and show how to derive a problem with linear constraints and a quadratic objective that approximates the exact problem. The solution to the LQ approximate problem represents a local linear approximation to optimal policy from the "timeless perspective" proposed in Benigno and Woodford (2004, 2005) [6,7], in the case of small enough stochastic disturbances. We also derive the second-order conditions for the LQ problem to have a solution, and show how to correctly rank alternative simple policy rules, again in the case of small enough shocks.

JEL classification: C61; C63; E61

Keywords: Perturbation methods; Timeless perspective; Policy rules

☆ An earlier version of this paper was presented as a Plenary Lecture at the 10th Annual Conference on Computing in Economics and Finance, Amsterdam, July 2004. We would like to thank Filippo Altissimo, Vasco Cúrdia, Wouter Den Haan, Ken Judd, Jinill Kim, Andy Levin, Paul Levine, Diego Rodriguez Palenzuela, Joseph Pearlman, Willi Semmler, Harald Uhlig, and two anonymous referees for many helpful comments, and the National Science Foundation for research support through a grant to the NBER.
* Corresponding author.
E-mail addresses: pbenigno@luiss.it (P. Benigno), mw2230@columbia.edu (M. Woodford).

0022-0531/$ – see front matter
doi:10.1016/j.jet.2011.10.012

Linear-quadratic (LQ) optimal-control problems have been the subject of an extensive literature.[1] General characterizations of their solutions and useful numerical algorithms to compute them are now available, allowing models with fairly large state spaces, complicated dynamic linkages, and a range of alternative informational assumptions to be handled. And the extension of the classic results of the engineering control literature to the case of forward-looking systems of the kind that naturally arise in economic policy problems when one allows for rational expectations on the part of the private sector has proven to be fairly straightforward.[2]

An important question, however, is whether optimal policy problems of economic interest should take this convenient form. It is easy enough to apply LQ methodology if one specifies an *ad hoc* quadratic loss function on the basis of informal consideration of the kinds of instability in the economy that one would like to reduce, and posits linear structural relations that capture certain features of economic time series without requiring these relations to have explicit choice-theoretic foundations, as in early applications to problems of monetary policy. But it is highly unlikely that the analysis of optimal policy in a DSGE model will involve either an exactly quadratic utility function or exactly linear constraints.

We shall nonetheless argue that LQ problems can usefully be employed as approximations to exact optimal policy problems in a fairly broad range of cases. Since an LQ problem necessarily leads to an optimal decision rule that is linear, the most that one could hope to obtain with any generality would be for the solution to the LQ problem to represent a *local linear approximation* to the actual optimal policy — that is, a first-order Taylor approximation to the true, nonlinear optimal policy rule. In this paper we present conditions under which this will be the case, and show how to derive an LQ approximate problem corresponding to any member of a general class of optimal policy problems.

The conditions under which the solution to an LQ approximate problem will yield a correct local linear approximation to optimal policy are in fact more restrictive than might be expected, as noted for example by Judd [28, pp. 536–539], [29, pp. 507–508]. In particular, it does *not* suffice that the objective and constraints of the exact problem be continuously differentiable a sufficient number of times, that the solution to the LQ approximate problem imply a stationary evolution of the endogenous variables, and that the exogenous disturbances be small enough (though each of these conditions is obviously *necessary*, except in highly special cases). An approach that simply computes a second-order Taylor-series approximation to the utility function and a first-order Taylor-series approximation to the model structural relations in order to define an approximate LQ problem — the approach criticized by Judd [28,29] that we have elsewhere (Benigno and Woodford [9]) called "naive LQ approximation" — may yield a linear policy rule with coefficients very different from those of a correct linear approximation to the optimal policy in the case of small enough disturbances.[3]

The discussion by Judd [28, pp. 536–553] might seem to imply that LQ approximation is an inherently mistaken idea — that it cannot be expected, other than in cases so special as to represent an essentially fortuitous result, to yield a correct approximation to optimal policy at all. Nonetheless, it is quite generally possible to construct an alternative quadratic objective function that *will* result in a correct local LQ approximation, in the sense that the linear solution to the LQ problem is a correct linear approximation to the solution to the exact problem. The correct

[1] See Kendrick [32] for an overview of the use of LQ methods in economics.

[2] See, *e.g.*, Backus and Driffill [2] for a useful review.

[3] For an example illustrating this possibility, see Benigno and Woodford [9]. The same problem can also result in incorrect welfare rankings of alternative simple policies, as discussed by Kim and Kim [34,35].

method was illustrated in the important paper of Magill [40], that applied results of Fleming [23] from the optimal-control literature to derive a local LQ approximation to a continuous-time multi-sector optimal growth model. Here we show how the method of Magill can be used in the context of discrete-time dynamic optimization problems where some of the structural relations are forward-looking, as is almost inevitably the case in optimal monetary or fiscal policy problems.[4]

Of course the problems that can arise as a result of "naive" LQ optimization can also be avoided through the use of alternative perturbation techniques, as explained by Judd. Approaches that are widely used in the recent literature on policy analysis in DSGE models include either (i) deriving the first-order conditions that characterize optimal (Ramsey) policy using the exact (nonlinear) objective and constraints, and then log-linearizing *these conditions* in order to obtain an approximate solution to them, rather than separately approximating the objective and constraints before deriving the first-order conditions[5]; or (ii) obtaining a higher-order (at least second-order) perturbation solution for the equilibrium implied by a given policy by solving a higher-order approximation to the constraints, and then evaluating welfare under the policy using this approximate solution.[6] These methods can also be used to correctly calculate a linear approximation to the optimal policy rule, and when applied to the problem considered here, provide alternative approaches to calculating the same solution.[7]

Despite the existence of these alternative perturbation approaches to the analysis of optimal policy, we believe that it remains useful to show how a correct form of LQ analysis is possible in the case of a fairly general class of problems. One reason is that the ability to translate a policy problem into this form allows one to use the extensive body of theoretical analysis and numerical techniques that have been developed for LQ problems. Another is that casting optimal policy analysis in DSGE models in this form can allow comparisons between welfare-based policy analysis and analyses of optimal policy based on *ad hoc* stabilization objectives (which have often been expressed as LQ problems). We also show that the LQ formulation of the approximate policy problem makes it possible to rank suboptimal policy rules by a criterion that is consistent with the characterization given of optimal policy. And finally, the LQ approximation makes it straightforward to analyze whether a solution to the first-order conditions for optimal policy also satisfies the relevant second-order conditions for optimality; one need simply check the algebraic conditions required for concavity of the quadratic objective in the approximating LQ problem, discussed in Section 3 of Benigno and Woodford [10]. Essentially, the calculations required in order to derive the LQ approximation are ones that would be required in any event to check

[4] Levine et al. [39] also discuss the application of the method of Magill to discrete-time problems. Their formal results, however, apply only to deterministic problems with purely backward-looking constraints. They argue (correctly) that their results can also be applied to problems with forward-looking constraints, under an assumption that the policymaker can credibly commit itself in advance, but they do not develop in the detail that we do here the technicalities involved in such applications. They also present only the Lagrangian approach that we describe in Section 1.3; in Section 1.2 we present an alternative approach to the derivation of a correct LQ approximation, which may provide additional insights, and we establish the equivalence of these two approaches. Finally, Levine et al. do not discuss the use of an LQ approximation to rank non-optimal policy rules; we treat this issue below in Section 3.

[5] See, *e.g.*, King and Wolman [38], Khan et al. [33], or Schmitt-Grohé and Uribe [46].

[6] For methods for executing computations of this kind, see Jin and Judd [27], Schmitt-Grohé and Uribe [45], and Kim et al. [36]. For an application to the analysis of optimal policy, see, *e.g.*, Schmitt-Grohé and Uribe [47].

[7] As shown in Section 1.3 below, our method computes the same coefficients for a linear policy rule as are obtained by linearization of the first-order conditions for the exact policy problem. The general intuition for this result is discussed in Section 1 of Benigno and Woodford [10].

whether a solution to the (exact, nonlinear) first-order conditions for optimal represents at least a local welfare maximum; these calculations only appear to be unnecessary under other numerical approaches because it is so common for authors to neglect the issue of second-order conditions.

In Section 1, we present a general class of dynamic optimization problems with forward-looking constraints, and derive an LQ approximate problem associated with any problem in this class. Section 2 discusses the general algebraic form of the first- and second-order conditions for optimality in the LQ approximate problem. Section 3 shows how the quadratic objective for stabilization policy derived in Section 1 can also be used to compute welfare comparisons between alternative sub-optimal policies, in the case that the stochastic disturbances are small enough. Finally, Section 4 discusses applications of the method described here and concludes.

1. LQ approximation of a problem with forward-looking constraints

We wish to consider an abstract discrete-time dynamic optimal policy problem of the following sort. Suppose that the policy authority wishes to determine the evolution of an (endogenous) state vector $\{y_t\}$ for $t \geqslant t_0$ to maximize an objective of the form

$$V_{t_0} \equiv E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \pi(y_t, \xi_t), \tag{1}$$

where $0 < \beta < 1$ is a discount factor, the period objective $\pi(y, \xi)$ is a concave function of y, and ξ_t is a vector of exogenous disturbances. The evolution of the endogenous states must satisfy a system of backward-looking structural relations

$$F(y_t, \xi_t; y_{t-1}) = 0 \tag{2}$$

and a system of forward-looking structural relations

$$E_t g(y_t, \xi_t; y_{t+1}) = 0, \tag{3}$$

that both must hold for each $t \geqslant t_0$, given the vector of initial conditions y_{t_0-1}.

Conditions of the form (2) allow current endogenous variables to depend on lagged states; for example, these relations could include a technological relation between the capital stock carried into the next period, current investment expenditure, and the capital stock carried into the current period. Conditions of the form (3) instead allow current endogenous variables to depend on current expectations regarding future states; for example, these relations could include an Euler equation for the optimal timing of consumer expenditure, relating current consumption to expected consumption in the next period and the expected rate of return on saving. While the most general notation would allow both leads and lags in all of the structural equations, supposing that there are equations of these two types will make clearer the different types of complications arising from the two distinct types of intertemporal linkages. We shall suppose that the number n_F of constraints of the first type each period plus the number n_g of constraints of the second type is less than the number n_y of endogenous state variables each period, so that there is at least one dimension along which policy can continuously vary the outcome y_t each period, given the past and expected future evolution of the endogenous variables. A t_0-optimal commitment (the standard Ramsey policy problem) is then the state-contingent evolution $\{y_t\}$ consistent with Eqs. (2)–(3) for all $t \geqslant t_0$ that maximizes (1).

1.1. Optimal policy from a "Timeless perspective"

As is well known, the presence of the forward-looking constraints (3) implies that a t_0-optimal commitment is not generally time-consistent. If, however, we suppose that a policy to apply from period t_0 onward must be chosen subject to an additional set of constraints on the acceptable values of y_{t_0}, it is possible for the resulting policy problem to have a recursive structure.[8] While this is not necessary for the method of LQ approximation to be applicable, it is necessary in order for both our approximate quadratic objective and approximate linear constraints to involve coefficients that are time-invariant, and correspondingly for our derived linear approximation to optimal policy to involve time-invariant coefficients, as is discussed further in Section 2.2 below.

As discussed in Benigno and Woodford [6,7], in order to obtain a problem with a recursive structure (the solution to which can be described by a time-invariant policy rule), we must choose initial pre-commitments regarding y_{t_0} that are *self-consistent*, in the sense that the policy that is chosen subject to these constraints would also satisfy constraints of exactly the same form in all later periods as well. The required initial pre-commitments are of the form

$$g(y_{t_0-1}, \xi_{t_0-1}; y_{t_0}) = \bar{g}_{t_0}, \tag{4}$$

where $\bar{g}_{t_0}$ may depend on the exogenous state at date t_0. Note that we assume the existence of a pre-commitment only about those aspects of y_{t_0} the anticipation of which back in period $t_0 - 1$ should have been relevant to equilibrium determination then; there is no need for any stronger form of commitment in order to render optimal policy time-consistent.

We are thus interested in characterizing the state-contingent policy $\{y_t\}$ for $t \geqslant t_0$ that maximizes (1) subject to constraints (2)–(4). Such a policy is *optimal from a timeless perspective* if $\bar{g}_{t_0}$ is chosen, as a function of predetermined or exogenous states at t_0, according to a self-consistent rule.[9] This means that the initial pre-commitment is determined by past conditions through a function

$$\bar{g}_{t_0} = \bar{g}(\xi_{t_0}, \mathbf{y_{t_0-1}}), \tag{5}$$

where $\mathbf{y_t}$ is an extended state vector[10]; this function has the property that under optimal policy, given this initial pre-commitment, the state-contingent evolution of the economy will satisfy

$$g(y_{t-1}, \xi_{t-1}; y_t) = \bar{g}(\xi_t, \mathbf{y_{t-1}}) \tag{6}$$

in each possible state of the world at each date $t > t_0$ as well. Thus the initial constraint is of a form that one would optimally commit oneself to satisfy at all subsequent dates.

[8] Marcet and Marimon [41] propose an alternative approach that modifies the policy objective by adding additional multiplier terms; the additional terms in the objective of the modified problem of Marcet and Marimon lead to the same additional terms in the Lagrangian for the policy problem as the additional constraints that we introduce here. We prefer to introduce initial pre-commitments because of the more transparent connection of the modified problem to the original policy problem under the present exposition. The first-order conditions for optimal policy in the recursive policy problem that we propose are the same as those derived by Marcet and Marimon, except in the initial period.

[9] See Benigno and Woodford [7], Giannoni and Woodford [24], or Woodford [52] for further discussion.

[10] The extended state vector may include both endogenous and exogenous variables, the values of which are realized in period t or earlier. More specific assumptions about the nature of the extended state vector are made below; see the discussion of Eq. (8).

Let $V(\bar{g}_{t_0}; y_{t_0-1}, \xi_{t_0}, \xi_{t_0-1})$ be the maximum achievable value of the objective (1) in this problem.[11] Then the infinite-horizon problem just defined is equivalent to a sequence of one-period decision problems in which, in each period $t \geqslant t_0$, a value of y_t is chosen and state-contingent one-period-ahead pre-commitments $\bar{g}_{t+1}(\xi_{t+1})$ (for each of the possible states ξ_{t+1} in the following period) are chosen so as to maximize

$$\pi(y_t, \xi_t) + \beta E_t V(\bar{g}_{t+1}; y_t, \xi_{t+1}, \xi_t), \tag{7}$$

subject to the constraints

$$F(y_t, \xi_t; y_{t-1}) = 0,$$
$$g(y_{t-1}, \xi_{t-1}; y_t) = \bar{g}_t,$$
$$E_t \bar{g}_{t+1} = 0,$$

given the values of $\bar{g}_t$, y_{t-1}, ξ_{t-1}, and ξ_t, all of which are predetermined and/or exogenous in period t. It is this recursive policy problem that we wish to study; note that it is only when we consider this problem (as opposed to the unconstrained Ramsey problem) that it is possible, in general, to obtain a deterministic steady state as an optimum in the case of suitable initial conditions, and hence only in this case that we can hope to approximate the optimal policy problem around such a steady state.

The solution to the recursive policy problem just defined involves values for the endogenous variables y_t given by a policy function of the form

$$y_t = y^*(\bar{g}_t, y_{t-1}, \xi_t, \xi_{t-1}),$$

and a choice of the following period's pre-commitment $\bar{g}_{t+1}$ of the form

$$\bar{g}_{t+1} = g^*(\xi_{t+1}; \bar{g}_t, y_{t-1}, \xi_t, \xi_{t-1}),$$

where y^* and g^* are time-invariant functions. Let us suppose furthermore that the evolution of the extended state vector depends only on the evolution of the two vectors $\{y_t, \xi_t\}$, through a recursion of the form

$$\mathbf{y_t} = \psi(\xi_t, y_t, \mathbf{y_{t-1}}); \tag{8}$$

this system of identities *defines* the extended state vector, the elements of which consist essentially of linear combinations of current and lagged elements of the vectors y_t and ξ_t. (To simplify notation, we shall suppose that the current values y_t and ξ_t are among the elements of $\mathbf{y}_t$.) The initial pre-commitment (5) is then self-consistent if

$$g^*\big(\xi_{t+1}; \bar{g}(\xi_t, \mathbf{y_{t-1}}), y_{t-1}, \xi_t, \xi_{t-1}\big) = \bar{g}\big(\xi_{t+1}, \psi\big(\xi_t, y^*(\bar{g}_t, y_{t-1}, \xi_t, \xi_{t-1}), \mathbf{y_{t-1}}\big)\big) \tag{9}$$

for all possible values of ξ_{t+1}, ξ_t, and $\mathbf{y_{t-1}}$. Note that this implies that Eq. (6) is satisfied at all times.

[11] We assume, to economize on notation, that the exogenous state vector ξ_t evolves in accordance with a Markov process. Hence ξ_t summarizes not only all of the disturbances that affect the structural relations at date t, but all information at date t about the subsequent evolution of the exogenous disturbances. This is important in order for a time-invariant value function to exist with the arguments indicated.

1.2. A correct LQ local approximation

We now derive a corresponding LQ problem using local approximations to both the objective and the constraints of the above problem. In order for these local approximations to involve coefficients that remain the same over time, we compute them near the special case of an optimal policy that involves values of the state variables that are constant over time. This special case involves both zero disturbances and suitably chosen initial conditions; we then seek to approximately characterize optimal policy for nearby problems in which the disturbances are small and the initial conditions are *close* to satisfying the assumed special conditions. To be precise, we assume both an initial state y_{t_0-1} and initial pre-commitments $\bar{g}_{t_0}$ such that the optimal policy in the case of zero disturbances is a steady state, *i.e.*, such that $y_t = \bar{y}$ for all t, for some vector $\bar{y}$. (Our subsequent calculations then assume that both y_{t_0-1} and $\bar{g}_{t_0-1}$ are close enough to being consistent with this steady state.) In order to define the steady state, we must consider the nature of optimal policy in the exact problem just defined.

The first-order conditions for the exact policy problem can obtained by differentiating a Lagrangian of the form

$$\mathcal{L}_{t_0} = V_{t_0} + E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \big[\lambda_t' F(y_t, \xi_t; y_{t-1}) + \beta^{-1} \varphi_{t-1}' g(y_{t-1}, \xi_{t-1}; y_t)\big], \tag{10}$$

where λ_t and φ_t are Lagrange multipliers associated with constraints (2) and (3) respectively, for any date $t \geqslant t_0$, and we use the notation $\beta^{-1}\varphi_{t_0-1}$ for the Lagrange multiplier associated with the additional constraint (4). This last notational choice allows the first-order conditions to be expressed in the same way for all periods. Optimality requires that the joint evolution of the processes $\{y_t, \xi_t, \lambda_t, \varphi_t\}$ satisfy

$$\begin{aligned} &D_y\pi(y_t, \xi_t) + \lambda_t' D_y F(y_t, \xi_t; y_{t-1}) + \beta E_t \lambda_{t+1}' D_{\check{y}} F(y_{t+1}, \xi_{t+1}; y_t) \\ &\quad + E_t \varphi_t' D_y g(y_t, \xi_t; y_{t+1}) + \beta^{-1} \varphi_{t-1}' D_{\hat{y}} g(y_{t-1}, \xi_{t-1}; y_t) = 0 \end{aligned} \tag{11}$$

at each date $t \geqslant t_0$, where D_y denotes the vector of partial derivatives of any of the functions with respect to the elements of y_t, while $D_{\hat{y}}$ means the vector of partial derivatives with respect to the elements of y_{t+1} and $D_{\check{y}}$ means the vector of partial derivatives with respect to the elements of y_{t-1}.

An *optimal steady state* is then described by a collection of vectors $(\bar{y}, \bar{\lambda}, \bar{\varphi})$ satisfying

$$\begin{aligned} &D_y\pi(\bar{y}, 0) + \bar{\lambda}' D_y F(\bar{y}, 0; \bar{y}) + \beta \bar{\lambda}' D_{\check{y}} F(\bar{y}, 0; \bar{y}) + \bar{\varphi}' D_y g(\bar{y}, 0; \bar{y}) \\ &\quad + \beta^{-1} \bar{\varphi}' D_{\hat{y}} g(\bar{y}, 0; \bar{y}) = 0, \end{aligned} \tag{12}$$

$$F(\bar{y}, 0; \bar{y}) = 0, \tag{13}$$

$$g(\bar{y}, 0; \bar{y}) = 0. \tag{14}$$

We shall suppose that such a steady state exists, and assume (in the policy problem with random disturbances) an initial state y_{t_0-1} near $\bar{y}$, and an initial pre-commitment $\bar{g}_{t_0}$ near zero.[12] Once the optimal steady state has been computed, we make no further use of conditions (11); our proposed method does not require that we directly seek to solve these equations.

[12] Note that the steady-state value of $\bar{g}$ is equal to $g(\bar{y}, 0; \bar{y}) = 0$.

Instead, we now consider local approximations to the objective and constraints near an optimal steady state. These local approximations are of the following kind. We shall suppose that the vector of exogenous disturbances $\{\xi_t\}$ can be written as

$$\xi_t = \epsilon u_t \tag{15}$$

for all t, where $\{u_t\}$ is a bounded vector stochastic process, and we consider the one-parameter family of specifications corresponding to different values of the real number ϵ, holding fixed the stochastic process $\{u_t\}$. We are interested in computing approximations that become arbitrarily accurate in the case of any small enough value of ϵ, and we shall classify the order of any given approximation according to the rate at which the approximation error becomes small as ϵ approaches zero.

In the one-parameter family of problems indexed by ϵ, we shall also suppose that the initial conditions y_{t_0} and $\bar{g}_{t_0}$ vary with ϵ; specifically, we shall suppose that y_{t_0} and $\bar{g}_{t_0}$ are each equal to ϵ times some constant vector. Thus as ϵ approaches zero, the problem becomes one for which the optimal steady state is the solution. We wish to derive an approximate characterization of the solution near the optimal steady state in the case of any small enough value of ϵ.

We begin by considering local approximations to the objective and constraints near an optimal steady state. Suppose that under some policy of interest, the equilibrium evolution of the endogenous variables $\{y_t(\epsilon)\}$ in the case of the economy indexed by (any small enough value of) ϵ satisfies

$$y_t(\epsilon) = \bar{y} + \mathcal{O}(\epsilon) \tag{16}$$

at all times. An expression such as (16) means that the residual becomes small (in the sup norm[13]) as ϵ is made smaller, and approaches zero at (at least) the same rate as ϵ. Note that if we consider a policy rule consistent with the optimal steady state when $\epsilon = 0$, and the regularity condition is satisfied that allows the implicit function theorem to be used to solve the system of equations consisting of the structural equations plus the policy rule for the implied equilibrium evolution $\{y_t(\epsilon)\}$ for any small enough value of ϵ, then the solution will necessarily satisfy (16). Moreover, if the system of equations consisting of (11) together with the structural equations has a determinate solution, so that the implicit function theorem can be used to characterize optimal policy for small enough values of ϵ, (16) will also be satisfied. In what follows, we restrict our analysis to policies that satisfy (16); note that this allows us to consider any policy that is close enough to the optimal policy.

A second-order Taylor series expansion of the objective function π yields

$$\begin{aligned}\pi(y;\xi) &= \bar{\pi} + D_y\pi \cdot \tilde{y} + D_\xi\pi \cdot \xi + \frac{1}{2}\tilde{y}' D^2_{yy}\pi \cdot \tilde{y} + \frac{1}{2}\xi' D^2_{\xi\xi}\pi \cdot \xi + \tilde{y}' D^2_{y\xi}\pi \cdot \xi + \mathcal{O}(\epsilon^3) \\ &= D_y\pi \cdot \tilde{y} + \frac{1}{2}\tilde{y}' D^2_{yy}\pi \cdot \tilde{y} + \tilde{y}' D^2_{y\xi}\pi \cdot \xi + \text{t.i.p.} + \mathcal{O}(\epsilon^3),\end{aligned} \tag{17}$$

where $\tilde{y}_t \equiv y_t - \bar{y}$ and the various matrices of partial derivatives are each evaluated at $(\bar{y}; 0)$. (Here we use the fact that (15) and (16) imply that ξ_t and $\tilde{y}_t$ are each of order $\mathcal{O}(\epsilon)$.) The expression "t.i.p." refers to terms that are independent of the policy chosen (such as the constant

[13] Under this norm for the linear space of bounded stochastic processes, the norm of the residual stochastic process is the least upper bound such that the norm of the residual vector (under the usual Euclidean norm for finite-dimensional vectors) is within that bound almost surely at all dates.

term and terms that depend only on the exogenous disturbances); the form of these terms is irrelevant in obtaining a correct ranking of alternative policies.

Substituting (17) into (1), we obtain the approximate objective

$$V_{t_0} = E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \left[D_y \pi \cdot \tilde{y}_t + \frac{1}{2} \tilde{y}_t' D_{yy}^2 \pi \cdot \tilde{y}_t + \tilde{y}_t' D_{y\xi}^2 \pi \cdot \xi_t \right] + \text{t.i.p.} + \mathcal{O}\big(\epsilon^3\big). \tag{18}$$

This would be used as the quadratic objective in what we have called the "naive" LQ approximation. However, (18) is not the only valid quadratic approximation to (1). Taylor's theorem implies that it is the only quadratic function that correctly approximates (1) in the case of *arbitrary* (small enough) variations in the state variables, but there are others that will also correctly approximate (1) in the case of variations that are *consistent with the structural relations*. We can obtain an infinite number of alternative quadratic welfare measures by adding to (18) arbitrary multiples of quadratic (Taylor series) approximations to functions that must equal zero in order for the structural relations to be satisfied. Among these, we are able to find a welfare measure that is *purely quadratic, i.e.*, that contains no non-zero linear terms, as in Benigno and Woodford [6], so that a linear approximation to the equilibrium evolution of the endogenous variables under a given policy rule suffices to allow the welfare measure to be evaluated to second order. The key to this is using a second-order approximation to the structural relations to substitute purely quadratic terms for the linear terms $D_y \pi \cdot \tilde{y}_t$ in the sum (18), as in Sutherland [48].

A similar second-order Taylor series approximation can be written for each of the functions F^k. It follows that

$$\begin{aligned} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \bar{\lambda}' F(y_t, \xi_t; y_{t-1}) = \sum_{t=t_0}^{\infty} \beta^{t-t_0} \Big\{ & \bar{\lambda}' [D_y F + \beta D_{\check{y}} F] \cdot \tilde{y}_t \\ & + \frac{1}{2} \bar{\lambda}_k \big[\tilde{y}_t' D_{yy}^2 F^k \cdot \tilde{y}_t + 2 \tilde{y}_t' D_{y\xi}^2 F^k \cdot \xi_t + 2\beta \tilde{y}_t' D_{\check{y}\xi}^2 F^k \cdot \xi_{t+1} \\ & + \beta \tilde{y}_t' D_{\check{y}\check{y}}^2 F^k \cdot \tilde{y}_t + 2 \tilde{y}_t' D_{y\check{y}}^2 F^k \cdot \tilde{y}_{t-1} \big] \Big\} \\ & + \text{t.i.p.} + \mathcal{O}\big(\epsilon^3\big). \end{aligned} \tag{19}$$

Using a similar Taylor series approximation of each of the functions g^i, we correspondingly obtain

$$\begin{aligned} \sum_{t=t_0}^{\infty} \beta^{t-t_0-1} \bar{\varphi}' g(y_{t-1}, \xi_{t-1}; y_t) = \sum_{t=t_0}^{\infty} \beta^{t-t_0} \Big\{ & \bar{\varphi}' \big[D_y g + \beta^{-1} D_{\hat{y}} g \big] \cdot \tilde{y}_t \\ & + \frac{1}{2} \bar{\varphi}_i \big[\tilde{y}_t' D_{yy}^2 g^i \cdot \tilde{y}_t + 2 \tilde{y}_t' D_{y\xi}^2 g^i \cdot \xi_t \\ & + 2\beta^{-1} \tilde{y}_t' D_{\hat{y}\xi}^2 g^i \cdot \xi_{t-1} + \beta^{-1} \tilde{y}_t' D_{\hat{y}\hat{y}}^2 g^i \cdot \tilde{y}_t \\ & + 2\beta^{-1} \tilde{y}_t' D_{\hat{y}y}^2 g^i \cdot \tilde{y}_{t-1} \big] \Big\} \\ & + \text{t.i.p.} + \mathcal{O}\big(\epsilon^3\big). \end{aligned} \tag{20}$$

It then follows from constraints (2)–(4) that in the case of any admissible policy,[14]

$$\beta^{-1}\bar{\varphi}'\bar{g}_{t_0} = E_{t_0}\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left\{\Phi\cdot\tilde{y}_t + \frac{1}{2}\big[\tilde{y}_t'H\cdot\tilde{y}_t + 2\tilde{y}_t'R\tilde{y}_{t-1} + 2\tilde{y}_t'Z(L)\xi_{t+1}\big]\right\}$$
$$+\,\text{t.i.p.} + \mathcal{O}\big(\epsilon^3\big), \tag{21}$$

where

$$\Phi \equiv \bar{\lambda}'[D_yF + \beta D_{\check{y}}F] + \bar{\varphi}'\big[D_yg + \beta^{-1}D_{\hat{y}}g\big],$$
$$H \equiv \bar{\lambda}_k\big[D^2_{yy}F^k + \beta D^2_{\check{y}\check{y}}F^k\big] + \bar{\varphi}_i\big[D^2_{yy}g^i + \beta^{-1}D^2_{\hat{y}\hat{y}}g^i\big],$$
$$R \equiv \bar{\lambda}_kD^2_{y\check{y}}F^k + \bar{\varphi}_i\beta^{-1}D^2_{\hat{y}y}g^i,$$
$$Z(L) \equiv \beta\bar{\lambda}_kD^2_{\check{y}\xi}F^k + \big(\bar{\lambda}_kD^2_{y\xi}F^k + \bar{\varphi}_iD^2_{y\xi}g^i\big)\cdot L + \beta^{-1}\bar{\varphi}_iD^2_{\hat{y}\xi}g^i\cdot L^2.$$

Using (12), we furthermore observe that[15]

$$\Phi = -D_y\pi.$$

With this substitution in (21), we obtain an expression that can be solved for

$$E_{t_0}\sum_{t=t_0}^{\infty}\beta^{t-t_0}D_y\pi\cdot\tilde{y}_t,$$

which can in turn be used to substitute for the linear terms in (18). We thus obtain an alternative quadratic approximation to (1),[16]

$$V_{t_0} = \frac{1}{2}E_{t_0}\sum_{t=t_0}^{\infty}\beta^{t-t_0}\big[\tilde{y}_t'Q\cdot\tilde{y}_t + 2\tilde{y}_t'R\tilde{y}_{t-1} + 2\tilde{y}_t'B(L)\xi_{t+1}\big] + \text{t.i.p.} + \mathcal{O}\big(\epsilon^3\big), \tag{22}$$

where now

$$Q \equiv D^2_{yy}\pi + H,$$
$$B(L) \equiv Z(L) + D^2_{y\xi}\pi\cdot L. \tag{23}$$

Since (22) involves no linear terms, it can be evaluated (up to a residual of order $\mathcal{O}(\epsilon^3)$) using only a linear approximation to the evolution of $\tilde{y}_t$ under a given policy rule.

It follows that a correct LQ approximation to the original problem is given by the problem of choosing a state-contingent evolution $\{\tilde{y}_t\}$ for $t \geqslant t_0$ to maximize the objective

$$V^Q_{t_0}(\tilde{y};\xi) \equiv \frac{1}{2}E_{t_0}\sum_{t=t_0}^{\infty}\beta^{t-t_0}\big[\tilde{y}_t'A(L)\tilde{y}_t + 2\tilde{y}_t'B(L)\xi_{t+1}\big] \tag{24}$$

[14] Note that we here include (4) among the constraints that a policy must satisfy. We shall call any evolution that satisfies (2)–(3) a "feasible" policy. Under this weaker assumption, the left-hand side of (21) must instead be replaced by $\beta^{-1}\bar{\varphi}'g(y_{t_0-1},\xi_{t_0-1};y_{t_0})$.

[15] This is the point at which our calculations rely on the assumption that the steady state around which we compute our local approximations is optimal.

[16] Here we include $\bar{g}_{t_0}$ among the "terms independent of policy." If we consider also policies that are not necessarily consistent with the initial pre-commitment, the left-hand side of (22) should instead be written as $V_{t_0} + \beta^{-1}\bar{\varphi}'g(y_{t_0-1},\xi_{t_0-1};y_{t_0})$. This generalization of (22) is used in the derivation of Eq. (67) below.

subject to the constraints that

$$C(L)\tilde{y}_t = f_t, \tag{25}$$

$$E_t D(L)\tilde{y}_{t+1} = h_t \tag{26}$$

for all $t \geqslant t_0$, and the additional initial constraint that

$$D(L)\tilde{y}_{t_0} = \tilde{h}_{t_0}, \tag{27}$$

where now

$$A(L) \equiv Q + 2R \cdot L, \tag{28}$$

$$C(L) \equiv D_y F + D_{\check{y}} F \cdot L, \tag{29}$$

$$f_t \equiv -D_\xi F \cdot \xi_t,$$

$$D(L) \equiv D_{\hat{y}} g + D_y g \cdot L, \tag{30}$$

$$h_t \equiv -D_\xi g \cdot \xi_t, \tag{31}$$

$$\tilde{h}_{t_0} \equiv h_{t_0-1} + \bar{g}_{t_0}.$$

1.3. *Equivalence to linearization of the exact FOCs*

In the case that the objective (24) is concave,[17] the first-order conditions associated with the LQ problem just defined characterize the solution to that problem. We can show that these linear equations also correspond to a local linear approximation to the first-order conditions associated with the exact problem, *i.e.*, the modified Ramsey policy problem defined in Section 1.1.

Let the system of first-order conditions (11) be linearized around the optimal steady state, yielding a system of linear expectational difference equations of the form

$$E_t\big[J(L)\tilde{y}_{t+1}\big] + E_t\big[K(L)\xi_{t+1}\big] + E_t\big[M(L)\tilde{\lambda}_{t+1}\big] + N(L)\tilde{\varphi}_t = 0 \tag{32}$$

for each $t \geqslant t_0$, where

$$\tilde{\lambda}_t \equiv \lambda_t - \bar{\lambda}, \qquad \tilde{\varphi}_t \equiv \varphi_t - \bar{\varphi},$$

and where $J(L)$ and $K(L)$ are each matrix lag polynomials of second order, and $M(L)$ and $N(L)$ are matrix lag polynomials of first order. A local linear characterization of the solution to the exact problem[18] can then be obtained by solving the system of linear equations consisting of (32) together with (25)–(27). These form a linear system to be solved for the joint evolution of the processes $\{\tilde{y}_t, \tilde{\lambda}_t, \tilde{\varphi}_t\}$ given the exogenous disturbance processes $\{\xi_t\}$ and the initial conditions $\tilde{y}_{t_0-1}$ and the initial pre-commitment $\bar{g}_{t_0}$ (or $\hat{h}_{t_0}$).

The following result establishes the connection between this method of local approximation and the LQ approach proposed above.

[17] The algebraic conditions under which this is so are discussed in the next section.

[18] Here we assume that a solution of the system consisting of the exact FOCs and the structural equations corresponds to the optimum; of course, if this is not true, local methods of the kind used in this paper cannot be used to provide even an approximate characterization of optimal policy.

Proposition 1. *In Eq. (32), the matrix polynomials are given by*

$$J(L) \equiv \frac{1}{2}\left[A(L) + A'\left(\beta L^{-1}\right)\right]L, \qquad K(L) \equiv B(L),$$
$$M(L) \equiv C'\left(\beta L^{-1}\right)L, \qquad N(L) \equiv \beta^{-1} D'\left(\beta L^{-1}\right)L.$$

It follows that the linear system (32) is exactly the set of FOCs for the LQ problem of maximizing (24) subject to constraints (25)–(27).

The proof of this result is given in Appendix A. As discussed there, the reason for this identity between the two systems of linear equations is that a local quadratic approximation to the Lagrangian for the exact policy problem is precisely the Lagrangian for the LQ problem[19]; indeed, an alternative derivation of the correct quadratic objective for the LQ problem would proceed precisely from a second-order Taylor expansion of the Lagrangian for the exact problem.[20] It follows from this result that the solution to the LQ problem represents a local linear approximation to optimal policy from a timeless perspective.

1.4. Qualifications

While the conditions under which a valid LQ approximation is possible are fairly general, several qualifications to our results are in order. First of all, the LQ approximation, when valid, is purely *local* in character; it can only provide an approximate characterization of optimal policy to the extent that disturbances are sufficiently small. Whether the disturbances are small enough for this to be a useful approximation will depend upon the application; and a judgment about how accurate the approximation is likely to be is not possible on the basis of the coefficients of the LQ approximate problem alone. And like all perturbation approaches, it depends on sufficient differentiability of the problem[21]; it cannot be applied, for example, to problems in which there are inequality constraints that sometimes bind but at other times do not. Moreover, the LQ approximation provides at best a linear approximation to optimal policy. More general perturbation methods can instead be used to compute approximations of any desired order, assuming sufficient differentiability of the objective and constraints, as stressed by Judd [28].

Second, a correct LQ approximation yields a correct linear approximation to the optimal policy in the case that linearization of the exact first-order conditions yields a system of linear equations that can be solved to obtain a linear approximation to optimal policy. If the regularity condition required for the linearized system to have a determinate solution fails, the implicit function theorem cannot be applied to obtain a linear approximation in this way, and the LQ approach similarly fails to provide a correct linear approximation to optimal policy. This is a problem that can certainly arise in cases of economic interest, such as the portfolio problem

[19] It is worth noting that this equivalence of the two quadratic Lagrangians holds in the case of all feasible policies, whether or not the policy is consistent with the initial pre-commitment (4). This is important for our discussion of the welfare evaluation of suboptimal policies using the Lagrangian for the LQ problem in Section 3.

[20] This is the approach to the derivation of an approximate LQ problem used by Levine et al. [39]. It is essentially a discrete-time version of the approach used by Magill [40] in the context of a continuous-time stochastic growth model with purely backward-looking constraints.

[21] Of course, this may depend on a choice of variables. Kim et al. [37] provide an example in which the first-order conditions are not differentiable, if one of the state variables is the square root of a measure of price dispersion. If instead the state variable is taken to be the measure of price dispersion, the conditions are differentiable and the LQ method is applicable, as shown in Woodford [52].

treated by Judd and Guu [30], and singular perturbation methods can still be employed in such cases, as Judd and Guu show. But it is a problem the existence of which can be diagnosed within the LQ analysis itself: for when the regularity fails, the first-order conditions of the LQ problem fail to determine a unique solution. And an identical caveat applies to the method of linearization of the exact first-order conditions.

Third, a correct LQ approximation yields a correct linear approximation to the optimal policy only in the case that the perturbed (local) solution to the exact first-order conditions characterized by the implicit function theorem is in fact an optimum. It cannot be taken as obvious that the first-order conditions suffice for optimality, since in applications of interest, the structural relations (2)–(3) often define a non-convex set. The question of convexity can be addressed at least locally by evaluating the relevant second-order conditions, as discussed further in the next section. But of course, verification of the second-order conditions for optimality still only guarantees that the solution to the LQ problem approximates a *local* welfare optimum. The question of global optimality of the solution cannot be treated using purely local methods, and is often quite difficult in dynamic stochastic models.

1.5. Comparison with Ramsey policy

Our focus on the problem of optimal policy from a "timeless perspective" deserves further comment. It might seem more natural to be interested in the unconstrained (Ramsey) problem, or what we have above called t_0-optimal policy. It is worth noting that the unconstrained Ramsey policy problem can be given a sequential formulation (as for example in Benigno and Woodford, [6,7]), in which the *continuation* problem, in any period after the initial period t_0, is a problem with initial pre-commitments of the form (4). Thus in any case in which Ramsey policy implies asymptotic convergence to a deterministic steady state, in the absence of exogenous disturbances, and to bounded fluctuations around such a steady state in the case of small enough disturbances, Ramsey policy will eventually coincide with optimal policy from a timeless perspective, and can be locally approximated using the method expounded above. Hence even if one's interest is in the Ramsey policy problem, the LQ analysis presented here can be useful as a description of the asymptotic character of optimal policy, after an initial transition period.

In principle, one might compute a local approximation to Ramsey policy during the transition period as well. But the t_0-optimal policy generally does not imply constant values of the endogenous variables, even when there are no random disturbances and the functions π, F and g are all time-invariant, as assumed above; hence a correct local approximation to Ramsey policy in the case of small disturbances would involve derivatives evaluated along this non-constant path, so that the coefficients of the linear approximation would generally be time-varying.[22] It is true that in the literature on Ramsey policy, one sometimes sees approximate characterizations of optimal policy computed by log-linearizing around a steady state that Ramsey policy approaches asymptotically in the absence of random disturbances. But in such cases, there is no guarantee that the approximate characterization will be accurate even in the case of arbitrarily small dis-

[22] In the special case in which the Lagrange multipliers $\bar{\varphi}$ are small — the case of "small steady-state distortions" discussed in Woodford [51, Chapter 6], [52] — the Ramsey policy will also be near the optimal steady state during the transition period, and a linear approximation with constant coefficients is possible. In this case, an LQ analysis can also be used to characterize unconstrained Ramsey policy; see, for example, the discussion of Figs. 7.1 and 7.2 in Woodford [51].

turbances, as the transition dynamics need not be sufficiently near the steady state for the local approximation to be accurate.[23]

Nor is it obvious that unconstrained Ramsey policy should be the monetary policy design problem of greatest interest. No policy authority should ever choose a policy commitment at some date t_0 with the expectation that it will simply be enforced forever after.

Policy rules, even if adopted, will surely be reconsidered from time to time, if only because of changes in the structure of the economy and in economists' understanding of that structure; so a practical theory of policy design should explain how a rule should be chosen *on the occasion of each such reconsideration*. Choice of the t_j-optimal policy at each date t_j at which the policy rule is reconsidered is the not the best approach to such a problem, for the same reason that discretionary policy is not the best approach to the conduct of policy in general. Choice of a policy that is optimal from a timeless perspective each time that policy is reconsidered instead has the appealing feature that, even if policy were to be continually reconsidered without any change in one's model of the economy, the criterion would allow one to choose to continue one's previously chosen policy commitment.[24]

2. Characterizing optimal policy

We now study necessary and sufficient conditions for a policy to solve the LQ problem of maximizing (24) subject to constraints (25)–(27). Let $\mathcal{H}$ be the Hilbert space of (real-valued) stochastic processes $\{\tilde{y}_t\}$ such that

$$E_{t_0}\sum_{t=t_0}^{\infty}\beta^{t-t_0}\tilde{y}_t'\tilde{y}_t < \infty. \tag{33}$$

We are interested in solutions to the LQ problem that satisfy the bound (33) because it guarantees that the objective V^Q is well defined (and is generically required for it to be so). Of course, our LQ approximation to the original problem is only guaranteed to be accurate in the case that $\tilde{y}_t$ is always sufficiently small; hence a solution to the LQ problem in which $\tilde{y}_t$ grows without bound, but at a slow enough rate for (33) to be satisfied, need not correspond (even approximately) to any optimum (or local optimum) of the exact problem. In this section, however, we take the LQ problem at face value, and discuss the conditions under which it has a solution, despite the fact that we should in general only be interested in bounded solutions.

2.1. A Lagrangian approach

The Lagrangian for the LQ problem is given by

$$\mathcal{L}_{t_0}^{Q} = \frac{1}{2}\Bigg\{E_{t_0}\sum_{t=t_0}^{\infty}\beta^{t-t_0}\big[\tilde{y}_t'A(L)\tilde{y}_t + 2\tilde{y}_t'B(L)\xi_{t+1} + 2\tilde{\lambda}_t'C(L)\tilde{y}_t$$

[23] Levine et al. [39] assert that the local LQ approximation method that is possible in the case of backward-looking constraints applies equally in the case of forward-looking constraints, as the differences in the latter case "only affect the boundary conditions and not the steady state of the optimum, which is all we require for LQ approximation" [39, p. 3318]. This ignores the fact that the initial conditions associated with unconstrained Ramsey policy are necessarily not near the steady state, in the case of a steady state with "large distortions" in the sense of Benigno and Woodford [7].

[24] See Giannoni and Woodford [24] or Woodford [52] for further discussion.

$$+ 2\beta^{-1}\tilde{\varphi}'_{t-1}D(L)\tilde{y}_t\big]\Bigg\}. \tag{34}$$

Differentiation of the Lagrangian yields a system of linear first-order conditions, of the form (32) with the matrices of coefficients stated in Proposition 1. These conditions, together with (25)–(27), form a linear system to be solved for the joint evolution of the processes $\{\tilde{y}_t, \tilde{\lambda}_t, \tilde{\varphi}_t\}$ given the exogenous disturbance processes $\{\xi_t\}$ and the initial conditions $\tilde{y}_{t_0-1}$ and the initial pre-commitment $\bar{g}_{t_0}$ (or $\hat{h}_{t_0}$).

These FOCs are easily shown to be *necessary* for optimality, but they are not generally *sufficient* for optimality as well; one must also verify that second-order conditions for optimality are satisfied. We now consider these additional conditions.

Let us consider the subspace $\mathcal{H}_1 \subset \mathcal{H}$ of processes $\hat{y} \in \mathcal{H}$ that satisfy the additional constraints

$$C(L)\hat{y}_t = 0, \tag{35}$$

$$E_t D(L)\hat{y}_{t+1} = 0 \tag{36}$$

for each date $t \geqslant t_0$, along with the initial commitments

$$D(L)\hat{y}_{t_0} = 0, \tag{37}$$

where we define $\hat{y}_{t_0-1} \equiv 0$ in writing (35) for period $t = t_0$ and in writing (37). This subspace is of interest because if a process $\tilde{y} \in \mathcal{H}$ satisfies constraints (25)–(27), another process $y \in \mathcal{H}$ with $y_{t_0-1} = \tilde{y}_{t_0-1}$ satisfies those constraints as well if and only if $y - \tilde{y} \in \mathcal{H}_1$. We may now state our next main result.

Proposition 2. *For $\{\tilde{y}_t\} \in \mathcal{H}$ to maximize the quadratic form* (24)*, subject to the constraints* (25)–(27) *given initial conditions $\tilde{y}_{t_0-1}$ and $\bar{g}_{t_0}$, it is necessary and sufficient that* (i) *there exist Lagrange multiplier processes*[25] $\tilde{\varphi}, \tilde{\lambda} \in \mathcal{H}$ *such that the processes $\{\tilde{y}_t, \tilde{\varphi}_t, \tilde{\lambda}_t\}$ satisfy the FOCs* (32) *for each $t \geqslant t_0$; and* (ii)

$$V^Q(\hat{y}) \equiv V^Q_{t_0}(\hat{y}; 0) = \frac{1}{2}E_{t_0}\sum_{t=t_0}^{\infty}\beta^{t-t_0}\big[\hat{y}'_t A(L)\hat{y}_t\big] \leqslant 0 \tag{38}$$

for all processes $\hat{y} \in \mathcal{H}_1$, where in evaluating (38) *we define $\hat{y}_{t_0-1} \equiv 0$. A process $\{\tilde{y}_t\}$ with these properties is furthermore* uniquely *optimal if and only if*

$$V^Q(\hat{y}) < 0 \tag{39}$$

for all processes $\hat{y} \in \mathcal{H}_1$ that are non-zero almost surely.

The proof is given in Appendix A. The case in which the stronger condition (39) holds — *i.e.*, the quadratic form $V^Q(\hat{y})$ is negative definite on the subspace $\mathcal{H}_1$ — is the one of primary interest to us, since it is in this case that we know that the process $\{\tilde{y}_t\}$ represents at least a local welfare maximum in the exact problem. In this case we can also show that pure randomization of policy reduces the welfare objective (24), and hence is locally welfare-reducing in the exact problem as well, as is discussed further in Benigno and Woodford [7].

[25] Note that $\tilde{\varphi}_t$ is also assumed to be defined for $t = t_0 - 1$.

2.2. A dynamic programming approach

We can furthermore establish a useful characterization of the algebraic conditions under which the second-order conditions (39) are satisfied. These are most easily developed by considering the recursive formulation of our optimal policy problem presented in Section 1.1.[26] Let us suppose that the exogenous state vector ξ_t evolves according to a linear law of motion

$$\xi_{t+1} = \Gamma \xi_t + \epsilon_{t+1}, \tag{40}$$

where Γ is a matrix, all of the eigenvalues of which have modulus less than $\beta^{-1/2}$, and $\{\epsilon_t\}$ is an i.i.d. vector-valued random sequence, drawn each period from a distribution with mean zero and a variance-covariance matrix Σ.[27] In this case, our LQ approximate policy problem has a recursive formulation, in which the continuation problem from any period t forward depends on the extended state vector

$$\mathbf{z}_t \equiv \begin{bmatrix} \tilde{y}_{t-1} \\ \tilde{h}_t \\ \xi_t \\ \xi_{t-1} \end{bmatrix}. \tag{41}$$

Let $\bar{V}^Q(\mathbf{z}_t)$ denote the maximum attainable value of the continuation objective V_t^Q, if the process $\{\tilde{y}_\tau\}$ from date t onward is chosen to satisfy constraints (25)–(26) for all $\tau \geqslant t$, an initial pre-commitment of the form

$$D(L)\tilde{y}_t = \tilde{h}_t, \tag{42}$$

and the bound (33). As usual in an LQ problem of this form, it can be shown that the value function is a quadratic function of the extended state vector,

$$\bar{V}^Q(\mathbf{z}_t) = \frac{1}{2}\mathbf{z}_t' P \mathbf{z}_t, \tag{43}$$

where P is a symmetric matrix to be determined. In characterizing the solution to the problem, it is useful to introduce notation for partitions of the matrix P. Let P_{ij} (for $i, j = 1, 2, 3, 4$) be the 16 blocks obtained when P is partitioned in both directions conformably with the partition of $\mathbf{z}_t$ in (41), and let

$$\mathbf{P}_i \equiv [P_{i1}\ P_{i2}\ P_{i3}\ P_{i4}]$$

(for $i = 1, 2, 3, 4$) be the four blocks obtained when P is partitioned only vertically.

In the recursive formulation of the approximate LQ problem, in each period t, $\tilde{y}_t$ is chosen, and a pre-commitment $\tilde{h}_{t+1}(\xi_{t+1})$ is chosen for each possible state in the period $t+1$ continuation, so as to maximize

$$\frac{1}{2}\tilde{y}_t' A(L)\tilde{y}_t + E_t\big[\tilde{y}_t' B(L)\xi_{t+1}\big] + \beta E_t \bar{V}^Q(\mathbf{z}_{t+1}), \tag{44}$$

[26] This section has been improved by the suggestions of Paul Levine and Joe Pearlman.

[27] These assumptions ensure that the process $\{\xi_t\}$ satisfies a bound of the form (33). If we further wish to ensure that the disturbances are bounded, so that our local approximations can be expected to be accurate in the event of small enough disturbances, we may assume further that all eigenvalues of Γ have a modulus less than 1, and that ϵ_{t+1} is drawn from a distribution with bounded support.

subject to the constraints that $\tilde{y}_t$ satisfy (25) and (42), and that the choices of $\{\tilde{h}_{t+1}(\xi_{t+1})\}$ satisfy

$$E_t \tilde{h}_{t+1} = h_t. \tag{45}$$

To simplify the discussion, we further assume that

$$\operatorname{rank}\begin{bmatrix} C_0 \\ D_0 \end{bmatrix} = n_F + n_g, \tag{46}$$

where here and below we write lag polynomials in the form $X(L) = \sum_j X_j L^j$. This condition implies that the constraints (25) and (42) include neither any redundant constraints nor any constraints that are inconsistent in the case of a generic state $\mathbf{z}_t$.

It is then possible to show that the matrix P is given by the solution to a system of Riccati-type equations. In particular, we show in Appendix A that the block P_{11} can be determined as the solution to the system of equations[28]

$$P_{11} = -G_1' M(P_{11})^{-1} G_1, \tag{47}$$

where for any matrix P_{11}, the matrix M is defined as

$$M \equiv \begin{bmatrix} A_0 + \beta P_{11} & C_0' & D_0' \\ C_0 & 0 & 0 \\ D_0 & 0 & 0 \end{bmatrix}, \qquad y_t^\dagger \equiv \begin{bmatrix} \tilde{y}_t \\ \tilde{\lambda}_t \\ \tilde{\psi}_t \end{bmatrix}. \tag{48}$$

The block P_{22} is correspondingly given by

$$P_{22} = -G_2' M^{-1} G_2. \tag{49}$$

In these equations we use the notation

$$G_1 \equiv \begin{bmatrix} (1/2)A_1 \\ C_1 \\ D_1 \end{bmatrix}, \qquad G_2 \equiv \begin{bmatrix} 0 \\ 0 \\ -I \end{bmatrix}. \tag{50}$$

Given the solution for P, it is straightforward to derive FOCs for the optimization problem (44). As shown in Appendix A, these imply dynamics for the state vector of the form

$$\mathbf{z}_{t+1} = \Phi \mathbf{z}_t + \Psi \epsilon_{t+1}, \tag{51}$$

for certain matrices Φ and Ψ. Here we note that if we partition Φ in the same way as P, the block Φ_{11} is given by

$$\Phi_{11} \equiv [-I\ 0\ 0] M^{-1} G_1. \tag{52}$$

Because the single-period problem (44) is finite-dimensional, it is also straightforward to characterize the second-order conditions for an optimum. In fact, the second-order conditions for the single-period problem are necessary and sufficient for strict concavity of the infinite-horizon optimal policy problem, as established by the following result.

[28] This is a system of n_y^2 equations to solve for the n_y^2 elements of P_{11}. Actually, because P_{11} is symmetric, and the system (47) has the same symmetry, we need only solve a system of $n(n+1)/2$ equations for $n(n+1)/2$ independent quantities.

Proposition 3. *Suppose that the exogenous disturbances have a law of motion of the form (40), where Γ is a matrix the eigenvalues of which all have modulus less than $\beta^{-1/2}$, and that the constraints satisfy the rank condition (46), where $n_F + n_g < n_y$. Then the LQ policy problem has a determinate solution, given by (51), if and only if* (i) *there exists a solution P_{11} to Eq. (47) such that for each of the minors of the matrix M defined in (48),* $\det M_r$ *has the same sign as $(-1)^r$, for each $n_F + n_g + 1 \leqslant r \leqslant n_y$;* (ii) *the eigenvalues of the matrix Φ_{11} defined in (52) all have modulus less than $\beta^{-1/2}$; and* (iii) *the matrix P_{22} defined in (49) is negative definite, i.e., is such that its rth principle minor has the same sign as $(-1)^r$, for each $1 \leqslant r \leqslant n_g$.*

The proof of this proposition is also given in Appendix A. Note that the conditions stated in the proposition are necessary and sufficient both for the existence of a determinate solution to the first-order conditions, and for the quadratic form $V^Q(\psi)$ to satisfy the strict concavity condition (39). In the case that either condition (i) or (iii) is violated, there may exist a determinate solution to the first-order conditions, but it will not represent an optimum, owing to violation of the second-order conditions.

The fact that condition (iii) is needed in addition to conditions (i)–(ii) in order to ensure that we have a concave problem indicates an important respect in which the theory of LQ optimization with forward-looking constraints is not a trivial generalization of the standard theory for backward-looking problems, since conditions (i)–(ii) are sufficient in a backward-looking problem of the kind treated by Magill [40].[29] It also shows that the second-order conditions for a stochastic problem are more complex than they would be in the case of a deterministic policy problem (again, unlike what is true of purely backward-looking LQ problems). For in a deterministic version of our problem with forward-looking constraints, conditions (i)–(ii) would also be sufficient for concavity, and thus for the solution to the first-order conditions to represent an optimum.

In a deterministic version of the problem — where we not only assume that $\xi_t = 0$ each period, but we restrict our attention to policies under which the evolution of the variables $\{\tilde{y}_t\}$ is purely deterministic — the constraints on possible equilibria are the purely backward-looking constraints (25) and

$$D(L)\tilde{y}_t = \tilde{h}_t \tag{53}$$

for each $t \geqslant t_0$, where we specify $\tilde{h}_t = h_{t-1} = 0$ for all $t \geqslant t_0 + 1$. This is a purely backward-looking problem, so that the standard second-order conditions apply. And it should be obvious that, as there is no longer a choice of $\tilde{h}_{t+1}(\xi_{t+1})$ to be made each period, our argument above for the necessity of condition (iii) would not apply.

But conditions (i)–(ii) are not generally a sufficient condition to guarantee that (39) is satisfied, in the presence of forward-looking constraints (26), if policy randomization is allowed.[30]

[29] See Levine et al. [39] for a derivation of the second-order conditions for a backward-looking, deterministic LQ problem, using what is essentially a discrete-time version of the approach of Magill. In some cases, conditions (i)–(ii) are both necessary and sufficient for concavity, even in the presence of forward-looking constraints. The problem treated in Benigno and Woodford [7] is an example of this kind.

[30] Our remarks here apply even in the case that the “fundamental” disturbances $\{\xi_t\}$ are purely deterministic; what matters is whether policy may be contingent upon random events. As is discussed further in Benigno and Woodford [7], when the second-order conditions fail to hold, policy randomization can be welfare-improving, even when the random variations in policy are unrelated to any variation in fundamentals.

Because constraints (26) need hold only in expected value, random policy may be able to vary the paths of the endogenous variables (in some states of the world) in directions that would not be possible in the corresponding deterministic problem, and this makes the algebraic conditions required for (39) to hold more stringent. Specifically, the value function for the continuation problem must be a strictly concave function of the state-contingent pre-commitment $\tilde{h}_{t+1}$ made for the following period, or it is possible to randomize $\tilde{h}_{t+1}$ (requiring a corresponding randomization of subsequent policy) without changing the fact that constraint (26) is satisfied in period t. Hence condition (iii) is necessary in the stochastic case.[31] It can also easily be shown (see Appendix A) that condition (iii) is not implied in general by conditions (i)–(ii).

3. Welfare evaluation of alternative policy rules

Our approach can be used not only to derive a linear approximation to a fully optimal policy commitment, but also to compute approximate welfare comparisons between suboptimal rules, that will correctly rank these rules in the case that random disturbances are small enough. Because empirically realistic models are inevitably fairly complex, a fully optimal policy rule is likely to be too complex to represent a realistic policy proposal; hence comparisons among alternative simple rules are of considerable practical interest. Here we discuss how this can be done.

We do not propose to simply evaluate (a local approximation to) expected discounted utility V_{t_0} under a candidate policy rule, because the optimal policy locally characterized above (*i.e.*, optimal policy "from a timeless perspective") does not maximize this objective; hence ranking rules according to this criterion would lead to the embarrassing conclusion that there exist policies better than the optimal policy. Thus we wish to use a criterion that ranks rules according to how close they come to solving the recursive policy problem defined in Section 1.1, rather than how close they come to maximizing V_{t_0}.

Of course, if we restrict our attention to policies that necessarily satisfy the initial pre-commitment (4), there is no problem; our optimal rule will be the one that maximizes V_{t_0}, or (in the case of small enough shocks) the one that maximizes $V_{t_0}^Q$. But *simple* policy rules are unlikely to precisely satisfy (4); thus in order to be able to select the best rule from some simple class, we need an alternative criterion, one that is defined for *all* policies that are close enough to being optimal, in a sense that is to be defined. At the same time, we wish it to be a criterion the maximization of which implies that one has solved the constrained optimization problem defined in Section 1.1.

---

[31] Levine et al. (2008) provide a different argument for a condition similar to our condition (iii) as a necessary condition for optimality in a model with a forward-looking constraint, which does not require a consideration of stochastic policy. They consider Ramsey-optimal policy rather than optimality from a timeless perspective; that is, they assume no initial pre-commitment (27). In this case, the deterministic optimal policy problem is like the one considered above, except that (53) need hold only in periods $t \geqslant t_0 + 1$; the optimal policy is then the same as in the backward-looking problem just discussed, except that instead of taking $\tilde{h}_{t_0}$ as given, one is free to choose $\tilde{h}_{t_0}$ so as to maximize (24). This latter problem has a solution only if the value function $\bar{V}_{t_0}^Q$ is bounded above, for a given vector $\tilde{y}_{t_0-1}$, and this is true in general only if it is a strictly concave function of $\tilde{h}_{t_0}$. The validity of this argument, however, depends on considering an exact LQ problem, rather than an LQ local approximation to a problem that may have different global behavior.

3.1. A Lagrangian approach

Our Lagrangian characterization of optimal policy suggests such a criterion. The timelessly optimal policy from date t_0 onward — that is, the policy that maximizes V_{t_0} subject to the initial constraint (4) in addition to the feasibility constraints (2)–(3) — is also the policy that maximizes the Lagrangian

$$V_{t_0}^{mod} \equiv V_{t_0} + \beta^{-1}\varphi'_{t_0-1} g(y_{t_0-1}, \xi_{t_0-1}; y_{t_0}), \tag{54}$$

where φ_{t_0-1} is the vector of Lagrange multipliers associated with the initial constraint (4). This is a function that coincides (up to a constant) with the objective V_{t_0} in the case of policies satisfying the constraint (4), but that is defined more generally, and that is maximized over the broader class of feasible policies by the timelessly optimal policy. Hence an appropriate criterion to use in ranking alternative policies is the value of $V_{t_0}^{mod}$ associated with each one. This criterion penalizes policies that fail to satisfy the initial pre-commitment (4), by exactly the amount by which a previously *anticipated* deviation of that kind would have reduced the expected utility of the representative household.

In the case of any policy that satisfies the feasibility constraints (2)–(3) for all $t \geqslant t_0$, we observe that

$$\begin{aligned} V_{t_0}^{mod} &= \bar{\mathcal{L}}_{t_0} + \beta^{-1}\tilde{\varphi}'_{t_0-1} g(y_{t_0-1}, \xi_{t_0-1}; y_{t_0}) \\ &= V_{t_0}^{Q} + \beta^{-1}\tilde{\varphi}'_{t_0-1} D_{\hat{y}} g \cdot \tilde{y}_{t_0} + \text{t.i.p.} + \mathcal{O}(\epsilon^3). \end{aligned}$$

This suggests that in the case of small enough shocks, the ranking of alternative policies in terms of $V_{t_0}^{mod}$ will correspond to the ranking in terms of the welfare measure

$$W_{t_0} \equiv V_{t_0}^{Q} + \beta^{-1}\tilde{\varphi}'_{t_0-1} D_{\hat{y}} g \cdot \tilde{y}_{t_0}. \tag{55}$$

Note that in this derivation we have assumed that $\tilde{y}_t = \mathcal{O}(\epsilon)$. We shall restrict attention to policy rules of this kind. Note that while this is an important restriction, it does not preclude consideration of extremely simple rules; and it is a property of the simple rules of greatest interest, *i.e.*, those that come closest to being optimal among rules of that degree of complexity.

In expression (54), and hence in (55), φ_{t_0-1} is the Lagrange multiplier associated with constraint (4) under the optimal policy. However, in order to evaluate W_{t_0} to second-order accuracy, it suffices to have a first-order approximation to this multiplier. Such an approximation is given by the multiplier $\tilde{\varphi}_{t_0-1}$ associated with the constraint (27) of the LQ problem. Thus we need only solve the LQ problem, as discussed in the previous section — obtaining a value for $\tilde{\varphi}_{t_0-1}$ along with our solution for the optimal evolution $\{y_t\}$ — in order to determine the value of W_{t_0}.

Moreover, we observe that in the characterization given in the previous section of the solution to the LQ problem, $\tilde{\varphi}_{t_0-1} = \mathcal{O}(\epsilon)$.[32] Thus a solution for the equilibrium evolution $\{\tilde{y}_t\}$ under a given policy that is accurate to first order suffices to evaluate the second term in (55) to second-order accuracy. Hence W_{t_0} inherits this property of $V_{t_0}^{Q}$, and it suffices to compute a linear approximation to the equilibrium dynamics $\{\tilde{y}_t\}$ under each candidate policy rule in order to evaluate W_{t_0} to second-order accuracy. We can therefore obtain an approximation solution for $\{\tilde{y}_t\}$ under a given policy by solving the linearized structural equations (25)–(26), together with

[32] This follows from the solution given in Appendix A for the Lagrange multiplier associated with the initial pre-commitment.

the policy rule, and use this solution in evaluating W_{t_0}. In this way welfare comparisons among alternative policies are possible, to second-order accuracy, using linear approximations to the model structural relations and a quadratic welfare objective.

Moreover, we can evaluate W_{t_0} to second-order accuracy using only a linear approximation to the policy rule. This has important computational advantages. For example, if we wish to find the optimal policy rule from among the family of simple rules of the form $i_t = \phi(y_t)$, where i_t is a policy instrument, and we are content to evaluate $V_{t_0}^{mod}$ to second-order accuracy, then it suffices to search over the family of linear policy rules[33]

$$\tilde{\imath}_t = f'\tilde{y}_t,$$

parameterized by the vector of coefficients f. There are no possible second-order (or larger) welfare gains resulting from nonlinearities in the policy rule.

In expression (55), the value of the multiplier $\tilde{\varphi}_{t_0-1}$ depends on the economy's initial state and on the value of the initial pre-commitment $\bar{g}_{t_0}$. However, we wish to be able to rank alternative rules for an economy in which no such commitment may exist prior to the adoption of the policy rule. We can avoid having to make reference to any historically given pre-commitment by assuming a self-consistent constraint of the form (5). We show in Appendix A[34] that

$$\begin{aligned}\tilde{h}_{t_0} &= \hat{h}(\xi_{t_0}, \xi_{t_0-1}) \\ &\equiv h_{t_0-1} - P_{22}^{-1} P_{23}(\xi_{t_0} - \Gamma \xi_{t_0-1})\end{aligned} \tag{56}$$

is such a self-consistent specification, where the matrix P_{23} (another block of the matrix P introduced in (43)) is defined there.

If we assume this initial pre-commitment, we can then define the optimal dynamics from a timeless perspective as functions solely of the initial conditions $(\tilde{y}_{t_0-1}, \xi_{t_0-1})$ and the evolution of the exogenous states $\{\xi_t\}$ from period t_0 onward. We thus obtain a law of motion of the form

$$\bar{\mathbf{y}}_{t+1} = \bar{\Phi}\bar{\mathbf{y}}_t + \bar{\Psi}\epsilon_{t+1} \tag{57}$$

for the extended state vector

$$\bar{\mathbf{y}}_t \equiv \begin{bmatrix} \tilde{y}_t \\ \xi_t \end{bmatrix}. \tag{58}$$

This can be simulated to obtain the optimal dynamics from a timeless perspective, given initial conditions $\bar{\mathbf{y}}_{t_0-1}$.[35]

In Appendix A, we also establish the following property of the optimal dynamics subject to the pre-commitment (56).

---

[33] Here we restrict attention to rules that are consistent with the optimal steady state, so that the intercept term is zero when the rule is expressed in terms of deviations from steady-state values. Note that a rule without this property will result in lower welfare, in the case of any small enough disturbances.

[34] See the proof of Lemma 4.

[35] Note that it is possible to solve for the initial pre-commitment using only the values of ξ_{t_0}, $\tilde{y}_{t_0-1}$ and ξ_{t_0-1} (or equivalently, for the initial Lagrange multipliers φ_{t_0-1} using only the values of $\tilde{y}_{t_0-1}$ and ξ_{t_0-1}, as shown below.) Thus it is not necessary to simulate the optimal equilibrium dynamics over a lengthy "estimation period" prior to the date t_0 in order to compute the optimal dynamics from a timeless perspective, as proposed by Juillard and Pelgrin [31].

Lemma 4. *The Lagrange multiplier associated with the initial pre-commitment (56) is equal to*

$$\beta^{-1}\tilde{\varphi}_{t_0-1} = \varphi^*(\mathbf{y_{t_0-1}}) \equiv \tilde{\psi}(\tilde{y}_{t_0-1}, \xi_{t_0-1}), \tag{59}$$

where $\tilde{\psi}(\tilde{y}_{t-1}, \xi_{t-1})$ *is a function defined in* Appendix A.

Then we can write[36]

$$W_{t_0} = W(\tilde{y}; \xi_{t_0}, \mathbf{y_{t_0-1}}) \equiv V^Q_{t_0} + \varphi^*(\mathbf{y_{t_0-1}})' D_{\hat{y}} g \cdot \tilde{y}_{t_0}. \tag{60}$$

This gives us an expression for our welfare measure purely in terms of the history and subsequent evolution of the extended state vector.

3.2. A time-invariant criterion for ranking alternative rules

Let us suppose that we are interested in evaluating a policy rule r that implies an equilibrium evolution of the endogenous variables of the form[37]

$$y_t = \phi_r(\xi_t, \mathbf{y_{t-1}}).$$

This (together with the law of motion for the exogenous disturbances) then implies a law of motion for the complete extended state vector

$$\mathbf{y}_t = \psi_r(\xi_t, \mathbf{y_{t-1}}). \tag{61}$$

Using this law of motion, we can evaluate (60), obtaining

$$W_{t_0} = W_r(\xi_{t_0}, \mathbf{y_{t_0-1}}).$$

We can do this for any rule r of the assumed type, and hence we can define an optimization problem

$$\max_{r \in \mathcal{R}} W_r(\xi_{t_0}, \mathbf{y_{t_0-1}}) \tag{62}$$

in order to determine the optimal rule from among the members of some family of rules $\mathcal{R}$.

However, the solution to problem (62) may well depend on the initial conditions $\mathbf{y_{t_0-1}}$ and ξ_{t_0} for which W_{t_0} is evaluated. This implies an unappealing degree of arbitrariness of the choice that would be recommended from within some family of simple rules, as well as time inconsistency of the policy recommendation. The criterion that we find most appealing is to integrate over a distribution of possible initial conditions, rather than evaluating W_r at the economy's actual state at the time of the choice, or at any other single state (such as the optimal steady state).[38]

Suppose that in the case of the optimal policy rule r^*, the law of motion (61) implies that the evolution of the extended state vector $\{\mathbf{y_t}\}$ is *stationary*. In this case, there exists a well-defined invariant (or unconditional) probability distribution μ for the possible values of $\mathbf{y_t}$ under the

[36] In writing the function $W(\cdot)$, and others that follow, we suppress the argument ξ, as the evolution of the exogenous disturbances is the same in the case of each of the alternative policies under consideration.

[37] This assumption that y_t depends only on the state variables indicated is without loss of generality, as we can extend the vector $\mathbf{y_t}$ if necessary in order for this to be so.

[38] For example, the decision to evaluate W_r assuming initial conditions consistent with the steady state — when in fact the state of the economy will fluctuate on both sides of the steady-state position — favors rules r for which W_r is a less concave function of the initial condition.

optimal policy.[39] Then we can define the optimal policy rule within some class of simple rules $\mathcal{R}$ as the one that solves the problem

$$\max_{r\in\mathcal{R}} E_\mu\big[\bar{W}_r(\mathbf{y_t})\big], \tag{63}$$

where[40]

$$\bar{W}_r(\mathbf{y_t}) \equiv E_t W_r(\xi_{t+1}, \mathbf{y_t}). \tag{64}$$

Because of the linearity of our approximate characterization of optimal policy, the calculations required in order to evaluate $E_\mu[W_r]$ to second-order accuracy are straightforward; these are illustrated in Benigno and Woodford [7, Sec. 5].

The most important case in which the method just described cannot be applied is when some of the elements of $\{\mathbf{y_t}\}$ possess unit roots, though all elements are at least difference-stationary (and some of the non-stationary elements may be cointegrated). Note that it is possible for even the equilibrium under optimal policy to have this property, consistent with our assumption of the bound (33).[41] There is a question in such a case whether our local approximation to the problem should remain an accurate approximation, but this is not a problem in the case that random disturbances occur in only a *finite* number of periods, so LQ problems of this kind may be of practical interest.

Let us suppose that those elements which possess unit roots are pure random walks (*i.e.*, with zero drift).[42] We can in such a case decompose the extended state vector as

$$\mathbf{y_t} = \mathbf{y_t}^{tr} + \mathbf{y_t}^{cyc},$$

where

$$\mathbf{y_t}^{tr} \equiv \lim_{T\to\infty} E_t \mathbf{y_T}$$

is the Beveridge–Nelson [11] "trend" component, and the "cyclical" component $\mathbf{y_t}^{cyc}$ will still be a stationary process. Moreover, the evolution of the cyclical component as a function of the exogenous disturbances under the optimal policy will be independent of the assumed initial value of the trend component (though not of the initial value of the cyclical component). It follows that we can define an invariant distribution μ for the possible values of $\mathbf{y_t}^{cyc}$ under the optimal policy, that is independent of the assumed value for the trend component. Then for any assumed initial value for the trend component $\mathbf{y}^{tr}_{\mathbf{t_0-1}}$, we can define the optimal policy rule within the class $\mathcal{R}$ as the one that solves the problem

39 We discuss the computation of the relevant properties of this invariant measure in Appendix A.

40 Recall that we assume that the exogenous disturbance process $\{\xi_t\}$ is Markovian, and that ξ_t is included among the elements of $\mathbf{y_t}$. Hence $\mathbf{y_t}$ contains all relevant elements of the period t information set for the calculation of this conditional expectation.

41 Benigno and Woodford [6] provide an example of an optimal stabilization policy problem in which the LQ approximate problem has this property. In this example, the unit root is associated with the dynamics of the level of real public debt, which display a unit root under optimal policy for the same reason as in the classic analysis of optimal tax smoothing by Barro [3] and Sargent [44, Chapter XV].

42 We may suppose that any deterministic trend under optimal policy has been eliminated by local expansion around a deterministic solution with constant trend growth, so that there is zero trend in the state variables $\{\tilde{y}_t\}$ expressed as deviations from that deterministic solution.

$$\max_{r\in\mathcal{R}} \Omega_r\big(\mathbf{y}^{tr}_{\mathbf{t_0-1}}\big) \equiv E_\mu\big[\bar{W}_r(\mathbf{y}_{\mathbf{t_0-1}})\big], \tag{65}$$

a generalization of (63).[43]

It might seem in this case that our criterion is again dependent on initial conditions, just as with the criterion (62) proposed first. The following result shows that this is not the case.

Lemma 5. *Suppose that under optimal policy, the extended state vector* $\mathbf{y_t}$ *consists entirely of components that are either* (i) *stationary, or* (ii) *pure random walks. Suppose also that the class of policy rules* $\mathcal{R}$ *is such that each rule in the class implies convergence to the same long-run values of the state variables as under optimal policy, in the absence of stochastic disturbances, so that the initial value of the trend component* $\mathbf{y}^{tr}_{t_0-1}$ *is the same regardless of the rule* r *that is considered. Then for any rule* $r\in\mathcal{R}$*, the objective* $\Omega_r(\mathbf{y}^{tr}_{\mathbf{t_0-1}})$ *defined in* (65) *can be decomposed into two parts,*

$$\Omega_r\big(\mathbf{y}^{tr}_{\mathbf{t_0-1}}\big) = \Omega^1\big(\mathbf{y}^{tr}_{\mathbf{t_0-1}}\big) + \Omega^2_r, \tag{66}$$

where the first component is the same for all rules in this class, while the second component is independent of the initial condition $\mathbf{y}^{tr}_{\mathbf{t_0-1}}$.

Hence the criterion (65) establishes the same ranking of alternative rules, regardless of the initial condition. The proof of this result is given in Appendix A.

3.3. Comparison with alternative criteria

The criterion for ranking alternative simple policy rules proposed above differs from others sometimes used. Many authors prefer to evaluate alternative simple policy rules by computing the expected value of V_{t_0} (rather than $V^{mod}_{t_0}$) associated with each rule (*e.g.*, Schmitt-Grohé and Uribe [47]). As noted above, this alternative criterion is one under which the optimal rule from a timeless perspective can be dominated by other rules, a point stressed by Blake [13], Jensen and McCallum [26], and Dennis [17], among others. A consequence is that the use of such a criterion as the basis for a reconsideration of policy will not lead to continuation of an optimal policy commitment, even if there has been no change in the policymaker's objective or model of the economy, and even if the class of policies considered is flexible enough to include the continuation of the prior optimal plan.

Choice of a rule to maximize V_{t_0} will not even lead to a stable policy choice — renewal of one's commitment to the same rule, each time the issue is reconsidered, if one's model remains the same — if V_{t_0} is evaluated conditional on the economy's actual state at date t_0. Schmitt-Grohé and Uribe [47] avoid this by proposing that one should choose the rule that would be judged best in the case of initial conditions consistent with the optimal steady state, whether the economy's actual initial state is that one or not. But this choice is an arbitrary one, and in general a different simple rule would be favored if one were to arbitrarily a choose a different fictitious initial condition.

The alternative criterion is also one that cannot be evaluated to second-order accuracy using only a first-order solution for the equilibrium evolution under a given policy. For a general fea-

[43] In the case that all elements of $\mathbf{y_t}$ are stationary, $\mathbf{y}^{tr}_{\mathbf{t}}$ is simply a constant, and all variations in $\mathbf{y_t}$ correspond to variations in $\mathbf{y}^{cyc}_{\mathbf{t}}$. In this case, (65) is equivalent to the previous criterion (63).

sible policy — consistent with the optimal steady state, but not necessarily consistent with the initial pre-commitment (4) — we can show that[44]

$$V_{t_0} = V_{t_0}^{Q} - \beta^{-1}\bar{\varphi}' D_{\hat{y}} g \cdot \tilde{y}_{t_0} + \text{t.i.p.} + \mathcal{O}\bigl(\epsilon^3\bigr). \tag{67}$$

The first term on the right-hand side of this expression is purely quadratic (has zero linear terms), but this is not true of the second term, if the initial pre-commitment is binding under the optimal policy. Evaluation of the second term to second-order accuracy requires a second-order approximation to the evolution $\{y_t\}$ under the policy of interest; there is thus no alternative to the use of higher-order perturbation solution methods as illustrated by Schmitt-Grohé and Uribe [47], and nonlinear terms in the policy rule generally matter for welfare.

We therefore cannot agree with Levine et al. [39], who state, in the context of a discussion of the general possibility of LQ approximation, that "if … one adopts a *conditional* welfare loss measure, starting at the zero-inflation steady state, then the timeless perspective is not relevant" [39, p. 3319]. This is incorrect, unless "starting at the [optimal] steady state" means starting with initial pre-commitments that are consistent with that steady state,[45] and not simply evaluation conditional on a fictitious initial state. For even assuming initial values of the predetermined state variables consistent with the optimal steady state, the policy that maximizes the conditional welfare measure V_{t_0} will not in general be a continuation of the optimal steady state, nor need it even be near the steady state (at least initially). While a local quadratic approximation to the criterion V_{t_0} is possible, it would have to be computed around the unconstrained Ramsey policy (which would not correspond to the optimal steady state, even under the hypothesized initial conditions), and so would involve time-varying coefficients.

Authors such as Blake [13] and Jensen and McCallum [26] instead avoid time-inconsistency without adopting the timeless perspective, by proposing to rank alternative rules according to the unconditional expected value of V_{t_0}, that is, the expected value of V_{t_0} under a probability distribution for initial conditions corresponding to the ergodic distribution for the endogenous variables associated with the particular time-invariant policy that is to be evaluated.[46] Damjanovic et al. [16] show that one can use an LQ approximation (slightly different from the one derived here) to evaluate time-invariant policy rules under this criterion. Note, however, that unlike the approach proposed here, the probability distribution for initial conditions that is used is not independent of the policy rule that is considered. This criterion has the unappealing feature of giving a rule that leads to different long-run average values of an endogenous variable (*e.g.*, the capital stock) "credit" for a higher initial average value of the variable as well. It also cannot be applied to evaluate non-stationary policies, or even time-invariant policies that imply non-stationary dynamics of endogenous variables, such as the optimal policy in Benigno and Woodford [6].

Finally, some authors who seek to evaluate alternative simple rules from the timeless perspective have proposed alternative methods for computing the initial lagged Lagrange multipliers φ_{t_0-1} than the one described above. Juillard and Pelgrin [31] propose an approach in which one simulates the optimal equilibrium dynamics over a lengthy "estimation period" prior to the date

44 Here we use the more general form of (22) mentioned in footnote 25.

45 Because Levine et al. treat problems with forward-looking constraints as equivalent to problems with only backward-looking constraints, their understanding of initial conditions consistent with the optimal steady state may tacitly include such commitments. But if this is so, the proper criterion on which to rank alternative policies is $V_{t_0}^{mod}$ rather than V_{t_0}, which is precisely the criterion required by the timeless perspective.

46 This approach had previously been used by Rotemberg and Woodford [43] to rank alternative time-invariant policy rules, prior to the proposal of the timeless perspective by Woodford [49].

t_0 in order to compute the proper initial multipliers. This should in principle lead to results consistent with ours, but is an unnecessarily cumbersome approach, and introduces sampling error in the case of a simulation of only finite length. Dennis [17] proposes a recursive approach similar to ours, but which yields a non-unique solution for the initial Lagrange multipliers, since it involves the generalized inverse of a matrix for which the generalized inverse is not unique, and allows the initial Lagrange multipliers to be functions of redundant information.

4. Applications

The approach expounded here has already proven fruitful in a number of applications to problems of optimal monetary and fiscal policy. Benigno and Woodford [7] use this method to derive an LQ approximation to the problem of optimal monetary stabilization policy in a DSGE model with monopolistic competition, Calvo-style staggered price-setting, and a variety of exogenous disturbances to preferences, technology, and fiscal policy.[47] Unlike the simpler LQ method used by Rotemberg and Woodford [43] and Woodford [50], the present method is applicable even in the case of (possibly substantial) distortions even in the absence of shocks, owing to market power or distorting taxes. As in the simpler case considered in Woodford [50], the quadratic stabilization objective obtained is a sum of two terms per period, corresponding to an inflation stabilization and an output-gap stabilization objective respectively; but both the definition of the output gap and the relative weight on output-gap stabilization are now more complex.

Benigno and Woodford [8] extend the analysis to the case in which both wages and prices are sticky, obtaining a generalization of the utility-based loss function in which a third quadratic term appears, proportional to squared deviations of nominal wage inflation from zero. This shows that the analysis by Erceg et al. [20] of the tradeoff between stabilization of wage inflation and price inflation applies also to economies with distorted steady states. Montoro [42] extends the analysis to allow for real disturbances to the relative supply price of oil.

Benigno and Benigno [4] analyze policy coordination between two national monetary authorities which each seek to maximize the welfare of their own country's representative household, and show that it is possible to locally characterize each authority's aims by a quadratic stabilization objective. Engel [21] extends the analysis to a two-country model with local-currency pricing rather than the producer-currency pricing assumed by Benigno and Benigno [4]. De Paoli [18] similarly shows how the analysis of Benigno and Woodford [7] can be extended to a small open economy, requiring the addition of a terms-of-trade (or real-exchange-rate) stabilization objective to the quadratic loss function; De Paoli [19] extends the analysis to the case of a small open economy with incomplete international risk sharing.

Because the present method applies to economies with a distorted steady state, it allows the theory of tax smoothing to be integrated with the theory of monetary stabilization policy. Benigno and Woodford [6] extend the analysis of Benigno and Woodford [7] to the case of an economy with only distorting taxes, and show that the problem of choosing jointly optimal monetary and fiscal policies can also be treated within an LQ framework that nests standard analyses of tax smoothing (with flexible prices, so that real effects of monetary policy are ignored) and of monetary policy (with lump-sum taxes, so that fiscal effects of monetary policy can be ignored) as special cases. Berriel and Sinigaglia [12] extend the analysis to the case of an economy with

[47] See Giannoni and Woodford [24] for a demonstration that this policy problem can be cast in the general form assumed in this paper. Woodford [52, Sec. 2] provides a further exposition of the LQ analysis of this problem, and a comparison of the results of the LQ analysis to those obtained by linearization of the exact first-order conditions for optimal policy.

multiple sectors that differ in the degree of stickiness of prices, and Horvath [25] extends it to a model in which wages as well as prices are sticky. Benigno and De Paoli [5] use similar methods to analyze optimal fiscal policy in a small open economy, while Ferrero [22] analyzes optimal monetary and fiscal policy in a monetary union with separate national fiscal authorities.

All of the analyses just mentioned involve fairly simple DSGE models, in which it is possible to derive the coefficients of the LQ approximate policy problem by hand. In the case of larger (and more realistic) models, such calculations are likely to be tedious. Nonetheless, it is an advantage of our method that it is straightforward to apply it even to fairly complex models and fairly general specifications of disturbances. Altissimo et al. [1], Cúrdia [15], Levine et al. [39], and Coenen et al. [14] all provide examples of numerical analyses of optimal policy in more complex models using the LQ approximation method. We believe that it should similarly be practical to apply these methods to a wide variety of other models of interest to policy institutions.

Appendix A. Proofs and derivations

A.1. *Proposition 1*

Proposition 1. *In Eq. (32), the matrix polynomials are given by*

$$J(L) \equiv \frac{1}{2}\big[A(L) + A'\big(\beta L^{-1}\big)\big]L, \qquad K(L) \equiv B(L),$$
$$M(L) \equiv C'\big(\beta L^{-1}\big)L, \qquad N(L) \equiv \beta^{-1} D'\big(\beta L^{-1}\big)L.$$

It follows that the linear system (32) is exactly the set of FOCs for the LQ problem of maximizing (24) subject to constraints (25)–(27).

Proof. The identity of the matrix polynomials follows directly from differentiation of the nonlinear functions in the exact FOCs (11). These can then be observed to be the same coefficients that appear in the FOCs for the LQ problem. Differentiation of the Lagrangian (34) yields a system of linear first-order conditions

$$\frac{1}{2}E_t\big\{\big[A(L) + A'\big(\beta L^{-1}\big)\big]\tilde{y}_t\big\} + E_t\big[B(L)\xi_{t+1}\big] + E_t\big[C'\big(\beta L^{-1}\big)\tilde{\lambda}_t\big] + \beta^{-1} D'\big(\beta L^{-1}\big)\tilde{\varphi}_{t-1} = 0 \tag{A.1}$$

that must hold for each $t \geqslant t_0$ under the optimal policy. This is a system of the form (32), with the matrices of coefficients stated in the proposition.

An intuition for this result can be provided as follows. The FOCs (11) for the exact policy problem are obtained by differentiating the Lagrangian $\mathcal{L}_{t_0}$ defined in (10). The linearization (32) of the FOCs around the optimal steady state is in turn the set of linear equations that would be obtained by differentiating a quadratic approximation to $\mathcal{L}_{t_0}$ around that same steady state. Hence we are interested in computing such a local approximation, for the case in which $y_t - \bar{y}$, $\lambda_t - \bar{\lambda}$, and $\varphi_t - \bar{\varphi}$ are each of order $\mathcal{O}(\epsilon)$ for all t.

We may furthermore write the Lagrangian in the form

$$\mathcal{L}_{t_0} = \bar{\mathcal{L}}_{t_0} + \tilde{\mathcal{L}}_{t_0},$$

where

$$\bar{\mathcal{L}}_{t_0} = V_{t_0} + E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \big[\bar{\lambda}' F(y_t, \xi_t; y_{t-1}) + \beta^{-1} \bar{\varphi}' g(y_{t-1}, \xi_{t-1}; y_t)\big],$$

$$\tilde{\mathcal{L}}_{t_0} = E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \big[\tilde{\lambda}_t' F(y_t, \xi_t; y_{t-1}) + \beta^{-1} \tilde{\varphi}_{t-1}' g(y_{t-1}, \xi_{t-1}; y_t)\big].$$

We can then use Eqs. (18) and (21) to show that the local quadratic approximation to $\bar{\mathcal{L}}_{t_0}$ is given by

$$\bar{\mathcal{L}}_{t_0} = V_{t_0}^{Q} + \text{t.i.p.} + \mathcal{O}(\epsilon^3).$$

In addition, the fact that $\tilde{\lambda}_t, \tilde{\varphi}_t$ are both of order $\mathcal{O}(\epsilon)$ means that a local quadratic approximation to the other term is given by

$$\tilde{\mathcal{L}}_{t_0} = E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \big[\tilde{\lambda}_t' \tilde{F}(y_t, \xi_t; y_{t-1}) + \beta^{-1} \tilde{\varphi}_{t-1}' \tilde{g}(y_{t-1}, \xi_{t-1}; y_t)\big] + \mathcal{O}(\epsilon^3),$$

where $\tilde{F}$ and $\tilde{g}$ are local linear approximations to the functions F and g respectively.

Hence the local quadratic approximation to the complete Lagrangian is given by

$$\mathcal{L}_{t_0} = V_{t_0}^{Q} + E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \big[\tilde{\lambda}_t' \tilde{F}(y_t, \xi_t; y_{t-1}) + \beta^{-1} \tilde{\varphi}_{t-1}' \tilde{g}(y_{t-1}, \xi_{t-1}; y_t)\big] + \text{t.i.p.} + \mathcal{O}(\epsilon^3). \tag{A.2}$$

But this is identical (up to terms independent of policy) to the Lagrangian (34) for the LQ problem of maximizing $V_{t_0}^{Q}$ subject to the linearized constraints. Hence the first-order conditions obtained from this approximate Lagrangian (which coincide with the local linear approximation to the first-order conditions for the exact problem) are identical to the first-order conditions for the LQ problem, and their solutions are identical as well. □

A.2. Proposition 2

Recall that $\mathcal{H}$ is the Hilbert space of (real-valued) stochastic processes $\{\tilde{y}_t\}$ such that

$$E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \tilde{y}_t' \tilde{y}_t < \infty, \tag{A.3}$$

and $\mathcal{H}_1 \subset \mathcal{H}$ is the subspace of sequences $\hat{y} \in \mathcal{H}$ that satisfy the additional constraints

$$C(L)\hat{y}_t = 0, \tag{A.4}$$

$$E_t D(L)\hat{y}_{t+1} = 0 \tag{A.5}$$

for each date $t \geqslant t_0$, along with the initial commitments

$$D(L)\hat{y}_{t_0} = 0, \tag{A.6}$$

where we define $\hat{y}_{t_0-1} \equiv 0$ in writing (A.4) for period $t = t_0$ and in writing (A.6).

Proposition 2. *For $\{\tilde{y}_t\} \in \mathcal{H}$ to maximize the quadratic form (24), subject to the constraints (25)–(27) given initial conditions $\tilde{y}_{t_0-1}$ and $\bar{g}_{t_0}$, it is necessary and sufficient that* (i) *there exist Lagrange multiplier processes*[48] *$\tilde{\varphi}, \tilde{\lambda} \in \mathcal{H}$ such that the processes $\{\tilde{y}_t, \tilde{\varphi}_t, \tilde{\lambda}_t\}$ satisfy (A.1) for each $t \geqslant t_0$; and* (ii)

$$V^Q(\hat{y}) \equiv V_{t_0}^Q(\hat{y}; 0) = \frac{1}{2} E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \big[\hat{y}_t' A(L) \hat{y}_t\big] \leqslant 0 \tag{A.7}$$

for all processes $\hat{y} \in \mathcal{H}_1$, where in evaluating (A.7) we define $\hat{y}_{t_0-1} \equiv 0$. A process $\{\tilde{y}_t\}$ with these properties is furthermore uniquely *optimal if and only if*

$$V^Q(\hat{y}) < 0 \tag{A.8}$$

for all processes $\hat{y} \in \mathcal{H}_1$ that are non-zero almost surely.

Proof. We have already remarked on the necessity of the first-order conditions (i). To prove the necessity of the second-order condition (ii) as well, let $\{\tilde{y}_t\} \in \mathcal{H}$, and consider the perturbed process

$$y_t = \tilde{y}_t + \hat{y}_t \tag{A.9}$$

for all $t \geqslant t_0 - 1$, where $\{\hat{y}_t\}$ belongs to $\mathcal{H}_1$ and we define $\hat{y}_{t_0-1} \equiv 0$. This construction guarantees that if the process $\{\tilde{y}_t\}$ satisfies the constraints (25)–(27), so does the process $\{y_t\}$.

We note that

$$V_{t_0}^Q(y; \xi) = V_{t_0}^Q(\tilde{y}; \xi) + \frac{1}{2} E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \big[\hat{y}_t' A(L) \tilde{y}_t + \tilde{y}_t' A(L) \hat{y}_t + 2\hat{y}_t' B(L) \xi_{t+1}\big]$$
$$+ \frac{1}{2} E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \big[\hat{y}_t' A(L) \hat{y}_t\big].$$

The second term on the right-hand side is furthermore equal to

$$\frac{1}{2} E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \hat{y}_t' \cdot \big\{[A(L) + A'\big(\beta L^{-1}\big)]\tilde{y}_t + 2B(L)\xi_{t+1}\big\}$$
$$= -E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \hat{y}_t' \cdot \big\{C'\big(\beta L^{-1}\big)\tilde{\lambda}_t + \beta^{-1} D'\big(\beta L^{-1}\big)\tilde{\varphi}_{t-1}\big\}$$
$$= -E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \big\{\tilde{\lambda}_t' C(L)\hat{y}_t + \beta^{-1} \tilde{\varphi}_{t-1}' D(L)\hat{y}_t\big\},$$

where we use the first-order conditions (A.1) to establish the first equality, and conditions (35)–(37) to establish the final equality.

Thus for any feasible process $\tilde{y}$ and any perturbation (A.9) defined by a process $\hat{y}$ belonging to $\mathcal{H}_1$,

$$V_{t_0}^Q(y; \xi) = V_{t_0}^Q(\tilde{y}; \xi) + V^Q(\hat{y}). \tag{A.10}$$

[48] Note that $\tilde{\varphi}_t$ is also assumed to be defined for $t = t_0 - 1$.

It follows that if there were to exist any $\hat{y} \in \mathcal{H}_1$ for which $V^Q(\hat{y}) > 0$, the plan $\tilde{y}$ could not be optimal. But as this is true regardless of what plan $\tilde{y}$ may be, (A.7) is necessary for optimality. Furthermore, if there were to exist a non-zero $\hat{y}$ for which $V^Q(\hat{y}) = 0$, it would be possible to construct a perturbation y (not equal to $\tilde{y}$ almost surely at all dates) that would achieve an equally high level of welfare. Hence the stronger version of the second-order conditions (A.8) must hold for all $\hat{y}$ not equal to zero almost surely, in order for $\{\tilde{y}_t\}$ to be a unique optimum.

One easily sees from the same calculation that these conditions are also sufficient for an optimum. Let $\{\tilde{y}_t\}$ be a process consistent with the constraints of the LQ problem. Then any alternative process $\{y_t\}$ that is also consistent with those constraints can be written in the form (A.9), where $\hat{y}$ is some element of $\mathcal{H}_1$. If the first-order conditions (A.1) are satisfied by the process $\{\tilde{y}_t\}$, we can again establish (A.10). Condition (A.7) then implies that no alternative process is preferable to $\{\tilde{y}_t\}$, while (A.8) would imply that $\{\tilde{y}_t\}$ is superior to any alternative that is not equal to $\tilde{y}$ almost surely. $\square$

A.3. Dynamic programming formulation of the LQ problem

In the recursive formulation of the approximate LQ problem, in each period t, $\tilde{y}_t$ is chosen, and a pre-commitment $\tilde{h}_{t+1}(\xi_{t+1})$ is chosen for each possible state in the period $t+1$ continuation, so as to maximize (44) subject to the constraints that $\tilde{y}_t$ satisfy (25) and (42), and that the choices of $\{\tilde{h}_{t+1}(\xi_{t+1})\}$ satisfy (45).

The first-order conditions for the optimal choice of $\tilde{y}_t$ in this single-period problem are of the form

$$\left[A_0 + (1/2)A_1 L\right]\tilde{y}_t + E_t\left[B(L)\xi_{t+1}\right] + \beta \mathbf{P}_1 E_t \mathbf{z}_{t+1} + C_0'\tilde{\lambda}_t + D_0'\tilde{\psi}_t = 0, \tag{A.11}$$

where $\tilde{\lambda}_t$, $\tilde{\psi}_t$ are the Lagrange multipliers associated with constraints (25) and (42) respectively. Condition (A.11) together with the constraints (25) and (42) constitute a system of $n = n_y + n_F + n_g$ linear equations to solve for $\tilde{y}_t$, $\tilde{\lambda}_t$, and $\tilde{\psi}_t$ as functions of $\mathbf{z}_t$. This system can be written in the matrix form $My_t^\dagger = -G\mathbf{z}_t$, where the matrix M is defined by (48), and the first two column blocks of the matrix G are the matrices G_1, G_2 defined in (50).

This has a determinate solution if and only if M is non-singular. This is evidently a necessary condition for strict concavity of the policy problem, and we shall assume that it holds in the remainder of this discussion.[49] Given this assumption, the unique solution is

$$y_t^\dagger = -M^{-1}G\mathbf{z}_t. \tag{A.12}$$

The first-order conditions for the optimal choice of the pre-commitments $\{\tilde{h}_{t+1}(\xi_{t+1})\}$ are that

$$\beta \mathbf{P}_2 \mathbf{z}_{t+1} = -\tilde{\varphi}_t \tag{A.13}$$

in each possible state ξ_{t+1} that can succeed the given state in period t, where $\tilde{\varphi}_t$ is the Lagrange multiplier associated with constraint (45); note that the value of $\tilde{\varphi}_t$ depends only on the state in period t. The fact that the left-hand side of (A.13) must be the same in each state ξ_{t+1} implies that

[49] We are actually only interested in whether there exists a unique solution for $\tilde{y}_t$. However, condition (46) implies that there can be no vector $y^\dagger \neq 0$ such that $My^\dagger = 0$, unless it involves $\tilde{y} \neq 0$. Thus if M is singular, there are necessary multiple solutions for $\tilde{y}_t$ if there are any solutions at all, and not just multiple solutions for the Lagrange multipliers.

$$P_{22}[\tilde{h}_{t+1} - h_t] + P_{23}\epsilon_{t+1} = 0$$

in each state. This allows a determinate solution for $\tilde{h}_{t+1}$ if and only if P_{22} is non-singular; this too is evidently a necessary condition for concavity, and is assumed from here on.[50] Under this assumption, (A.13) together with (42) implies that

$$\tilde{h}_{t+1} = h_t - P_{22}^{-1} P_{23}\epsilon_{t+1}. \tag{A.14}$$

We can also solve uniquely for the Lagrange multiplier,

$$\begin{aligned}\tilde{\varphi}_t &= -\beta \mathbf{P}_2 E_t \mathbf{z}_{t+1} \\ &= -\beta P_{21}\tilde{y}_t - \beta P_{22} h_t - \beta[P_{23}\Gamma + P_{24}]\xi_t.\end{aligned} \tag{A.15}$$

Eqs. (A.12) and (A.14) completely describe the optimal dynamics of the variables $\{\tilde{y}_t, \tilde{h}_t\}$, starting from some initial conditions $(\tilde{y}_{t_0-1}, \tilde{h}_{t_0})$, given the evolution of the exogenous states $\{\xi_t\}$. The system consisting of these solutions for $\tilde{y}_t$ and $\tilde{h}_{t+1}(\xi_{t+1})$, together with the law of motion (40), can be written in the form (51), for certain matrices Φ and Ψ. If we partition Φ in the same way as P, it follows from the form of the solutions obtained above that $\Phi_{ij} = 0$ for all $i \geqslant 2$, $j \leqslant 2$. From this (together with our assumption about the eigenvalues of Γ) it follows that all eigenvalues of Φ have modulus less than $\beta^{-1/2}$ if and only if all eigenvalues of the block Φ_{11}, defined by (52), have this property. Hence there exists a determinate solution to the first-order conditions for optimal policy, *i.e.*, a unique solution satisfying the bound (33), if and only if M and P_{22} are non-singular matrices, and all eigenvalues of Φ_{11} have modulus less than $\beta^{-1/2}$.

Note that the solution (51) involves elements of the matrix P. We can solve for those elements of P in the following way. It follows from the assumed representation (43) for the value function that the vector of partial derivatives with respect to $\tilde{y}_{t-1}$ will equal

$$\bar{V}_1^Q = \mathbf{P}_1 \mathbf{z}_t.$$

On the other hand, application of the envelope theorem to the problem (44) implies that

$$\bar{V}_1^Q = G_1' y_t^\dagger = -G_1' M^{-1} G \mathbf{z}_t. \tag{A.16}$$

Equating the corresponding coefficients in these two representations, we observe that

$$P_{1j} = -G_1' M^{-1} G_j$$

for $j = 1, 2, 3, 4$. A similar argument implies that

$$P_{2j} = -G_2' M^{-1} G_j \tag{A.17}$$

for $j = 1, 2, 3, 4$.

These expressions involve the matrix M, which depends on P_{11}; but the subsystem (47) of these equations represents a set of n_y^2 equations to solve for the n_y^2 elements of P_{11}. Once we have solved for P_{11}, we know the matrix M, and can solve for the other elements of P. In particular, we can solve for P_{22} using (49), and check whether it is non-singular, as required in (A.14). The other elements of P can be solved for using the same method.[51]

[50] If P_{22} is singular, it is obvious that there are multiple solutions for $\tilde{h}_{t+1}(\xi_{t+1})$ consistent with the first-order conditions, but one might wonder if these correspond to multiple state-contingent evolutions $\{\tilde{y}_t\}$. In fact they do, for a single state-contingent evolution $\{\tilde{y}_t\}$ is consistent with only one process $\{\tilde{h}_t\}$, which can be determined from (42).

[51] Details of the algebra are provided in a note on computational issues available from the authors.

Thus far, we have discussed only the implications of the FOCs for the single-period optimization problem. Again, the question arises whether a solution to the first-order conditions corresponds to a maximum of (44). The second-order conditions for a finite-dimensional optimization problem are well known. First, the objective is strictly concave in $\tilde{y}_t$ if and only if the matrix $A_0 + \beta P_{11}$ is such that

$$\tilde{y}'[A_0 + \beta P_{11}]\tilde{y} < 0$$

for all $\tilde{y} \neq 0$ such that

$$C_0\tilde{y} = 0, \qquad D_0\tilde{y} = 0.$$

Using a result of Debreu (1952),[52] we can state algebraic conditions on these matrices that are easily checked. For each r such that $n_F + n_g + 1 \leqslant r \leqslant n_y$, let M_r be the lower-right square block of M of size $n_F + n_g + r$.[53] Then the concavity condition stated above holds if and only if $\det M_r$ has the same sign as $(-1)^r$, for each $n_F + n_g + 1 \leqslant r \leqslant n_y$. Note that in the case that policy is *unidimensional* — meaning that there is a single instrument to set each period, which suffices to determine the evolution of the endogenous variables, so that $n_F + n_g = n_y - 1$ — then this requirement reduces to the single condition that the determinant of M have the same sign as $(-1)^{n_y}$.

Second, in each possible state ξ_{t+1} in the following period, the continuation objective $\bar{V}^Q(\mathbf{z}_{t+1})$ is a concave function of $\tilde{h}_{t+1}(\xi_{t+1})$ if and only if the submatrix P_{22} is *negative definite, i.e.*, such that $\tilde{h}' P_{22}\tilde{h} < 0$ for all $\tilde{h} \neq 0$. This condition is also straightforward to check using the Debreu theorem: the principal minors of P_{22} must have alternating signs.

These two conditions are obviously necessary for strict concavity of the single-period problem, and hence for strict concavity of the infinite-horizon optimal policy problem. In fact, they are also sufficient, as established by Proposition 3.

A.4. Proposition 3

Proposition 3. *Suppose that the exogenous disturbances have a law of motion of the form* (40), *where Γ is a matrix the eigenvalues of which all have modulus less than $\beta^{-1/2}$, and that the constraints satisfy the rank condition* (46), *where $n_F + n_g < n_y$. Then the LQ policy problem has a determinate solution, given by* (51), *if and only if* (i) *there exists a solution P_{11} to Eq.* (47) *such that for each of the minors of the matrix M defined in* (48), $\det M_r$ *has the same sign as* $(-1)^r$, *for each $n_F + n_g + 1 \leqslant r \leqslant n_y$*; (ii) *the eigenvalues of the matrix Φ_{11} defined in* (52) *all have modulus less than $\beta^{-1/2}$; and* (iii) *the matrix P_{22} defined in* (49) *is negative definite, i.e., is such that its rth principle minor has the same sign as* $(-1)^r$, *for each $1 \leqslant r \leqslant n_g$.*

Proof. (1) The discussion in the text has already established the necessity of each of conditions (i)–(iii), so it remains only to show that they are also sufficient for the solution (51) to represent a solution to the original infinite-horizon optimal policy problem. We shall do this by establishing that conditions (i)–(iii) imply that the sufficient conditions of Proposition 2 are satisfied by this solution.

[52] See also Theorem 1.E.17 of Takayama (1985).

[53] Given (46), we can order the elements of $\tilde{y}_t$ so that the left $(n_F + n_g) \times (n_F + n_g)$ block of the matrix in (46) is non-singular, and we assume that this has been done when forming these submatrices.

We begin by establishing that the processes $\{\tilde{y}_t, \tilde{\lambda}_t, \tilde{\varphi}_t\}$ associated with the solution (51) satisfy the first-order conditions (A.1) for the infinite-horizon problem. We have already shown in the text that under conditions (i)–(iii), there exists a determinate solution (51) for the dynamics of $\{\mathbf{z}_t\}$, that it satisfies the bound (33) along with the constraints (25)–(27), and that associated with it are a unique system of Lagrange multipliers $\{\tilde{\lambda}_t, \tilde{\psi}_t, \tilde{\varphi}_t\}$, the solution for which has also been explained in the text. We wish to show that these processes must satisfy (A.1) for each $t \geqslant t_0$.

By construction, the processes $\{y_t^\dagger\}$ satisfy the first-order conditions (A.11) for each $t \geqslant t_0$. Moreover, it follows from (A.16) that

$$\mathbf{P}_1 E_t \mathbf{z}_{t+1} = G_1' E_t y_{t+1}^\dagger.$$

Substituting this into (A.11), we obtain

$$\frac{1}{2} E_t\{[A(L) + A'(\beta L^{-1})]\tilde{y}_t\} + E_t[B(L)\xi_{t+1}] + E_t[C'(\beta L^{-1})\tilde{\lambda}_t] + E_t[D'(\beta L^{-1})\tilde{\psi}_t] = 0 \tag{A.18}$$

for each $t \geqslant t_0$.

Differentiating $\bar{V}^Q(\mathbf{z}_t)$ with respect to $\tilde{h}_t$, and using the envelope theorem as in the derivation of (A.16), we obtain $\bar{V}_2^Q = -\tilde{\psi}_t$, from which we conclude that

$$\mathbf{P}_2 \mathbf{z}_t = -\tilde{\psi}_t$$

for each $t \geqslant t_0$. Comparison with first-order condition (A.13) for the optimal choice of $\tilde{h}_{t+1}$ in the recursive policy problem indicates that

$$\tilde{\psi}_t = \beta^{-1}\tilde{\varphi}_{t-1} \tag{A.19}$$

for each $t \geqslant t_0 + 1$. We may assume (as a definition of $\tilde{\varphi}_{t_0-1}$[54]) that (A.19) holds when $t = t_0$ as well. Then use of (A.19) to substitute for the process $\{\tilde{\psi}_t\}$ in (A.18) yields (A.1), which accordingly must hold for each $t \geqslant t_0$. Hence the processes constructed to satisfy the first-order conditions of the recursive policy problem must satisfy the first-order conditions for the infinite-horizon policy problem characterized in Section 3.1 as well.

(2) It remains to show that conditions (i)–(iii) also imply that the strict concavity condition (A.8) is satisfied. Let us consider an arbitrary process $\tilde{y} \in \mathcal{H}_1$, and associated with it define the process $\tilde{h}$ by

$$\tilde{h}_t = D(L)\tilde{y}_t \tag{A.20}$$

for each $t \geqslant t_0 + 1$, and by the stipulation that $\tilde{h}_{t_0} = 0$. We thus obtain a pair of processes satisfying

$$C(L)\tilde{y}_t = 0, \tag{A.21}$$

$$D(L)\tilde{y}_t = \tilde{h}_t, \tag{A.22}$$

$$E_t\tilde{h}_{t+1} = 0 \tag{A.23}$$

for all $t \geqslant t_0$. These are furthermore an example of a process $\{\mathbf{z}_t\}$ consistent with the constraints of the recursive policy problem, in the case that $\xi_t = 0$ at all times and the initial pre-commitment is given by $\tilde{h}_{t_0} = 0$.

[54] Note that $\tilde{\varphi}_{t_0-1}$ has no other meaning in the analysis of the recursive policy problem presented in Section 3.2.

We note that the analysis given in the text of the single-period problem of maximizing (44), applied to the special case in which $\xi_t = 0$ at all times,[55] implies that for any values of $\tilde{y}_{t-1}$ and $\tilde{h}_t$, the maximum possible attainable value of the objective

$$\frac{1}{2}\tilde{y}_t' A(L)\tilde{y}_t + \frac{\beta}{2}E_t\big[\mathbf{z}_{t+1}' P\mathbf{z}_{t+1}\big]$$

consistent with constraints (A.21)–(A.23) is equal to

$$\frac{1}{2}\mathbf{z}_t' P\mathbf{z}_t;$$

and this value is attained only if

$$\mathbf{z}_{t+1} = \Phi\mathbf{z}_t$$

with certainty, which is to say, only if

$$\tilde{y}_t = \Phi_{11}\tilde{y}_{t-1} + \Phi_{12}\tilde{h}_t \tag{A.24}$$

and

$$\tilde{h}_{t+1} = 0 \tag{A.25}$$

in each possible state in period $t+1$.

Thus the fact that the processes $\{\tilde{y}_t, \tilde{h}_t\}$ satisfy (A.21)–(A.23) for all $t \geqslant t_0$ implies that

$$\frac{1}{2}\tilde{y}_t' A(L)\tilde{y}_t + \frac{\beta}{2}E_t\big[\mathbf{z}_{t+1}' P\mathbf{z}_{t+1}\big] \leqslant \frac{1}{2}\mathbf{z}_t' P\mathbf{z}_t$$

for all $t \geqslant t_0$, and that the inequality is strict unless (A.24)–(A.25) hold. Now if conditions (A.24)–(A.25) hold for all $t \geqslant t_0$, $\tilde{y}_t = 0$ at all times. Thus in the case that $\tilde{y}_t$ is not equal to zero almost surely for all t, there must be at least one date t_1 such that at least one of these conditions is violated with positive probability when $t = t_1$. In that case, there must be some $k > 0$ such that

$$E_{t_0}\left\{\frac{1}{2}\tilde{y}_{t_1}' A(L)\tilde{y}_{t_1} + \frac{\beta}{2}\mathbf{z}_{t_1+1}' P\mathbf{z}_{t_1+1}\right\}\frac{1}{2} \leqslant E_{t_0}\mathbf{z}_{t_1}' P\mathbf{z}_{t_1} - k.$$

It then follows, by summing these inequalities (appropriately discounted) for successive periods, that

$$E_{t_0}\sum_{t=t_0}^{T}\beta^{t-t_0}\frac{1}{2}\tilde{y}_t' A(L)\tilde{y}_t + \frac{\beta^{T+1-t_0}}{2}E_{t_0}\mathbf{z}_{T+1}' P\mathbf{z}_{T+1} \leqslant \frac{1}{2}\mathbf{z}_{t_0}' P\mathbf{z}_{t_0} - k = -k, \tag{A.26}$$

for all $T \geqslant t_1$.

As we have stipulated that the process $\tilde{y}$ is an element of $\mathcal{H}_1$, and thus satisfies the bound (33), we necessarily have

$$\lim_{T\to\infty}\beta^{T+1}E_{t_0}\mathbf{z}_{T+1}' P\mathbf{z}_{T+1} = 0.$$

[55] It follows from the usual principle of certainty equivalence for LQ problems that the matrices characterizing the solution to this problem do not depend on the value of the variance-covariance matrix Σ for the disturbances. In fact, it is easily observed that the derivations given in the text would apply equally to a problem in which $\xi_t = 0$ at all times.

(Note that it follows from (A.20) that the elements of $\tilde{h}$ cannot grow asymptotically at a faster rate than do the elements of $\tilde{y}$.) It then follows from (A.26) that

$$\limsup_{T\to\infty} E_{t_0} \sum_{t=t_0}^{T} \beta^{t-t_0} \frac{1}{2} \tilde{y}_t' A(L) \tilde{y}_t \leqslant -k. \tag{A.27}$$

But since it follows from the assumption that $\tilde{y}$ satisfies (33) that the series in (A.27) has a limit, this limit must be no greater than $-k$. Hence $\tilde{y}$ satisfies (A.8), and all of the sufficient conditions of Proposition 2 have been verified. This establishes the proposition. □

Example. Suppose that y_t has two elements, that the objective of policy is to maximize

$$V_{t_0}^{Q}(\tilde{y}) \equiv \frac{1}{2} E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \tilde{y}_t' A \tilde{y}_t, \tag{A.28}$$

where A is a symmetric 2×2 matrix, and that the only constraint on what policy can achieve is a single, forward-looking constraint

$$E_t[\delta \tilde{y}_{1,t} - \tilde{y}_{1,t+1}] = 0 \tag{A.29}$$

for all $t \geqslant t_0$, where $|\delta| < \beta^{-1/2}$. There are no exogenous disturbances, but the expectations appear because we wish to consider the possibility of (arbitrarily) randomized policies. We assume an initial pre-commitment of the form

$$\tilde{y}_{1,t_0} = \delta \tilde{y}_{1,t_0-1} + \tilde{h}_{t_0}, \tag{A.30}$$

for some quantity $\tilde{h}_{t_0}$.

In the case that policy is restricted to be deterministic, the constraint completely determines the path of $\{\tilde{y}_{1t}\}$; the only (perfect foresight) sequence consistent with the initial pre-commitment and the forward-looking constraint is the one in which

$$\tilde{y}_{1,t} = [\delta \tilde{y}_{1,t_0-1} + \tilde{h}_{t_0}] \delta^{t-t_0}$$

for all $t \geqslant t_0$. The problem then reduces to the choice of a sequence $\{\tilde{y}_{2,t}\}$, constrained only by the bound (33), so as to maximize the objective. This is obviously a concave problem if and only if $\tilde{y}' A \tilde{y}$ is a concave function of $\tilde{y}_2$ for a given value of $\tilde{y}_1$. This in turn is true if and only if $A_{22} < 0$; the other elements of A are irrelevant.

If instead we allow random policies, the condition just derived is no longer sufficient for concavity (though still necessary). One can show that the problem is concave if and only if A is a negative definite matrix. This is obviously a sufficient condition (as it implies that (A.28) is concave for arbitrary sequences). To show that it is also necessary, suppose instead that it is not true. Then there exists a vector $v \neq 0$ such that $v' A v \geqslant 0$. Now let $\{\bar{y}_t\}$ be any process satisfying the constraints (33), (A.29), and (A.30), and consider the alternative process $\{\tilde{y}_t\}$ generated by the law of motion

$$\tilde{y}_t = \bar{y}_t + \delta(\tilde{y}_{t-1} - \bar{y}_{t-1}) + v\epsilon_t$$

for each $t \geqslant t_0 + 1$, starting from the initial condition (A.30), where $\{\epsilon_t\}$ is a (scalar-valued) martingale-difference sequence satisfying the bound (33). One can easily show that the process $\{\tilde{y}_t\}$ satisfies (33), (A.29), and (A.30) as well; moreover, the value of the objective in the case of this process satisfies

$$V_{t_0}^Q(\tilde{y}) = V_{t_0}^Q(\bar{y}) + \left(1 - \beta\delta^2\right)^{-1} v' A v \, E_{t_0} \sum_{t=t_0+1}^{\infty} \beta^t \epsilon_t^2$$
$$\geqslant V_{t_0}^Q(\bar{y}).$$

Since we can construct an alternative policy that is at least as good in the case of *any* policy, there is no uniquely optimal policy in such a case; and in addition, we have shown that arbitrary randomization of policy is possible without welfare loss.

Let us examine how these results compare with the conditions stated in Proposition 3. In this example, condition (47) states that

$$P_{11} = \alpha \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix},$$

where

$$\alpha = -\delta^2 \left[M^{-1}\right]_{33}.$$

This form for P_{11} implies in turn that M is invertible as long as $A_{22} \neq 0$, and that in that case,

$$\left[M^{-1}\right]_{33} = -\alpha\beta - \frac{|A|}{A_{22}}.$$

Hence we obtain a unique solution,

$$\alpha = \frac{\delta^2}{1 - \beta\delta^2} \frac{|A|}{A_{22}}.$$

Since $n_F = 0$, $n_g = 1$, $n_y = 2$, condition (i) of the proposition holds if and only if $\det M_2 = \det M > 0$, and under the above solution for P_{11}, $\det M = -A_{22}$; hence condition (i) reduces to the requirement that $A_{22} < 0$.

This solution for P_{11}, and hence for M, also implies that

$$\Phi_{11} = \begin{bmatrix} \delta & 0 \\ -\delta A_{21}/A_{22} & 0 \end{bmatrix}.$$

Hence the eigenvalues of Φ_{11} are 0 and δ. Thus under our assumption about δ, condition (ii) is necessarily satisfied, as long as $A_{22} \neq 0$ (so that Φ_{11} exists).

We observe that both conditions (i) and (ii) hold if and only if $A_{22} < 0$, which is just the concavity condition derived above for the deterministic policy problem.

The solution for P_{11} similarly implies that

$$P_{22} = -G_2' M^{-1} G_2 = -\left[M^{-1}\right]_{33} = \frac{1}{1 - \beta\delta^2} \frac{|A|}{A_{22}}.$$

Since the numerator in this last expression is positive, condition (iii) holds (in addition to the other two conditions) if and only if we also have $\det A > 0$. Since A is negative definite if and only if $A_{22} < 0$ and $\det A > 0$, we can alternatively state that condition (iii) holds (in addition to the other two) if and only if A is also negative definite. This is the additional condition derived above for concavity in the case of stochastic policies.

A.5. Lemma 4

We begin by explaining the proposed form (56) for the initial pre-commitment. Let us define a new extended state vector

$$\hat{\mathbf{z}}_t \equiv \begin{bmatrix} \tilde{y}_{t-1} \\ \hat{h}(\xi_t, \xi_{t-1}) \\ \xi_t \\ \xi_{t-1} \end{bmatrix},$$

where[56]

$$\hat{h}(\xi_t, \xi_{t-1}) \equiv h_{t-1} - P_{22}^{-1} P_{23}(\xi_t - \Gamma \xi_{t-1}),$$

then it follows from (A.14) that under the solution to the recursive policy problem, $\mathbf{z}_t = \hat{\mathbf{z}}_t$ for each $t \geqslant t_0 + 1$. (However, $\hat{\mathbf{z}}_t$, unlike $\mathbf{z}_t$, is a function solely of $\tilde{y}_{t-1}$ and the history of the exogenous disturbances.) Hence

$$\tilde{h}_{t_0} = \hat{h}(\xi_{t_0}, \xi_{t_0-1}) \tag{A.31}$$

is a self-consistent constraint of the form (5). This is just the initial pre-commitment (56) proposed in the text.

Lemma 4. *The Lagrange multiplier associated with the initial pre-commitment (A.31) is equal to*

$$\beta^{-1} \tilde{\varphi}_{t_0-1} = \varphi^*(\mathbf{y_{t_0-1}}) \equiv \tilde{\psi}(\tilde{y}_{t_0-1}, \xi_{t_0-1}), \tag{A.32}$$

where

$$\tilde{\psi}(\tilde{y}_{t-1}, \xi_{t-1}) \equiv [0\ 0\ -I] M^{-1} G \begin{bmatrix} \tilde{y}_{t-1} \\ h_{t-1} \\ \Gamma \xi_{t-1} \\ \xi_{t-1} \end{bmatrix}. \tag{A.33}$$

Proof. If we assume the initial pre-commitment, we can then define the optimal dynamics from a timeless perspective as functions solely of the initial conditions $(\tilde{y}_{t_0-1}, \xi_{t_0-1})$ and the evolution of the exogenous states $\{\xi_t\}$ from period t_0 onward. Substituting (A.14) for the pre-commitment $\tilde{h}_{t+1}$ in the solution (A.12) for the optimal choice of $\tilde{y}_{t+1}$, we observe that under the solution to the recursive policy problem (and hence under the solution to the original problem as well), $\tilde{y}_{t+1}$ is a linear function of $\tilde{y}_t$, ξ_{t+1}, and ξ_t, for each $t \geqslant t_0$. This solution together with the process (40) for the exogenous disturbances imply a law of motion of the form (57) for the extended state vector $\bar{\mathbf{y}}_t$ defined in the text.

If we assume an initial pre-commitment (A.31), it also follows from (A.12) that

$$\tilde{\psi}_t = [0\ 0\ -I] M^{-1} G \hat{\mathbf{z}}_t \tag{A.34}$$

is the Lagrange multiplier associated with the pre-commitment each period in the recursive problem. Moreover, because the only constraint on the way in which $\tilde{h}_{t+1}(\xi_{t+1})$ can be chosen for

[56] Here it should be recalled that h_{t-1} is a linear function of ξ_{t-1}, defined in (31).

the following period is given by the expected-value constraint (45), the first-order conditions for optimal policy imply that $\tilde{\psi}_t = E_{t-1}\tilde{\psi}_t$ for each $t \geqslant t_0 + 1$,[57] and hence that

$$\begin{aligned}\tilde{\psi}_t &= [0\ 0\ -I]M^{-1}GE_{t-1}\hat{\mathbf{z}}_t \\ &= \tilde{\psi}(\tilde{y}_{t-1}, \xi_{t-1})\end{aligned}$$

where the function $\tilde{\psi}(\tilde{y}_{t-1}, \xi_{t-1})$ is defined as in (A.33).

Consistency of this result with (A.34) implies that the right-hand side of (A.34) must be equivalent to $\tilde{\psi}(\tilde{y}_{t-1}, \xi_{t-1})$; that is, that the coefficients multiplying $\tilde{y}_{t-1}$, ξ_t, and ξ_{t-1} must be the same in both expressions. But since (A.34) must hold at $t = t_0$ as well, in the case of an initial pre-commitment (A.31), and not only for $t \geqslant t_0 + 1$, it follows that under such a pre-commitment,

$$\tilde{\psi}_t = \tilde{\psi}(\tilde{y}_{t-1}, \xi_{t-1})$$

for all $t \geqslant t_0$.

In the case that $t = t_0$, the multiplier $\tilde{\psi}_{t_0}$ associated with the initial pre-commitment is the one that is denoted $\beta^{-1}\tilde{\varphi}_{t_0-1}$ in (A.2) and in (55). From this, (A.32) is immediately established. □

A.6. Lemma 5

Lemma 5. *Suppose that under optimal policy, the extended state vector* $\mathbf{y_t}$ *consists entirely of components that are either* (i) *stationary, or* (ii) *pure random walks. Suppose also that the class of policy rules* $\mathcal{R}$ *is such that each rule in the class implies convergence to the same long-run values of the state variables as under optimal policy, in the absence of stochastic disturbances, so that the initial value of the trend component* $\mathbf{y}^{tr}_{t_0-1}$ *is the same regardless of the rule* r *that is considered. Then for any rule* $r \in \mathcal{R}$*, the objective*

$$\Omega_r\left(\mathbf{y}^{tr}_{\mathbf{t_0-1}}\right) \equiv E_\mu\left[\bar{W}_r(\mathbf{y_{t_0-1}})\right], \tag{A.35}$$

can be decomposed into two parts,

$$\Omega_r\left(\mathbf{y}^{tr}_{\mathbf{t_0-1}}\right) = \Omega^1\left(\mathbf{y}^{tr}_{\mathbf{t_0-1}}\right) + \Omega^2_r, \tag{A.36}$$

where the first component is the same for all rules in this class, while the second component is independent of the initial condition $\mathbf{y}^{tr}_{\mathbf{t_0-1}}$.

Proof. We restrict attention to a class of rules $\mathcal{R}$ with the property that each rule in the class implies convergence to the same long-run values of the state variables as under optimal policy, in the absence of stochastic disturbances. Because we analyze the dynamics under a given policy using a linearized version of the structural relations, certainty-equivalence obtains, and it follows that the limiting behavior (as $T \to \infty$) of the long-run forecast $E_{t_0}[\mathbf{y_T}]$ must also be the same under any rule $r \in \mathcal{R}$, given the initial conditions $\mathbf{y_{t_0-1}}$. Thus given these initial conditions, the decomposition of the initial extended state vector into components $\mathbf{y}^{tr}_{\mathbf{t_0-1}}$ and $\mathbf{y}^{cyc}_{\mathbf{t_0-1}}$ is the same under any rule $r \in \mathcal{R}$.

Let us consider the decomposition

$$\tilde{y}_t = \bar{y}_t + \hat{y}_t,$$

[57] In fact, one can show that $\tilde{\psi}_t = \beta^{-1}\tilde{\varphi}_{t-1}$ for each $t \geqslant t_0 + 1$. This follows from differentiation of the value function $V^Q(\mathbf{z}_{t+1})$ with respect to $\tilde{h}_{t+1}$ using the envelope theorem, and comparison of the result with (A.13).

where $\{\bar{y}_t\}$ is the deterministic sequence

$$\bar{y}_t \equiv E_{t_0-1}\tilde{y}_t$$

and $\hat{y}_t$ is the component of $\tilde{y}_t$ that is unforecastable as of date $t_0 - 1$. Then if we evaluate

$$\bar{W}(\tilde{y}; \mathbf{y}_{t_0-1}) \equiv E_{t_0-1} W(\tilde{y}; \xi_{t_0}, \mathbf{y}_{\mathbf{t_0-1}}),$$

where W is the quadratic form defined in (60), under the evolution implied by any rule r, we find that

$$\bar{W}(\tilde{y}; \mathbf{y}_{t_0-1}) = \bar{W}(\bar{y}; \mathbf{y}_{t_0-1}) + \bar{W}(\hat{y}; \mathbf{y}_{t_0-1}). \tag{A.37}$$

Here all the cross terms in the quadratic form have conditional expectation zero because $\bar{y}$ is deterministic while $\hat{y}$ is unforecastable.

Moreover, under any rule r, the value of $\hat{y}_t$ is a linear function of the sequence of unexpected shocks between periods t_0 and t, that is independent of the initial state. (This independence follows from the linearity of the law of motion (61), under the linear approximation that we use to solve for the equilibrium dynamics under a given policy rule.) Hence the second term on the right-hand side of (A.37),[58]

$$\bar{W}(\hat{y}; \mathbf{y}_{t_0-1}) = E_{t_0-1} V^Q_{t_0}(\hat{y}),$$

is independent of the initial state $\mathbf{y}_{t_0-1}$ as well. Let $\bar{W}^2_r$ denote the value of this expression associated with a given rule r.

Instead, the value of $\bar{y}_t$ will be a linear function of $\mathbf{y}_{t_0-1}$, again as a result of the linearity of (61). And in our LQ problem with a self-consistent initial pre-commitment, the function (59) is linear as well. It follows that the first term on the right-hand side of (A.37) is a quadratic function of $\mathbf{y}_{t_0-1}$,

$$\bar{W}(\bar{y}; \mathbf{y}_{t_0-1}) = \mathbf{y}'_{t_0-1} \Xi_r \mathbf{y}_{t_0-1},$$

where the subscript r indicates that the matrix of coefficients Ξ_r can depend on the policy rule that is chosen. Then substituting $\mathbf{y}^{tr}_{t_0-1} + \mathbf{y}^{cyc}_{t_0-1}$ for $\mathbf{y}_{t_0-1}$ in the above expression, and integrating over possible initial values of the cyclical component, for a given initial value of the trend component, we observe that

$$E_\mu\big[\bar{W}(\bar{y}; \mathbf{y}_{t_0-1})\big] = \mathbf{y}^{tr\prime}_{t_0-1} \Xi_r \mathbf{y}^{tr}_{t_0-1} + E_\mu\big[\mathbf{y}^{cyc\prime} \Xi_r \mathbf{y}^{cyc}\big], \tag{A.38}$$

using the fact that $E_\mu[\mathbf{y}^{cyc}] = 0$.

Finally, we observe that under any rule r, the linearity of the law of motion (61) implies that conditional forecasts of the evolution of the endogenous variables take the form

$$E_{t_0-1} y_T = y^{tr}_{t_0-1} + B_{T+1-t_0} \mathbf{y}^{cyc}_{\mathbf{t_0-1}},$$

where the sequence of matrices $\{B_j\}$ may depend on the rule r, but the first term on the right-hand side is the same for all rules in the class $\mathcal{R}$. Using this solution for the sequence $\bar{y}$ to evaluate $\bar{W}(\bar{y}; \mathbf{y}_{t_0-1})$, we find that the first term in (A.38) must be a quadratic function of $y^{tr}_{t_0-1}$ that is the same for all rules r, that can be denoted $y^{tr\prime}_{t_0-1} \bar{\Xi} y^{tr}_{t_0-1}$. Thus if we integrate (A.37) over the invariant distribution μ, we obtain

[58] Here the expected value of the second term on the right-hand side of (60) vanishes because of the unforecastability of $\hat{y}_{t_0}$.

$$E_\mu\big[\bar{W}_r(\mathbf{y_{t_0-1}})\big] = y_{t_0-1}^{tr\prime}\bar{\Xi}y_{t_0-1}^{tr} + E_\mu\big[\mathbf{y}^{cyc\prime}\Xi_r\mathbf{y}^{cyc}\big] + \bar{W}_r^2,$$

which is precisely a decomposition of the asserted form (A.36). This proves that the criterion (A.35) establishes the same ranking of alternative rules, regardless of the initial condition. □

A.7. Computing the invariant measure μ

We need to know the invariant distribution μ over possible initial conditions under optimal policy, in order to compute the proposed welfare criterion (65). Because $\bar{W}_r(\cdot)$ is a quadratic function, we only need to compute the unconditional mean and variance-covariance matrix of $\mathbf{y_t^{cyc}}$ under optimal policy.

Let us recall from (57)–(58) that the optimal dynamics imply a law of motion

$$\mathbf{y}_{t+1} = \bar{\Phi}\mathbf{y}_t + \bar{\Psi}\epsilon_{t+1} \tag{A.39}$$

for the extended state vector

$$\mathbf{y}_t \equiv \begin{bmatrix} \tilde{y}_t \\ \xi_t \end{bmatrix}.$$

Under the law of motion (A.39) implied by a pre-commitment of the form (A.14), the trend component of the extended state vector $\bar{\mathbf{y}}_t$ is given by $\bar{\mathbf{y}}_t^{tr} = \Pi\bar{\mathbf{y}}_t$, where Π is the matrix[59]

$$\Pi \equiv \lim_{j\to\infty} \bar{\Phi}^j,$$

and the cyclical component is correspondingly given by $\bar{\mathbf{y}}_t^{cyc} = [I - \Pi]\bar{\mathbf{y}}_t$. It then follows that the law of motion for the cyclical component is

$$\bar{\mathbf{y}}_{t+1}^{cyc} = \bar{\Phi}\bar{\mathbf{y}}_t^{cyc} + [I - \Pi]\bar{\Psi}\epsilon_{t+1}. \tag{A.40}$$

We note furthermore that (A.40) describes a jointly stationary set of processes, since the matrix $\bar{\Phi}$ is stable on the subspace of vectors $\mathbf{v}$ of the form $\mathbf{v} = [I - \Pi]\bar{\mathbf{y}}$ for some vector $\bar{\mathbf{y}}$.[60] Hence there exist a well-defined vector of unconditional means $\mathbf{E}$ and an unconditional variance-covariance matrix $\mathbf{V}$. The unconditional means are all zero, while the matrix V is given by the solution to the linear equation system

$$\mathbf{V} = \bar{\Phi}\mathbf{V}\bar{\Phi}' + [I - \Pi]\bar{\Psi}\Sigma\bar{\Psi}'\big[I - \Pi'\big].$$

In the case of some policy rules, it may be necessary to include additional lags of $\tilde{y}_t$ or ξ_t in the extended state vector $\mathbf{y}_t$, in order for the equilibrium dynamics under the rule r to have a representation of the form (61). However, in this case, the additional elements of $\mathbf{y}_t^{cyc}$ will all be lags of elements of the vector considered above. Hence the law of motion (A.40) can be used to derive the relevant unconditional moments in this case as well (though we omit the algebra).

[59] Under the assumption (made in the text) that the extended state vector is difference-stationary, this limit must be well defined.

[60] When restricted to this subspace, the operator $\bar{\Phi}$ has eigenvalues consisting of those eigenvalues of $\bar{\Phi}$ that are less than one in modulus; these are in turn a subset of the eigenvalues of Φ that are less than one in modulus (some zero eigenvalues have been dropped).

References

[1] F. Altissimo, V. Cúrdia, D. Rodriguez Palenzuela, Linear-quadratic approximation to optimal policy: an algorithm and two applications, paper presented in: The Conference "Quantitative Analysis of Stabilization Policies", Columbia University, September 2005.

[2] D. Backus, J. Driffill, The consistency of optimal policy in stochastic rational expectations models, CEPR Discussion Paper No. 124, 1986.

[3] R.J. Barro, On the determination of public debt, J. Polit. Economy 87 (1979) 940–971.

[4] G. Benigno, P. Benigno, Designing targeting rules for international monetary policy cooperation, J. Monet. Econ. 53 (3) (2006) 473–506.

[5] G. Benigno, B. De Paoli, On the international dimension of fiscal policy, J. Money, Credit, Banking 42 (2010) 1523–1542.

[6] P. Benigno, M. Woodford, Optimal monetary and fiscal policy: a linear-quadratic approach, in: M. Gertler, K. Rogoff (Eds.), NBER Macroecon. Ann., vol. 2003, MIT Press, Cambridge, 2004, pp. 271–333.

[7] P. Benigno, M. Woodford, Inflation stabilization and welfare: the case of a distorted steady state, J. Europ. Econ. Assoc. 3 (6) (2005) 1185–1236.

[8] P. Benigno, M. Woodford, Optimal stabilization policy when wages and prices are sticky: the case of a distorted steady state, in: J. Faust, A. Orphanides, D. Reifschneider (Eds.), Models Monet. Pol., Board of Governors of the Federal Reserve System, Washington, 2005, pp. 127–180.

[9] P. Benigno, M. Woodford, Optimal taxation in an RBC model: a linear quadratic approach, J. Econ. Dynam. Control 30 (9–10) (2006) 1445–1489.

[10] P. Benigno, M. Woodford, Linear-quadratic approximation of optimal policy problems, NBER Working Paper No. 12672, revised draft, 2008.

[11] S. Beveridge, C.R. Nelson, A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle', J. Monet. Econ. 7 (2) (1981) 151–174.

[12] T. Berriel, D. Sinigaglia, Optimal fiscal and monetary policy under sectoral heterogeneity, unpublished, Princeton University, 2008.

[13] A.P. Blake, A 'timeless perspective' on optimality in forward-looking rational expectations models, NIESR Discussion Papers 188, National Institute of Economic and Social Research, 2001.

[14] G. Coenen, G. Lombardo, F. Smets, R. Straub, International transmission and monetary policy cooperation, in: J. Gali, M. Gertler (Eds.), International Dimensions of Monetary Policy, University of Chicago Press, Chicago, 2009.

[15] V. Cúrdia, Optimal monetary policy under sudden stops, Federal Reserve Bank of New York Staff Report No. 323, 2008.

[16] T. Damjanovic, V. Damjanovic, C. Nolan, Linear-quadratic approximation to unconditionally optimal policy: the distorted steady state, J. Monet. Econ. 55 (2008) 491–500.

[17] R. Dennis, When is discretion superior to timeless perspective policymaking?, J. Monet. Econ. 57 (2010) 266–277.

[18] B. De Paoli, Monetary policy and welfare in a small open economy, J. Int. Econ. 77 (2009) 11–22.

[19] B. De Paoli, Monetary policy in a small open economy: the role of the asset market structure, J. Money, Credit, Banking 41 (2009) 1301–1330.

[20] C.J. Erceg, D.W. Henderson, A.T. Levin, Optimal monetary policy with staggered wage and price contracts, J. Monet. Econ. 46 (2000) 281–313.

[21] C. Engel, Currency misalignments and optimal monetary policy: a reexamination, NBER Working Paper No. 14829, 2009.

[22] A. Ferrero, Fiscal and monetary rules for a currency union, J. Int. Econ. 77 (2009) 1–10.

[23] W.H. Fleming, Stochastic control for small noise intensities, SIAM J. Control 9 (1971) 473–517.

[24] M. Giannoni, M. Woodford, Optimal target criteria for stabilization policy, NBER Working Paper No. 15757, 2010.

[25] M. Horvath, Simple monetary-fiscal targeting rules, CDMA Working Paper No. 0801, University of St. Andrews, 2008.

[26] C. Jensen, B.C. McCallum, The non-optimality of proposed monetary policy rules under timeless-perspective commitment, NBER Working Paper No. 8882, 2002.

[27] H. Jin, K.L. Judd, Perturbation methods for general dynamic stochastic models, unpublished, Stanford University, 2002.

[28] K.L. Judd, Approximation, perturbation and projection methods in economic analysis, in: H.M. Amman, D.A. Kendrick, J. Rust (Eds.), Handbook of Computational Economics, vol. 1, North-Holland, Amsterdam, 1996.

[29] K.L. Judd, Numerical Methods in Economomics, MIT Press, Cambridge, MA, 1998.

[30] K.L. Judd, S. Guu, Asymptotic methods for asset market equilibrium analysis, Econ. Theory 18 (2001) 127–157.
[31] M. Juillard, F. Pelgrin, Computing optimal policy in a timeless perspective: an application to a small open economy, Bank of Canada Working Paper No. 2007-32, 2007.
[32] D.A. Kendrick, Stochastic control for economic models: past, present and the paths ahead, J. Econ. Dynam. Control 29 (2005) 3–30.
[33] A. Khan, R.G. King, A.L. Wolman, Optimal Monetary Policy, Rev. Econ. Stud. 70 (4) (2003) 825–860.
[34] J. Kim, S. Kim, Spurious welfare reversal in international business cycle modelsm, J. Int. Econ. 60 (2) (2003) 471–500.
[35] J. Kim, S. Kim, Two pitfalls of linearization methods, J. Money, Credit, Banking 39 (2006) 995–1001.
[36] J. Kim, S. Kim, E. Schaumburg, C.A. Sims, Calculating and using second order accurate solutions of discrete time dynamic equilibrium models, J. Econ. Dynam. Control 32 (2008) 3397–3414.
[37] J. Kim, A.T. Levin, T. Yun, Diagnosing and treating bifurcations in perturbation analysis of dynamic macro models, FEDS Discussion Paper No. 2007-14, 2007.
[38] R.G. King, A.L. Wolman, What should the monetary authority do when prices are sticky?, in: J.B. Taylor (Ed.), Monetary Policy Rules, University of Chicago Press, Chicago, 1999.
[39] P. Levine, J. Pearlman, R. Pierse, Linear-quadratic approximation, external habit, and targeting rules, J. Econ. Dynam. Control 32 (2008) 3315–3349.
[40] M.J.P. Magill, A local analysis of N-sector capital accumulation under uncertainty, J. Econ. Theory 15 (1977) 211–218.
[41] A. Marcet, R. Marimon, Recursive contracts, Universitat Pompeu Fabra Working Paper No. 337, 1988.
[42] C. Montoro, Oil shocks and optimal monetary policy, BIS Working Paper No. 307, 2010.
[43] J.J. Rotemberg, M. Woodford, An optimization-based econometric framework for the evaluation of monetary policy, in: B.S. Bernanke, J.J. Rotemberg (Eds.), NBER Macroecon. Ann., vol. 12, MIT Press, Cambridge, 1997, pp. 297–346.
[44] T.J. Sargent, Macroeconomic Theory, 2nd ed., Academic Press, New York, 1987.
[45] S. Schmitt-Grohé, M. Uribe, Solving dynamic general equilibrium models using a second-order approximation to the policy function, J. Econ. Dynam. Control 28 (2004) 755–775.
[46] S. Schmitt-Grohé, M. Uribe, Optimal fiscal and monetary policy under sticky prices, J. Econ. Theory 114 (2004) 198–230.
[47] S. Schmitt-Grohé, M. Uribe, Optimal simple and implementable monetary and fiscal rules, J. Monet. Econ. 54 (2007) 1702–1725.
[48] A. Sutherland, A simple second-order solution method for dynamic general equilibrium models, CEPR Discussion Paper No. 3554, 2002.
[49] M. Woodford, Commentary: how should monetary policy be conducted in an era of price stability?, in: New Challenges for Monetary Policy, Federal Reserve Bank of Kansas City, Kansas City, 1999.
[50] M. Woodford, Inflation stabilization and welfare, Contrib. to Macroecon. 2 (1) (2002), Article 1, www.bepress.com.
[51] M. Woodford, Interest and Prices: Foundations of a Theory of Monetary Policy, Princeton Univ. Press, Princeton, 2003.
[52] M. Woodford, Optimal monetary stabilization policy, in: B.M. Friedman, M. Woodford (Eds.), Handbook of Monetary Economics, vol. 3, Elsevier, 2010, pp. 723–828.