Robustly Optimal Monetary Policy with Near-Rational Expectations *

Michael Woodford
Columbia University

November 11, 2005

Abstract

The paper considers optimal monetary stabilization policy in a forward-looking model, when the central bank recognizes that private-sector expectations need not be precisely model-consistent, and wishes to choose a policy that will be as good as possible in the case of any beliefs that are close enough to model-consistency. It is found that commitment continues to be important for optimal policy, that the optimal long-run inflation target is unaffected by the degree of potential distortion of beliefs, and that optimal policy is even more history-dependent than if rational expectations are assumed.

*I would like to thank Tom Sargent for comments, Mauro Roca for research assistance and the National Science Foundation for research support through a grant to the NBER.

An extensive literature has considered the optimal conduct of monetary policy under the assumption of rational (or model-consistent) expectations. This literature has found that it is quite important to take account of the effects of the systematic (and hence predictable) component of monetary policy on expectations. For example, it is found quite generally that an optimal policy commitment differs from the policy that would be chosen through a sequential optimization procedure with no advance commitment of future policy. It is also found quite generally that optimal policy is *history-dependent* — a function of past conditions that no longer affect the degree to which it would be possible to achieve stabilization aims from the present time onward.[1]

Both of these conclusions, however, depend critically on the idea that an advance commitment of future policy should change people's expectations at earlier dates. This may lead to the fear that analyses that assume rational expectations (RE) exaggerate the degree to which a policy authority can rely upon private-sector expectations to be shaped by its policy commitments in precisely the way that it expects them to be. What if the relation between what a central bank plans to do and what the public will expect to happen is not quite so predictable? Might both the case for advance commitment of policy and the case for history-dependent policy be considerably weakened under a more skeptical view of the precision with which the public's expectations can be predicted?

One way of relaxing the assumption of rational expectations is to model agents as forecasting using an econometric model, the coefficients of which they must estimate using data observed prior to some date; sampling error will then result in forecasts that depart somewhat from precise consistency with the analyst's model.[2] However, selecting a monetary policy rule on the basis of its performance under a specific model of "learning" runs the risk of exaggerating the degree to which the policy analyst can predict and hence exploit the forecasting errors that result from a particular way of extrapolating from past observations. One might even conclude that the optimal policy under learning achieves an outcome better than any possible rational-expectations equilibrium, by inducing systematic forecasting errors of a kind that happen to serve the central bank's stabilization objectives. But if such a policy were shown to be possible under some model of learning considered to be plausible (or

[1] Both points are discussed extensively in Woodford (2003, chap. 7).

[2] Examples of monetary policy analysis under assumptions of this kind about private-sector expectations include Orphanides and Williams (2005a, 2005b) and Gaspar *et al.* (2005).

even consistent with historical data), would it really make sense to conduct policy accordingly, relying on the public to continue making precisely the mistakes that the policy is designed to exploit?

It was exactly this kind of assumption of superior knowledge on the part of the policy analyst that the rational expectations hypothesis was intended to prevent. Yet as just argued, the assumption of RE also implies an extraordinary ability on the part of the policy analyst to predict exactly what the public will be expecting when policy is conducted in a particular way. In this paper, I propose instead an approach to policy analysis that does *not* assume that the central bank can be certain exactly what the public will expect if it chooses to conduct policy in a certain way. Yet neither does it neglect the fact that people are likely to catch on, at least to some extent, to systematic patterns created by policy, in analyzing the effects of alternative policies. In this approach, the policy analyst assumes that private-sector expectations should not be *too different* from what her model would predict under the contemplated policy — people are assumed to have *near-rational expectations* (NRE). But it is recognized that a range of different beliefs would all qualify as NRE. The policymaker is then advised to choose a policy that would not result in too bad an outcome under *any* NRE, *i.e.*, a *robustly* optimal policy given the uncertainty about private-sector expectations.[3]

1 Near-Rational Expectations

I can expound the general conception of robust policy that I wish to propose using an abstract two-period policy game. A vector of endogenous variables x_t is determined in two successive periods ($t = 0, 1$); there are many possible states of the world s in period 1, and x_1 may depend on s. The policymaker chooses a vector of controls u_t in each period; the value of u_1 may be contingent on the state s. As a result of optimizing behavior by the private sector, in any equilibrium, the endogenous variables $(x_0, x_1(\cdot))$ must satisfy a system of functional equations

$$F(y_0, y_1(\cdot); \mu) = 0, \tag{1.1}$$

[3]The conception of policy robustness here is similar to the one explored in detail in Hansen and Sargent (2005c), though they do not consider the particular source of uncertainty about policy outcomes treated here.

where y_t is the vector obtained by stacking x_t and u_t, and μ is the element of $\mathcal{M}$ (the set of measures over possible values of (s, y_1)) that indicates private-sector expectations in the initial period.

The policymaker wishes to choose a policy $(u_0, u_1(\cdot))$ so as to minimize an expected loss function

$$\mathrm{E}[L(y_0, y_1, s)], \tag{1.2}$$

where the expectation $\mathrm{E}[\cdot]$ is with respect to the measure $\bar{\mu} \in \mathcal{M}$ indicating the policymaker's expectations in the initial period.[4] In the case of any measure π over the possible states of the world s and any measurable function $g(\cdot)$, let $\nu_{\pi,g}$ denote the element of $\mathcal{M}$ with marginal distribution π and such that zero probability is assigned to any outcomes in which $y_1 \neq g(s)$. Then the policymaker evaluates the objective (1.2) in the case of an equilibrium $(y_0, y_1(\cdot))$ using the measure

$$\bar{\mu} = \nu_{\bar{\pi}, y_1}, \tag{1.3}$$

where $\bar{\pi}$ indicates the policymaker's beliefs about the probability of different possible states of nature s. (In the beliefs of the policymaker, $\bar{\pi}$ is given independently of the policy chosen, while $y_1(\cdot)$ and hence $\bar{\mu}$ will depend on policy.)

In rational-expectations (RE) policy analysis, the analyst assumes that in any equilibrium, the expectations of the private sector will *also* correspond to the measure $\bar{\mu} = \nu_{\bar{\pi}, y_1}$. Hence the analyst associates to any policy $(u_0, u_1(\cdot))$ an equilibrium $(y_0, y_1(\cdot))$ that satisfies

$$F(y_0, y_1(\cdot); \nu_{\bar{\pi}, y_1}) = 0,$$

and then evaluates (1.2) using the implied measure (1.3).

I shall suppose instead that the analyst recognizes that private agents may *not* have rational expectations, *i.e.*, that beliefs $\mu \neq \bar{\mu}$ are possible. But I shall suppose that he nonetheless assumes that μ is not *too different* from $\bar{\mu}$. One reasonable kind of conformity to demand is to assume that private beliefs be *absolutely continuous* with respect to the analyst's beliefs, which means that private agents will agree with

[4]The policymaker is here assumed not to entertain doubts about the correctness of her own expectations; thus I am here not concerned with the main kind of uncertainty emphasized by Hansen and Sargent (2005c).

the analyst about which outcomes have zero probability.[5] This requires that private-sector beliefs should be of the form

$$\mu = \nu_{\pi, y_1} \tag{1.4}$$

for *some* measure π, even if π is not necessarily the same as $\bar{\pi}$. (In effect, agents are assumed to correctly understand the equilibrium mapping from states of the world to outcomes, even if they do not also correctly assign probabilities to states of the world, as would be required for an RE equilibrium.)

The assumption of absolute continuity also requires that π be absolutely continuous with respect to $\bar{\pi}$. A consequence of this is that there must exist a measurable function $m(\cdot)$, with the property that $\mathrm{E}[m] = 1$, such that for any measurable function $g(\cdot)$ (specifying a random variable at date 1), the expectation $\hat{\mathrm{E}}[g]$ of this random variable under the distorted probability beliefs of the private sector is equal to[6]

$$\hat{\mathrm{E}}[g] = \mathrm{E}[mg].$$

This representation of the distorted beliefs of the private sector is useful in defining a measure of the distance of the private-sector beliefs π from those of the policy analyst, $\bar{\pi}$. As discussed in Hansen and Sargent (2005a, b, c), the *relative entropy*

$$R(\pi, \bar{\pi}) \equiv \mathrm{E}[m \;\log m]$$

is a distance measure with a number of appealing properties.[7] In particular, distorted beliefs π that are not too different from $\bar{\pi}$ in the sense that $R(\pi, \bar{\pi})$ is small are ones

[5]In the dynamic problem treated in the application to monetary stabilization policy below, I actually assume only that private beliefs be absolutely continuous *over finite time intervals*, as in Hansen *et al.* (2005). This means that I allow for misspecifications that should be detected in the case of a data sample of infinite length, as long as they are not easy to detect using a finite data set. As Hansen *et al.* discuss, this is necessary if one wants the policy analyst to be concerned about possible misspecifications that continue to matter far in the future. Absolute continuity over finite time intervals suffices for the representation of distorted beliefs proposed in this section to continue to apply in the dynamic setting.

[6]The existence of the function $m(\cdot)$ is guaranteed by the Radon-Nikodym theorem. In the case of a discrete set of states s, $m(s)$ is simply the ratio $\pi(s)/\bar{\pi}(s)$. This way of describing distorted beliefs is used, for example, by Hansen and Sargent (2005a, b) and Hansen *et al.* (2005).

[7]For example, $R(\pi, \bar{\pi})$ is a positive-valued, convex function of π, uniquely minimized (with the value zero) when $m_{t+1} = 1$ almost surely (the case of RE).

that (according to the beliefs of the analyst) private agents would not be expected to be able to disconfirm by observing the outcome of repeated plays of the game, except in the case of a very large number of repetitions (the number expected to be required being larger the smaller the relative entropy). One might thus view the distorted beliefs π as more plausible the smaller is $R(\pi, \bar{\pi})$.

One way to incorporate a concern on the part of the policy analyst for robustness with regard to this type of uncertainty is to suppose that the analyst wishes to choose a policy $(x_0, x_1(\cdot))$ that is not too bad (does not imply too high a value of (1.2) under *any* equilibrium (solution to (1.1) associated with private-sector beliefs of the form (1.4) for which the relative entropy is not too large. Thus we might assume that the policy is chosen to minimize

$$\bar{L}(x_0, x(\cdot)) = \max \mathrm{E}[L(y_0, y_1, s)], \tag{1.5}$$

where the maximization in (1.5) is over triples $(y_0, y_1(\cdot), \pi)$ such that (1.1) is satisfied when μ is given by (1.4), and such that

$$R(\pi, \bar{\pi}) \leq \bar{R}, \tag{1.6}$$

for some finite bound $\bar{R} > 0$. In this case, the concern for robustness would be modeled in a way analogous to the formalization of ambiguity aversion by Gilboa and Schmeidler (1989).

Alternatively, we can model a concern for robustness in a way analogous to the one that is primarily used by Hansen and Sargent (2005c), who follow the lead of the engineering literature on robust control. Instead of supposing that the "worst-case" near-rational expectations (NRE) contemplated by the analyst are those that maximize (1.2) over a set of possible beliefs defined by the constraint (1.6), we may suppose that the worst-case beliefs (and associated equilibrium outcomes) associated with a given policy are the triple $(\hat{y}_0, \hat{y}_1(\cdot), \hat{\pi})$ that maximize

$$\mathrm{E}[L(y_0, y_1, s)] - \theta R(\pi, \bar{\pi}), \tag{1.7}$$

for some penalty coefficient $\theta > 0$, over all possible triples $(y_0, y_1(\cdot), \pi)$ such that (1.1) is satisfied when μ is given by (1.4). Here *no* constraint such as (1.6) is imposed on the distorted beliefs that may be considered, but beliefs that are less plausible (from the point of view of the analyst) are more heavily penalized in the objective (1.7).

Thus the analyst will only worry about possible distorted private-sector beliefs that ought to be easy to disconfirm in the case that this particular kind of difference in beliefs would be especially problematic for the particular policy under consideration.[8]

This is the definition of worst-case NRE that I shall use here. The policy analyst is assumed to choose a policy $(x_0, x_1(\cdot)$ that minimizes $\hat{L}(x_0, x(\cdot))$, the maximized value of(1.7) under beliefs (1.3), obtained when $(y_0, y_1(\cdot))$ are the worst-case NRE beliefs consistent with the policy $(x_0, x_1(\cdot))$.[9] One can think of this as the Stackelberg equilibrium of a game between the policymaker and a "malevolent agent" who chooses the private-sector beliefs π that will most embarrass the policymaker.[10] Robust policy in this sense approaches the optimal policy commitment under RE in the limit as θ is made unboundedly large, so that the beliefs of the private sector are assumed to be given by $\pi = \bar{\pi}$ regardless of the policy chosen.

The robust policy problem considered here is related to, though not quite the same as, the type of problem considered by Hansen and Sargent (2005c, chap. 16). Hansen and Sargent discuss a class of "Stackelberg problems" in which a "leader" chooses

[8]Maccheroni *et al.* (2004, 2005) show that choosing so as to minimize an objective of the form (1.7) is consistent with a set of axioms for choice under ambiguity aversion only slightly weaker than those of Gilboa and Schmeidler. Both the objective (1.7) and this one are only two members of a broader family that they characterize; the Hansen-Sargent "multiplier preferences" are convenient for my purposes.

[9]Alternatively, one might suppose that the policy analyst is assumed to choose a policy $(x_0, x_1(\cdot)$ that minimizes $L^\dagger(x_0, x(\cdot))$, the value of (1.2) under the same worst-case NRE beliefs. The case assumed in the text corresponds to "variational preferences" of the kind discussed by Maccheroni *et al.* (2004, 2005), and also to the kind of "multiplier robust control problem" treated extensively by Hansen and Sargent (2005c). Apart from the appeal of the axiomatic foundations offered by Maccheroni *et al.* for their representation of preferences, this formulation has the advantage of making the objectives of the policy analyst and the "malevolent agent" perfectly opposed, so that the "policy game" between them is a zero-sum game. This can have advantages when characterizing the solution, though I have not relied on this aspect of the game in the analysis below. The monetary stabilization policy problem is analyzed under the alternative assumption in Woodford (2005), and the same qualitative results are obtained.

[10]Under the assumption made here about the policymaker's objective, the game is zero-sum, and so under certain regularity conditions (that apply in the application below, for example), the Stackelberg equilibrium is also the Nash equilibrium; one could then analyze a "multiplier game" analogous to the one treated in Hansen and Sargent (2005c, chap. 6). Such a change in the timing of moves by the two "players" is not innocuous, instead, under the alternative objective for the policymaker mentioned in the previous footnote.

a policy taking into account not only the optimizing response of the "follower" to the policy, but also the fact that the follower optimizes under distorted beliefs (*i.e.,* distorted from the point of view of the leader), as a result of the follower's concern for possible model misspecification.[11] The problem considered here is similar, except that here the policy analyst is worried about the NRE beliefs that would be *worst for her own objectives,* while in the Hansen-Sargent game, the leader anticipates that the follower will act on the basis of the distorted beliefs that would imply *the worst outcome for the follower himself.*[12]

One might think that this difference should not matter in practice, if the policymaker's objective coincides with that of the "follower" — as one might think should be the case in an analysis of optimal policy from the standpoint of public welfare. But in the application to stabilization policy below, the private sector is not really a single agent, even though I assume that all price-setters share the same distorted beliefs. It is not clear that allowing for a concern for robustness on the part of individual price-setters would lead to their each optimizing in response to common distorted beliefs, that coincide with those beliefs under which *average* expected utility is lowest.

But more crucially, even in a case where the private sector is made up of identical agents who each solve precisely the same problem, the distorted beliefs that matter in the Hansen-Sargent analysis are those that result in an equilibrium $(y_0, y_1(\cdot))$ with the highest possible value of $\hat{\mathrm{E}}[L(y_0, y_1, s)]$, *i.e.,* the greatest *subjective losses* from the point of view of the private sector. In the problem considered here, instead, the NRE beliefs that matter are those that result in an equilibrium with the highest possible value of $\mathrm{E}[L(y_0, y_1, s)]$; even if the loss function is identical for the policymaker and the private sector, I assume that it is the *policymaker's* evaluation of expected losses that matters for robust policy analysis.

In the case that the objective of public policy is assumed to be private welfare, this choice might not be considered obvious; there is always some ambiguity about what

[11]Hansen and Sargent also allow for a concern with potential misspecification on the part of the leader, but in the limiting case of their setup in which $\Theta = \infty$ while $\theta < \infty$, only the follower contemplates that the common "approximating model" may be incorrect; the leader regards it as correct, but takes account of the effect on the follower's behavior of the follower's concern that the model may be incorrect.

[12]I also consider a different class of possible distorted probability beliefs (Hansen and Sargent allow only for shifts in the mean of the conditional distribution of possible values for the disturbances) and use a different measure of the degree of distortion of PS beliefs (relative entropy).

it should mean for policy to be welfare-maximizing in the case that private agents are regarded by the policy analyst as being mistaken about their situation. Here I take the view that the policy analyst (that in this paper, at least, has no doubts about the correct beliefs) should evaluate private welfare from the point of view of what she believes to be the true likelihood of alternative outcomes. One might also consider the alternative assumption, and define robustly optimal policy as the policy that minimizes $\hat{\mathrm{E}}[L(y_0, y_1, s)]$. In the application considered next, this alternative assumption would lead to a much more trivial problem: the robustly optimal policy commitment would be exactly like an optimal policy commitment under RE, if the worst-case NRE beliefs were treated as true.[13] Here I consider instead the harder problem of how to choose a robustly optimal policy from the point of view of the policy analyst's own probability beliefs.

2 An Application to Monetary Stabilization Policy

The example considered here weakens the assumption regarding private-sector expectations in the well-known analysis by Clarida *et al.* (1999) of optimal monetary policy in response to "cost-push shocks." It is assumed that the central bank can bring about any desired state-contingent evolution of inflation π_t and of the output gap x_t consistent with the aggregate-supply relation

$$\pi_t = \kappa x_t + \beta \hat{E}_t \pi_{t+1} + u_t, \tag{2.1}$$

where $\kappa > 0, 0 < \beta < 1$, $\hat{E}_t[\cdot]$ denotes the common (distorted) expectations of the private sector (more specifically, of price-setters — I shall call these *PS* expectations) conditional on the state of the world in period t, and u_t is an exogenous cost-push

[13]This would mean, for example, that the optimal policy commitment could be implemented through commitment to perfectly stabilize a certain linear combination of the log price level and the output gap, as discussed in Woodford (2003, chap. 7). The quantitative form of the optimal target criterion would be completely unaffected by the central bank's degree of concern for possible forecast error on the part of the private sector. The possibility of NRE beliefs would only have to be taken into account when implementing policy; for example, when evaluating the short-run tradeoff between inflation and the output gap, in order to produce an outcome that satisfies the target criterion.

shock. The analysis is here simplified by assuming that all PS agents have common expectations (though these may not be model-consistent); given this, the usual derivation[14] of (2.1) as a log-linear approximation to an equilibrium relation implied by optimizing price-setting behavior follows just as under the assumption of RE.

The central bank's (CB) policy objective is minimization of a discounted loss function

$$E_0 \sum_{t=0} \beta^t \frac{1}{2}[\pi_t^2 + \lambda(x_t - x^*)^2] \tag{2.2}$$

where $\lambda > 0, \quad x^* \geq 0$, and the discount factor β is the same as in (2.1). Here $E_t[\cdot]$ denotes the conditional expectation of a variable under the CB beliefs, which I shall treat as the "true" probabilities, since the analysis is conducted from the point of view of the CB, which wishes to consider the effects of alternative possible policies. I do not allow for any uncertainty on the part of the CB about the probability with which various "objective" states of the world (histories of exogenous disturbances) occur, in order to focus on the issue of uncertainty about PS expectations. The CB believes that the exogenous states s_t evolve according to a law of motion

$$s_{t+1} = As_t + Bw_{t+1} \tag{2.3}$$

for some matrices A, B, where the random vector w_{t+1} is i.i.d. with distribution $N(0, I)$; the cost-push shock each period is then given by $u_t = v's_t$. Thus the vector s_t describes all information available at time t about current or future "fundamentals". Note that the law of motion (2.3) is *not* assumed to be correctly understood by the PS.

I shall suppose that the central bank chooses (once and for all, at some initial date) a state-contingent policy $\pi_t = \pi(h_t)$, where $h_t \equiv (w_t, w_{t-1}, \ldots)$ is the history of realizations of the exogenous disturbances. I assume that commitment of this kind is possible, to the extent that it proves to be desirable; and we shall see that it *is* desirable to commit in advance to a policy different from the one that would be chosen *ex post*, once any effects of one's decision on prior inflation expectations could be neglected. I also assume that there is no problem for the central bank in *implementing* the state-contingent inflation rate that it has chosen, once a given situation h_t is reached.[15] This is likely to require that someone in the central bank

[14] See, *e.g.*, Woodford (2003, chap. 3).

[15] Even so, the assumption that the central bank commits itself to a state-contingent path for

can observe exactly what PS inflation expectations are at the time of implementation of the policy (in order to determine the nominal interest rate required to bring about a certain rate of inflation); I assume uncertainty about PS expectations only at the time of selection of the state-contingent policy commitment. Note that any such strategy $\pi(\cdot)$ implies a uniquely defined state-contingent evolution of both inflation and the output gap (given PS beliefs), using equation (2.1), and thus a well-defined value for CB expected losses (2.2).

As in section 1, I shall assume that NRE require that PS beliefs about the economy's evolution over any finite horizon (and in particular, PS beliefs about the probability of various states in the following period) be absolutely continuous with respect to those of the CB. Hence there exists a process $\{m_{t+1}\}$ with

$$m_{t+1} \geq 0 \quad \text{a.s.}, \qquad E_t[m_{t+1}] = 1,$$

such that

$$\hat{E}_t[X_{t+1}] = E_t[m_{t+1}X_{t+1}]$$

for any random variable X_{t+1}. The degree of distortion of PS beliefs can furthermore be measured by the (discounted) relative entropy

$$E_0 \sum_{t=0}^{\infty} \beta^t m_{t+1} \log m_{t+1},$$

as in Hansen and Sargent (2005a). The presence of the discount factor β^t in this expression means that the CB's concern with potential PS misunderstanding doesn't vanish asymptotically; this makes possible a time-invariant characterization of robustly optimal policy in which the concern for robustness has nontrivial consequences.[16]

Consequently, in the case of any policy commitment $\{\pi_t\}$ contemplated by the CB, the "worst-case" NRE beliefs considered by the CB are given by the process $\{m_{t+1}\}$ that maximizes

$$E_0 \sum_{t=0} \beta^t \frac{1}{2}[\pi_t^2 + \lambda(x_t - x^*)^2] - \theta E_0 \sum_{t=0}^{\infty} \beta^t m_{t+1} \log m_{t+1} \tag{2.4}$$

inflation, rather than to a Taylor rule or to the satisfaction of some other form of target criterion, is not innocuous. Using this representation of the policy commitment would be innocuous in a RE analysis like that of Clarida *et al.* (1999), since one is effectively choosing from among all possible REE. But here different representations of policy do not lead to the same set of equilibrium allocations being consistent with near-rational expectations.

[16]See Hansen *et al.* (2005) for discussion of this issue, in the context of a continuous-time analysis.

subject to the constraint that $E_t m_{t+1} = 1$ at all times, where at each date x_t is implied by the equation

$$\pi_t = \kappa x_t + \beta E_t[m_{t+1}\pi_{t+1}] + u_t. \tag{2.5}$$

Here $\theta > 0$ is a multiplier that indexes the degree of concern for robustness of policy with respect to non-RE beliefs.

This problem for the "malevolent agent" is in turn equivalent to a sequence of problems in which for each possible history h_t, a function specifying m_{t+1} as a function of the realization of w_{t+1} is chosen so as to maximize

$$\frac{1}{2}[\pi_t^2 + \lambda(x_t - x^*)^2] - \theta E_t[m_{t+1}\log m_{t+1}] \tag{2.6}$$

subject to the constraint that $E_t m_{t+1} = 1$, where again x_t is implied by (2.5). Worst-case NRE then determine the expected output effect of any given state-dependent inflation commitment, according to a time-invariant relation of the form

$$x_t = x^{pess}(u_t, \pi_t, \pi_{t+1}(\cdot)), \tag{2.7}$$

where $\pi_{t+1}(\cdot)$ specifies π_{t+1} as a measurable function of w_{t+1}. The degree of distortion of PS beliefs under the worst-case NRE is similarly indicated by a time-invariant function

$$E_t[m_{t+1}\log m_{t+1}] = R^{pess}(u_t, \pi_t, \pi_{t+1}(\cdot)) \tag{2.8}$$

indicating the relative entropy of the worst-case PS beliefs. A robustly optimal policy commitment by the CB is then one that minimizes the maximized value of (2.4), which is to say, that minimizes the objective function obtained by substituting (2.7) for the output gap and (2.8) for the relative-entropy term in (2.4).

This problem can be given a recursive structure if we add an additional constraint, assuming that the initial inflation commitment $\pi_0(w_0)$ is exogenously given.[17] Let $J(\pi_0; s_0)$ be the min-max value of (2.4), conditional on a particular initial state. Then under robustly optimal policy, each period the function $\pi_{t+1}(\cdot)$ is chosen, given the prior inflation commitment π_t and the state s_t, so as to minimize

$$\begin{aligned}\frac{1}{2}\pi_t^2 + \frac{\lambda}{2}(x^{pess}(u_t, \pi_t, \pi_{t+1}(\cdot)) - x^*)^2 \quad &- \quad \theta R^{pess}(u_t, \pi_t, \pi_{t+1}(\cdot)) \\ &+ \quad \beta E[J(\pi_{t+1}(w); As_t + Bw)],\end{aligned} \tag{2.9}$$

[17] The same kind of initial commitment defines an optimal policy "from a timeless perspective" in the RE analysis presented in Woodford (2003, chap. 7).

where the expectation is over possible realizations of w. The minimized value of (2.9) is then precisely the value function $J(\pi_t; s_t)$. This constrained version of the robust policy problem is of interest because (as a result of its recursive form) it results in a time-invariant rule for robustly optimal policy.

This recursive structure implies that there is no need for the CB to commit itself more than a period in advance. However, it *is* important that the state-contingent inflation commitments be chosen at least a period in advance, rather than waiting until the state s_{t+1} is known and then choosing π_{t+1} to minimize $J(\pi_{t+1}; s_{t+1})$. The latter (purely discretionary) approach to policy will not achieve as low a value of (2.9), and hence not as low a value of (2.4) under the worst-case NRE beliefs, as will the approach of choosing a state-contingent inflation commitment each period for the following period. The reason for this advantage of policy commitment is exactly the same, of course, as in the RE analysis of optimal policy in this model (treated in Clarida *et al.*, 1999, and Woodford, 2003, chap. 7).

3 Robustly Optimal Linear Policy

Rather than seeking to characterize fully optimal policy in the sense defined above, I shall here characterize the optimal policy within a more restrictive class of *linear policies*. By a linear policy I mean one in which each period's state-contingent inflation commitment is of the form

$$\pi_{t+1}(w_{t+1}) = p_t^0 + p_t^{1\prime} w_{t+1},$$

where p_t^0 is some function of h_t and p_t^1 depends only on t.[18] The optimal policy commitment under RE is linear in this sense; hence a consideration of this special family of policies suffices to indicate a direction in which it is desirable to change the CB's policy commitment as a result of concern for robustness.

We begin by characterizing the worst-case NRE in the case of an arbitrary linear policy. One notes that an interior solution to the problem of maximizing (2.6) exists

[18]It will turn out that in the case of the optimal linear policy, p_t^0 is also a *linear* function of h_t, but one does not need to impose that.

only if[19]

$$|p_t^1|^2 < \frac{\theta}{\beta^2}\frac{\kappa^2}{\lambda}. \tag{3.10}$$

Otherwise, the objective (2.6) is *convex,* and the worst-case expectations involve extreme distortion, resulting in unbounded losses for the CB. Obviously, it is optimal for the CB to choose a linear policy such that p_t^1 satisfies the bound (3.10) at all times. This provides an immediate contrast with optimal policy under RE, where the optimal vector p^1 (which is constant over time) is proportional to σ_u, the standard deviation of the cost-push shocks.[20] At least for large values of σ_u, it is evident that concern for robustness leads to *less sensitivity* of inflation to cost-push disturbances (smaller $|p_t^1|$). One also observes that it leads to a failure of *certainty equivalence,* as this would require $|p_t^1|$ to grow in proportion to σ_u.

In the case of a linear policy satisfying (3.10), under the worst-case NRE, the CB fears that the PS will expect w_{t+1} to be conditionally distributed as $N(\mu_t, I)$. If $p_t^1 = 0$, $\mu_t = 0$, while if $p_t^1 \neq 0$,

$$\mu_t = (\bar{\pi}_t - p_t^0)\,\frac{p_t^1}{|p_t^1|^2}, \tag{3.11}$$

where the worst-case inflation expectation (value of $\hat{\mathrm{E}}_t\pi_{t+1}$) is given by

$$\bar{\pi}_t = \Delta_t^{-1}\left[p_t^0 - (\pi_t - u_t - \kappa x^*)\frac{\beta\lambda}{\theta\kappa^2}|p_t^1|^2\right], \tag{3.12}$$

$$\Delta_t \equiv 1 - \frac{\beta^2}{\theta}\frac{\lambda}{\kappa^2}|p_t^1|^2 > 0. \tag{3.13}$$

The worst-case NRE beliefs distort PS inflation expectations with respect to p_t^0 (the CB's expectation) in the direction opposite to that needed to bring x_t closer to x^*; and this distortion is greater the larger is the sensitivity of (next period's) inflation to unexpected shocks, becoming unboundedly large as the bound (3.10) is approached. As a consequence of this possibility, the CB fears an output gap equal to

$$x_t^{pess} - x^* = \frac{(\pi_t - u_t - \kappa x^*) - \beta p_t^0}{\kappa\Delta_t}. \tag{3.14}$$

[19]See the Appendix, section A.1, for derivation of this condition, as well as the results stated in the following two paragraphs. Strictly speaking, it is possible for the inequality (3.10) to be only weakly satisfied, if p_t^0 satisfies a certain linear relation stated in the Appendix; the Appendix treats this case as well, omitted here for simplicity. It is shown in section A.2 that in the robustly optimal linear policy, the inequality is strict.

[20]See, *e.g.*, equation (3.20) below.

Note that $x_t - x^*$ is larger than it would be under RE by a factor Δ_t^{-1}, which exceeds 1 except in the limit in which θ is unboundedly large (the RE limit), or if $p_t^1 = 0$, so that inflation is perfectly predictable.

The probabilities assigned by the PS to different possible realizations of w_{t+1} are distorted by a factor m_{t+1} such that

$$\log m_{t+1} = c_t - \frac{\beta}{\theta}\frac{\lambda}{\kappa}(x_t - x^*)\pi_{t+1},$$

where the constant c_t takes the value necessary in order for $E_t m_{t+1}$ to equal 1. This implies that the degree of distortion of the worst-case NRE beliefs (as measured by relative entropy) is equal to

$$R_t^{pess} \equiv \hat{E}_t[\log m_{t+1}] = \frac{1}{2}\left[\frac{\beta}{\theta}\frac{\lambda}{\kappa}(x_t - x^*)\right]^2 |p_t^1|^2 \geq 0. \tag{3.15}$$

Note that the degree of distortion against which the policy analyst must guard is greater the larger the degree of inefficiency of the output gap (*i.e.*, the larger is $|x_t - x^*|$), as this increases the marginal cost to the CB's objectives of (maliciously chosen) forecast errors of a given size; and greater the larger the degree to which inflation is sensitive to disturbances (*i.e.*, the larger is $|p_t^1|$), as this increases the scope for misunderstanding of the probability distribution of possible future rates of inflation, for a given degree of discrepancy between CB and PS beliefs (as measured by relative entropy). Of course, it is also greater the smaller is θ, the penalty parameter that we use to index the CB's degree of concern for robustness to PS expectational error.

Substituting (3.14) for the output gap and (3.15) for the relative entropy term in (2.4), we obtain a loss function for the CB of the form

$$E_0 \sum_{t=0}^{\infty} \beta^t L(\pi_t; p_t; s_t), \tag{3.16}$$

where $p_t \equiv (p_t^0, p_t^1)$ and

$$\pi_t = p_{t-1}^0 + p_{t-1}^{1\prime} w_t. \tag{3.17}$$

Expression (3.16) indicates the CB's expected losses from a given linear policy $\{p_t\}$, under the worst-case NRE beliefs. We wish to minimize this subject to an initial constraint π_0. Moreover, because we do not wish to allow p_t^1 to vary in response to random shocks, we actually minimize the unconditional expectation of (3.16), integrating over alternative possible initial conditions (p^0_{-1}, s_{-1}) and over alternative possible realizations of w_0.

A *robustly optimal linear policy* (from a timeless perspective) is then a pair of sequences $\{p^0_t, p^1_t\}$ that minimize the expected value of (3.16) subject to the law of motion (3.17), given an initial commitment $p^1_{-1} = \bar{p}^1$ and integrating over initial conditions (p^0_{-1}, s_{-1}) using a measure ρ. The value of p^0_t is allowed to depend on the history h_t, as well as the particular initial conditions that are drawn from the support of ρ, but a value for p^1_t must be chosen that is independent of shock realizations and the same for all initial conditions (which is why the measure ρ matters). The initial constraints $(\bar{p}^1, \rho)$ are chosen to be *self-consistent*,[21] which means that under the optimal policy, $p^1_t = \bar{p}^1$ for all $t \geq 0$, and ρ is an invariant measure for (p^0_t, s_t). One can show that values of $(\bar{p}^1, \rho)$ exist with this property.

Given $p^1_t = \bar{p}^1$, the loss function $L(\pi_t; p_t; s_t)$ is a quadratic function of (π_t, p^0_t, s_t), and the laws of motion (2.3) and (3.17) are linear in these variables. Hence one has a linear-quadratic optimal control problem, and the optimal solution is a linear policy of the form

$$p^0_t = \mu p^0_{t-1} + a' s_t + \mu \bar{p}^{1\prime} w_t, \qquad (3.18)$$

just like the unconstrained optimal policy under RE. A concern for robustness affects the numerical magnitudes of μ, a, and $\bar{p}_1$. But one thing that is not affected is the fact that (3.18) implies stationary fluctuations in the inflation rate around a long-run inflation target of zero. Thus the optimal long-run target is unaffected by the degree of concern for robustness; in particular, allowance for NRE does *not* result in an inflation bias of the kind associated with discretionary policy.[22]

Here I illustrate the quantitative effects of a concern for robustness in an example in which the cost-push shock is purely transitory, so that w_t is a scalar and $u_t = \sigma_u w_t$. Under RE,

$$0 < \mu < 1, \qquad a' s_t = -\mu \sigma_u w_t, \qquad (3.19)$$

and

$$\bar{p}^1 = \mu \sigma_u. \qquad (3.20)$$

With a concern for robustness (finite θ), conditions (3.19) both still hold, but μ is

[21] See Woodford (2003, chap. 7) for the concept of self-consistency invoked here.

[22] On the inflation bias associated with discretionary policy, see Clarida *et al.* (1999) or Woodford (2003, chap. 7).

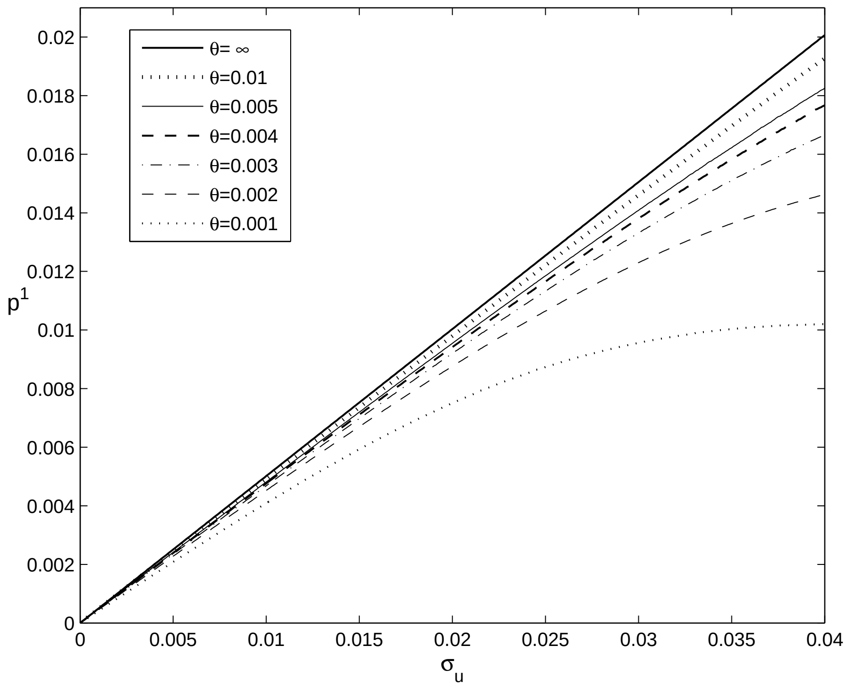

Figure 1: Variation of $\bar{p}^1$ with σ_u, under alternative degrees of concern for robustness.

now the smaller root of the quadratic equation

$$P(\mu) \equiv \beta\mu^2 - \left(1 + \beta + \frac{\kappa^2\bar{\Delta}}{\lambda}\right)\mu + 1 = 0, \tag{3.21}$$

where $0 < \bar{\Delta} \leq 1$ is the constant value of (3.13) associated with $\bar{p}^1$. It is evident from (3.21) that μ is larger the smaller is $\bar{\Delta}$; and since a concern for robustness lowers $\bar{\Delta}$, it raises μ relative to the RE case. Moreover, contrary to (3.20), one can show that

$$\bar{p}_1 < \mu\sigma_u \tag{3.22}$$

when θ is finite.

Figure 1 shows how $\bar{p}^1$ varies with σ_u for alternative values of θ.[23] In the RE case, $\bar{p}^1$ increases linearly with σ_u, as indicated by (3.20) and as required for certainty-equivalence. For any given amplitude of cost-push shocks, lower θ (greater concern

[23] In this figure, I assume parameter values $\beta = 0.99, \kappa = 0.05, \lambda = 0.08$, and $x^* = 0.2$. A low value of λ is justified by the welfare-theoretic foundations of the loss function (2.2) discussed in Woodford (2003, chap. 6).

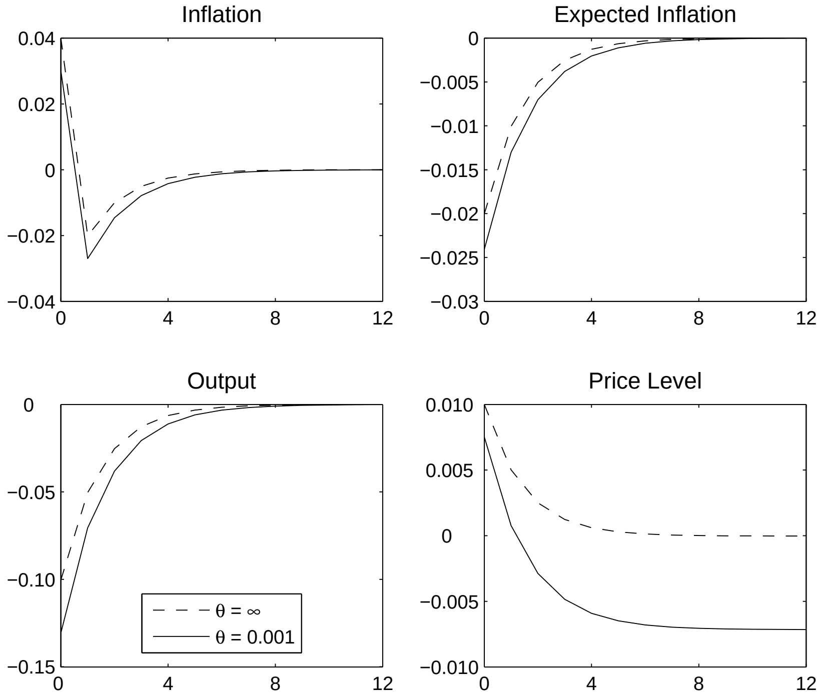

Figure 2: Optimal responses to a positive cost-push shock, with and without concern for robustness.

for robustness) results in a lower optimal $\bar{p}_1$, indicating less sensitivity of inflation to the current cost-push shock. The extent to which this is true increases in the case of larger shocks; in the case of any finite value of θ, $\bar{p}_1$ increases less than proportionally with σ_u, indicating a failure of certainty equivalence. In fact, $\bar{p}_1$ remains bounded above, as required by (3.10).

Thus a concern for robustness results in less willingness to let inflation increase in response to a positive cost-push shock. This is because larger surprise variations in inflation increase the extent to which PS agents may over-forecast inflation, worsening the output/inflation tradeoff facing the CB. This conclusion recalls the one reached by Orphanides and Williams (2005a) on the basis of a model of learning.

At the same time, a concern for robustness increases the degree to which optimal policy is history-dependent. As in the RE case, an optimal commitment involves a lower inflation rate (on average) in periods *subsequent* to a positive cost-push shock. Moreover, because μ is closer to 1 when θ is smaller, this effect on average inflation should *last longer,* so that the history-dependence of the optimal inflation commitment is even greater than under RE. And not only should the CB commit to eventually undo any price increases resulting from positive cost-push shocks (as in the RE case); when θ is finite, it should commit to eventually reduce the price level *below* the level it would have had in the absence of the shock. This is illustrated in Figure 2 in the case of the numerical example just discussed.[24] The lower right panel shows the impulse response of the log price level; while under rational expectations, the optimal commitment returns the price level eventually to precisely the level that it would have had in the absence of the shock, when $\theta = 0.001$, the optimal commitment eventually *reduces* the price level, by an amount about twice as large as the initial price-level increase in response to the shock. The result that the sign of the initial price-level effect is eventually reversed is quite general. Equations (3.17) – (3.19) imply that the cumulative log price increase due to a one-standard-deviation cost-push shock is equal to

$$(\bar{p}^1 - \mu\sigma_u)/(1 - \mu),$$

which is zero when (3.20) holds, but negative when θ is finite.

Allowance for NRE means that the CB cannot count on its intention to lower inflation (on average) following a positive cost-push shock to lower PS expectations of inflation by as much as the CB's own forecast of future inflation is reduced. (For example, Figure 3 compares the impulse response of PS expected inflation $\hat{\mathrm{E}}_t\pi_{t+1}$ to the response of CB expected inflation $\mathrm{E}_t\pi_{t+1}$, in the same numerical example as in Figure 2.) But the consequence of this for robustly optimal policy is not that the CB should not bother to try to influence inflation expectations through a history-dependent policy; instead, it is optimal to commit to adjust the subsequent inflation target to an even greater extent and in a more persistent, in order to ensure that inflation expectations are affected even if expectations are not perfectly model-consistent.

[24]In the figure, optimal impulse responses to a one-standard-deviation positive cost-push shock are shown, both in the case of infinite θ (the standard RE analysis) and for a value $\theta = 0.001$. Other parameter values are as in Figure 1; in addition, it is assumed here that $\sigma_u = 0.02$.

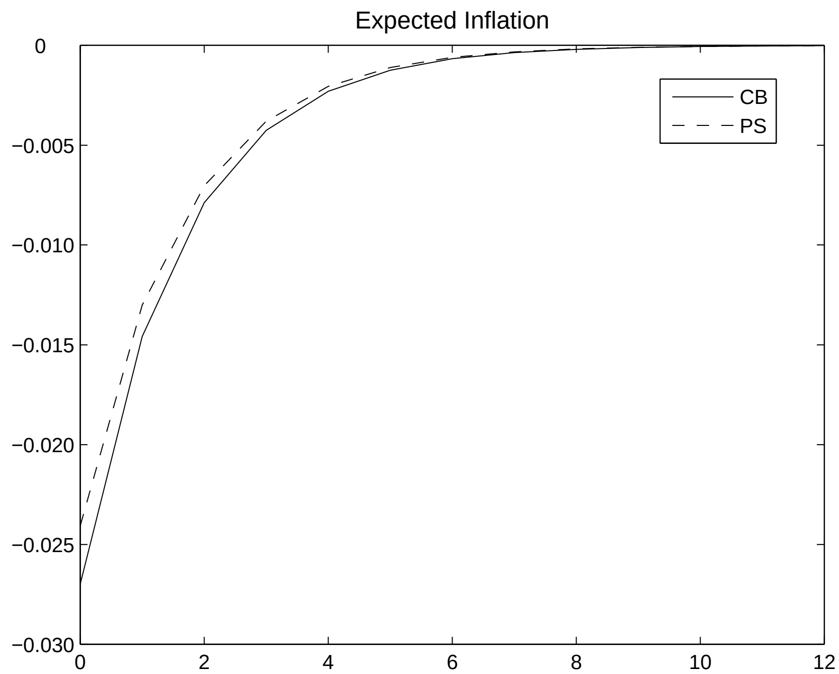

Figure 3: Distortion of PS beliefs, in the worst-case NRE contemplated by the CB when $\theta = 0.001$.

4 Conclusion

I have shown how it is possible to analyze optimal policy for a central bank that recognizes that private expectations may not be model-consistent, without committing oneself to a *particular* model of expectational error. The approach leads to a one-parameter family robustly-optimal linear policies, indexed by a parameter θ that measures the degree of concern for possible misunderstanding of equilibrium dynamics.

Even when the central bank's uncertainty about private expectations is considerable (the case of low θ), calculation of the effects of *anticipations* of the systematic component of policy is still quite an important factor in policy analysis. As in RE analysis, *commitment* is still important, and optimal policy is still *history-dependent* — indeed, a concern for robustness only increases the optimal degree of history-dependence. And, as in the RE analysis, a crucial feature of an optimal commitment is a guarantee that inflation will be low and fairly stable. The fact that private beliefs may be distorted does not provide any reason to aim for a higher average rate of inflation, while it does provide a reason for the central bank to resist even more firmly the inflationary consequences of "cost-push" shocks.

A Appendix: Details of Derivations

A.1 Worst-Case NRE Beliefs

The problem of the "malevolent agent" in any state of the world at date t (corresponding to a history h_t up to that point) is to choose a function specifying m_{t+1} as a function of the realization of w_{t+1} so as to maximize

$$\frac{1}{2}[\pi_t^2 + \lambda(x_t - x^*)^2] - \theta E_t[m_{t+1} \log m_{t+1}] \tag{A.23}$$

subject to the constraint that $E_t m_{t+1} = 1$, where at each date x_t is implied by the equation

$$\pi_t = \kappa x_t + \beta E_t[m_{t+1}\pi_{t+1}] + u_t. \tag{A.24}$$

Here I characterize the solution to this problem in the case that the CB follows a linear policy, so that in each possible state at date $t+1$ (given the history h_t), the inflation rate is given by $\pi_{t+1} = p_t^0 + p_t^{1\prime} w_{t+1}$, where p_t^0 depends only on h_t and p_t^1 depends only on the date t.

It is obvious that the choice of the random variable m_{t+1} matters only through its consequences for the relative entropy (which affects the objective (A.23)) on the one hand, and its consequences for PS expected inflation (which affects the constraint (A.24) on the other. Hence in the case of any $\theta > 0$, the worst-case beliefs will minimize the relative entropy $E_t[m_{t+1} \log m_{t+1}]$ subject to the constraints that

$$E_t m_{t+1} = 1, \qquad E_t[m_{t+1}\pi_{t+1}] = \bar{\pi}_t, \tag{A.25}$$

whatever degree of distortion the PS inflation expectation $\bar{\pi}_t$ may represent. I first consider this problem. Since $r(m) \equiv m \log m$ is a strictly convex function of m, such that $r'(m) \to -\infty$ as $m \to 0$ and $r'(m) \to +\infty$ as $m \to +\infty$, it is evident that there is a unique, interior optimum, such that the first-order condition

$$v'(m_{t+1}) = \phi_{1t} + \phi_{2t}\pi_{t+1}$$

holds in each state at date $t+1$, where ϕ_{1t}, ϕ_{2t} are Lagrange multipliers associated with the two constraints (A.25). This implies that

$$\log m_{t+1} = c_t + \phi_{2t}\pi_{t+1} \tag{A.26}$$

in each state, for some constant c_t. The two constants c_t and ϕ_{2t} in (A.26) are then the values that satisfy the two constraints (A.25).

Under the assumption of a linear policy, π_{t+1} is conditionally normally distributed, so that (A.26) implies that m_{t+1} is conditionally log-normal.[25] It follows that

$$\begin{aligned}\log E_t m_{t+1} &= E_t[\log m_{t+1}] + \frac{1}{2}\mathrm{var}_t[\log m_{t+1}] \\ &= c_t + \phi_{2t} p_t^0 + \frac{1}{2}\phi_{2t}^2 |p_t^1|^2.\end{aligned}$$

Hence the first constraint (A.25) is satisfied if and only if

$$c_t = -\phi_{2t} p_t^0 - \frac{1}{2}\phi_{2t}^2 |p_t^1|^2. \tag{A.27}$$

Under the worst-case beliefs, the PS perceives the conditional probability density for w_{t+1} to be $\tilde{f}(w_{t+1}) = m_{t+1}(w_{t+1}) f(w_{t+1})$, where $f(\cdot)$ is the density for a vector that is distributed as $N(0, I)$. It follows from (A.26) and (A.27) that $\tilde{f}(\cdot)$ is the density for a vector that is distributed as $N(\mu_t, I)$, where the bias in the perceived conditional expectation of w_{t+1} is $\mu_t = \phi_{2t} p_t^1$. Hence

$$\hat{E}_t \pi_{t+1} = p_t^0 + p_t^{1\prime} \mu_t = p_t^0 + \phi_{2t} |p_t^1|^2,$$

and the second constraint (A.25) is satisfied if and only if[26]

$$\phi_{2t} = \frac{\bar{\pi}_t - p_t^0}{|p_t^1|^2}. \tag{A.28}$$

Condition (A.27) then uniquely determines c_t as well, and m_{t+1} is completely described by (A.26), once we have determined the value of $\bar{\pi}_t$ that should be chosen by the "malevolent agent." Note that the bias μ_t is given by expression (3.11), as asserted in the text.

The relative entropy of the worst-case beliefs will then be equal to

$$\begin{aligned}R_t^{pess} = \hat{E}_t[\log m_{t+1}] &= c_t + \phi_{2t} \hat{E}_t \pi_{t+1} \\ &= \frac{1}{2}\frac{(\bar{\pi}_t - p_t^0)^2}{|p_t^1|^2},\end{aligned} \tag{A.29}$$

[25]This is one of the main reasons for the convenience of restricting our attention to linear policies.

[26]Here I assume that $p_t^1 \neq 0$. If $p_t^1 = 0$, the constraint is satisfied regardless of the distortion chosen by the "malevolent agent," as long as $\bar{\pi}_t = p_t^0$, which is necessarily the case. In this case, c_t and ϕ_{2t} are not separately identified, but (A.27) suffices to show that $m_{t+1} = 1$ with certainty.

using (A.27) and (A.28). This is proportional to the squared distance between the PS inflation forecast and that of the central bank; but for any given size of gap between the two, the size of the distortion of probabilities that is required is smaller the larger is $|p_t^1|$.[27]

It remains to determine the worst-case choice of $\bar{\pi}_t$.[28] It follows from (A.24) that

$$(x_t^{pess} - x^*)^2 = \frac{1}{\kappa^2}(\pi_t - u_t - \kappa x^* - \beta\bar{\pi}_t)^2. \qquad \text{(A.30)}$$

Substituting this for the squared output gap and (A.29) for the relative entropy in (A.23), we obtain an objective for the "malevolent agent" that is a quadratic function $Q(\bar{\pi}_t; u_t, \pi_t, p_t)$ of the distorted inflation forecast $\bar{\pi}_t$, and otherwise independent of the distorted beliefs; thus $\bar{\pi}_t$ is chosen to maximize this function. The function is strictly concave (because the coefficient multiplying $\bar{\pi}_t^2$ is negative) if and only if p_t^1 satisfies inequality (3.10). If the inequality is reversed, the function Q is instead *convex*, and is minimized rather than maximized at the value of $\bar{\pi}_t$ that satisfies the first-order condition $Q_{\bar{\pi}} = 0$. But in this case, the "malevolent agent" can achieve an unboundedly large positive value of the objective (A.23), as stated in the text; and a robustly optimal policy can never involve a value of p_t^1 this large.

In the case that (3.10) holds with equality, Q is linear in $\bar{\pi}_t$, and it is again possible for the "malevolent agent" to achieve an unboundedly large positive value of the objective through an extreme choice of $\bar{\pi}_t$, except in the special case that

$$p_t^0 = \beta^{-1}(\pi_t - u_t - \kappa x^*), \qquad \text{(A.31)}$$

so that the linear function has a slope of exactly zero. Thus unless p_t^0 satisfies (A.31), p_t^1 must satisfy the bound (3.10) in order for the objective (A.23) to have a finite maximum. Even in the special case that (A.31) holds exactly, p_t^1 must satisfy a variant of (3.10) in which the strict inequality is replaced by a weak inequality.

When (3.10) holds, the maximum value of Q occurs for the value of $\bar{\pi}_t$ such that $Q_{\bar{\pi}} = 0$. This implies that the worst-case value of $\bar{\pi}_t$ is the one given by (3.12) – (3.13) in the text. Substituting this solution into (A.29) and (A.30), one obtains the implied output gap (3.14) and and relative entropy (3.15) under the worst-case NRE

[27] Equation (A.29) again assumes that $p_t^1 \neq 0$. In the event that $p_t^1 = 0$, it follows from the previous footnote that the relative entropy of the worst-case beliefs will equal zero.

[28] The analysis here assumes that $p_t^1 \neq 0$. If $p_t^1 = 0$, there is no choice about the value of $\bar{\pi}_t$; it must equal p_t^0.

beliefs, as stated in the text. Substituting these expressions into the objective (A.23), one obtains an objective for the CB of the form (3.16), in which the period loss is given by

$$L(\pi_t; p_t; s_t) = \frac{1}{2}\pi_t^2 + \frac{\lambda}{2\kappa^2 \Delta_t}[\pi_t - u_t - \kappa x^* - \beta p_t^0]^2,$$

where Δ_t is the function of p_t^1 defined by (3.13).

When, instead, (3.10) holds with equality, and (A.31) holds as well, the worst-case value of $\bar{\pi}_t$ is indeterminate, but the maximized value of (A.23) is nonetheless well-defined, and equal to zero. In this case, the period loss function is equal to

$$L(\pi_t; p_t; s_t) = \frac{1}{2}\pi_t^2.$$

When neither this case nor the one discussed in the previous paragraph applies, we can define $L(\pi_t; p_t; s_t)$ as being equal to $+\infty$. The function is then defined (but possibly equal to $+\infty$) for all possible values of its arguments.

Note also that $L(\pi_t; p_t; s_t)$ is necessarily non-negative, since for any values of the arguments, it is possible for the "malevolent agent" to obtain a non-negative value of (A.23) by choosing $m_{t+1} = 1$ in all states; the maximized value of (A.23) is then necessarily at least this high. It follows that both the conditional expectations and the infinite sum in (3.16) are sums (or integrals) of non-negative quantities; hence both are well-defined (though possibly equal to $+\infty$) for all possible values of the arguments. Thus the CB objective (3.16) is well-defined for arbitrary state-contingent sequences $\{p_t\}$ and an arbitrary initial condition (π_0, s_0).

A.2 Robustly Optimal Linear Policy

Given the worst-case PS beliefs characterized in the previous section, the problem of the CB is to choose a sequence $\{p_t\}$ for all $t \geq 0$ so as to minimize

$$E_\rho E_0 \sum_{t=0}^{\infty} \beta^t L(\pi_t; p_t; s_t), \qquad \text{(A.32)}$$

where

$$\pi_{t+1} = p_t^0 + p_t^{1\prime} w_{t+1} \qquad \text{(A.33)}$$

and

$$s_{t+1} = A s_t + B w_{t+1} \qquad \text{(A.34)}$$

for each $t \geq 0$, and (π_0, s_0) are given as initial conditions. Here E_ρ indicates an integral over alternative possible initial conditions (p^0_{-1}, s_{-1}, w_0) using a measure ρ, the choice of which is explained in the next section; and it is assumed that $p^1_{-1} = \bar{p}^1$, where the choice of $\bar{p}^1$ (a single value) is also explained in the next section. I use the notation $E_t[\cdot]$ to indicate an expectation conditional upon the history h_t, by which I mean the particular initial conditions (p^0_{-1}, s_{-1}, w_0) that have been drawn, together with the subsequent realizations of the exogenous disturbances $(w_1, \ldots, w_t)$. I furthermore suppose that the CB's choice of p^1_t must depend only on the date t, while the choice of p^0_t may depend on the history h_t.

One can show that the objective (A.32) is a convex function of the sequence $\{p_t\}$. I begin by noting that (A.23) is a convex function of π_t and x_t, for any choice of $m_{t+1}(\cdot)$. Then since (A.24) is a linear relation among π_t, x_t, and $\pi_{t+1}(\cdot)$, it follows that, taking as given the choice of $m_{t+1}(cdot)$, the value of (A.23) implied by any choice of $\pi_{t+1}(\cdot)$ by the CB is a convex function of π_t and $\pi_{t+1}(cdot)$. Similarly, since (A.33) is linear, the value of (A.23) implied by any choice of p_t is a convex function of π_t and p_t, for any choice of $m_{t+1}(\cdot)$. Then since the maximum of a set of convex functions is a convex function, it follows that the maximized value of (A.23) is also a convex function of π_t and p_t. Thus $L(\pi_t; p_t; s_t)$ is a convex function of (π_t, p_t). Finally, a sum of convex functions is convex; this, together with the linearity of (A.33), implies that (A.32) is a convex function of the sequence $\{p_t\}$.

Convexity implies that the CB's optimal policy can be characterized by a system of first-order conditions, according to which

$$L_0(\pi_t; p_t; s_t) + \beta E_t L_\pi(\pi_{t+1}; p_{t+1}; s_{t+1}) = 0 \tag{A.35}$$

for each possible history h_t at any date $t \geq 0$, and

$$E_\rho E_0[L_1(\pi_t; p_t; s_t) + \beta L_\pi(\pi_{t+1}; p_{t+1}; s_{t+1}) w_{t+1}] = 0 \tag{A.36}$$

for each date $t \geq 0$. Here L_π denotes $\partial L/\partial \pi$, L_0 denotes $\partial L/\partial p^0$, and L_1 denotes $\partial L/\partial p^1$. Condition (A.35) is the first-order condition for the optimal choice of p^0_t, and (A.36) is the corresponding condition for the optimal choice of p^1_t. The latter condition is required to hold only in its *ex ante* (or unconditional) expected value, because I have defined a linear policy as one under which p^1_t does not depend on the history of realization of the shocks.

Note that it follows from the characterization in the previous section that for any plan satisfying (3.10), the partial derivatives just referred to are well-defined, and equal to

$$L_\pi(\pi_t; p_t; s_t) = \pi_t + \frac{\lambda}{\kappa^2} \frac{\pi_t - u_t - \kappa x^* - \beta p_t^0}{\Delta_t},$$

$$L_0(\pi_t; p_t; s_t) = -\beta \frac{\lambda}{\kappa^2} \frac{\pi_t - u_t - \kappa x^* - \beta p_t^0}{\Delta_t},$$

$$L_1(\pi_t; p_t; s_t) = \frac{\beta^2}{\theta} \left(\frac{\lambda}{\kappa^2}\right)^2 \left(\frac{\pi_t - u_t - \kappa x^* - \beta p_t^0}{\Delta_t}\right)^2 p_t^1.$$

Substituting (A.33) for π_t and (3.13) for Δ_t in these expressions, one obtains the first-order conditions (A.35) – (A.36) as restrictions upon the sequence $\{p_t\}$.

As explained in the text, I wish to find a sequence of functions $\{\varphi_t(\cdot)\}$ and a value $\bar{p}^1$ such that the linear policy under which

$$p_t^0 = \varphi_t(h_t), \qquad p_t^1 = \bar{p}^1$$

for all $t \geq 0$ satisfies the first-order conditions (A.35) – (A.36), when the initial measure ρ is the ergodic measure for the variables (p_t^0, s_t, w_{t+1}) under the policy just specified, and in addition $p_{-1}^1 = \bar{p}_1$. I first show that there exists a state-contingent evolution for $\{p_t^0\}$ that satisfies (A.35) in the case of an arbitrary constant value $\bar{p}^1$ that satisfies the bound (3.10), and for which there exists a well-defined ergodic measure. Using the ergodic measure ρ corresponding to a given value of $\bar{p}^1$, I then determine the nonlinear equation that $\bar{p}^1$ must satisfy in order for (A.36) to hold each period under the conjectured solution. Demonstration that a robustly optimal linear policy exists then requires only that one show that there exists a solution $\bar{p}^1$ to this equation that also satisfies the bound (3.10).

Under the assumption that $p_t^1 = \bar{p}^1$ for all $t \geq -1$, (A.35) is a stochastic linear difference equation for the process $\{p_t^0\}$ of the form

$$E_t[A(L)p_{t+1}^0] = v_t, \tag{A.37}$$

where

$$A(L) \equiv \beta - \left(1 + \beta + \frac{\kappa^2 \bar{\Delta}}{\lambda}\right) L + L^2,$$

$$v_t \equiv u_t - E_t u_{t+1} - \bar{p}^{1\prime} w_t.$$

(Here $\bar{\Delta}$ is the constant value of Δ_t implied by the constant value $\bar{p}^1$.) By factoring the lag polynomial in (A.37), one can easily show that (A.37) has a unique stationary solution, given by

$$p_t^0 = \mu p_{t-1}^0 - \mu E_t[(1 - \beta\mu L^{-1})^{-1} v_t], \tag{A.38}$$

where $0 < \mu < 1$ is the smaller root of the characteristic equation (3.21) given in the text. Note that a stationary solution exists regardless of the value assumed for $\bar{p}^1$. It is then straightforward to solve for the ergodic measure ρ associated with a given value of $\bar{p}^1$.

Equation (A.38) is a solution for the dynamics of $\{p_t^0\}$ of the kind indicated by equation (3.18) in the text. In the special case in which w_t is a scalar and $u_t = \sigma_u w_t$, $v_t = (\sigma_u - \bar{p}^1) w_t$, and (A.38) reduces to

$$p_t^0 = \mu p_{t-1}^0 - \mu(\sigma_u - \bar{p}^1) w_t. \tag{A.39}$$

Thus we have established conditions (3.19) given in the text. As noted in the text, it is evident from (3.21) that μ is monotonically decreasing in $\bar{\Delta}$. Since a concern for robustness results in $\bar{\Delta} < 1$, while $\bar{\Delta} = 1$ in the case of rational expectations, we see that a concern for robustness results in a value of μ that is larger (closer to 1), implying more persistence in the fluctuations in $\{p_t^0\}$.

It remains to determine when condition (A.36) is also satisfied. I first observe that

$$\begin{aligned} E_\rho E_0[L_1(\pi_t; p_t; s_t)] &= \frac{\beta^2}{\theta}\left(\frac{\lambda}{\kappa^2}\right)^2 \frac{\bar{p}^1}{\bar{\Delta}^2} E[(\pi_t - u_t - \kappa x^* - \beta p_t^0)^2] \\ &= \frac{\beta^2}{\theta}\left(\frac{\lambda}{\kappa^2}\right)^2 \frac{\bar{p}^1}{\bar{\Delta}^2}[a + 2b\bar{p}^1 + (\bar{p}^1)^2], \end{aligned}$$

where

$$\begin{aligned} a &\equiv E[(p_{t-1}^0 - u_t - \kappa x^* - \beta p_t^0)^2], \\ b &\equiv E[w_t(p_t^0 - u_t - \kappa x^* - \beta p_t^0)]. \end{aligned}$$

Here $E[\cdot]$ denotes the expectation under the ergodic measure associated with the dynamics for $\{p_t^0\}$ indicated by (A.38) — which measure is uniquely defined in the case of a given value of $\bar{p}^1$.

Similarly, one can show that

$$\begin{aligned} E_\rho E_0[L_\pi(\pi_{t+1}; p_{t+1}; s_{t+1}) w_{t+1}] &= E[\pi_{t+1} w_{t+1}] + \frac{\lambda}{\kappa^2 \bar{\Delta}} E[(\pi_{t+1} - u_{t+1} - \kappa x^* - \beta p_{t+1}^0) w_{t+1}] \\ &= \bar{p}^1 + \frac{\lambda}{\kappa^2 \bar{\Delta}}[\bar{p}^1 + b]. \end{aligned}$$

Hence condition (A.36) is equivalent to

$$f(\bar{p}^1) \equiv \frac{\beta^2}{\theta}\left(\frac{\lambda}{\kappa^2}\right)^2 \frac{c}{\bar{\Delta}^2}\bar{p}^1 + \bar{p}^1 + \frac{\lambda}{\kappa^2\bar{\Delta}}[\bar{p}^1 + b] = 0, \qquad \text{(A.40)}$$

where

$$c \equiv a + 2b\bar{p}^1 + (\bar{p}^1)^2.$$

A robustly optimal linear policy then exists if and only if (A.40) has a solution $\bar{p}^1$ that satisfies the bound (3.10). Of course, in defining the function $f(\cdot)$, one must take account of the dependence of c and $\bar{\Delta}$ on the value of $\bar{p}^1$.

When $\{p_t^0\}$ evolves in accordance with the stationary dynamics (A.39), the above definitions imply that

$$\begin{aligned} a &= (\kappa x^*)^2 + E\{[(1-\beta\mu)p_{t-1}^0 - (\sigma_u - \beta\mu(\sigma_u - \bar{p}^1))w_t]^2\} \\ &= (\kappa x^*)^2 + \frac{(1-\beta\mu)^2\mu^2}{1-\mu^2}(\sigma_u - \bar{p}^1)^2 + [(1-\beta\mu)\sigma_u + \beta\mu\bar{p}^1]^2, \end{aligned}$$

$$\begin{aligned} b &= -\sigma_u - \beta E[p_t^0 w_t] \\ &= -(1-\beta\mu)\sigma_u - \beta\mu\bar{p}^1. \end{aligned}$$

I furthermore observe that $a = a_0 + b^2$, where

$$a_0 \equiv (\kappa x^*)^2 + \frac{(1-\beta\mu)^2\mu^2}{1-\mu^2}(\sigma_u - \bar{p}^1)^2 > 0.$$

Hence

$$c = a_0 + (b + \bar{p}^1)^2 > 0$$

can be signed for all admissible values of $\bar{p}^1$. Substituting this function of $\bar{p}^1$ for c and (3.13) for $\bar{\Delta}$ in (A.40) yields a nonlinear equation in $\bar{p}^1$, that is solved numerically in order to produce Figure 1.

One can easily show that a solution to this equation in the admissible range must exist. Note first that (3.10) can alternatively be written in the form

$$|\bar{p}^1| < \hat{p}^1 \equiv \frac{\kappa}{\lambda^{1/2}}\frac{\theta^{1/2}}{\beta}.$$

I next observe that

$$f(0) = \frac{\lambda}{\kappa^2\bar{\Delta}}b = -\frac{\lambda}{\kappa^2}(1-\beta\mu)\sigma_u < 0.$$

On the other hand, in the case of any finite θ, as $p^1 \rightarrow \hat{p}^1$, the first term in the expression (A.40) becomes larger than the other two terms, so that $f(p^1) > 0$ for any value of p^1 close enough to (while still below) the bound, Since the function $f(\cdot)$ is well-defined and continuous on the entire interval $[0,\ \hat{p}^1)$, there must be an intermediate value $0 < \bar{p}^1 < \hat{p}^1$ at which $f(\bar{p}^1) = 0$. Such a value satisfies both (3.10) and (A.40), and so describes a robustly optimal linear policy.

It remains to establish (3.20) and (3.22). When evaluated at the value $p^1 = \mu\sigma_u$, the second two terms in (A.40) are equal to

$$-\frac{\lambda}{\kappa^2\bar{\Delta}}P(\mu)\sigma_u = 0,$$

where $P(\mu)$ is the polynomial defined in (3.21). Moreover, in the limiting case in which $\theta \rightarrow \infty$ (the RE case), the first term in condition (A.40) is identically zero, so that $f(\mu\sigma_u) = 0$, and $\bar{p}^1 = \mu\sigma_u$ is a solution.[29] Instead, when θ is finite, the first term is necessarily positive, so that $f(\mu\sigma_u) > 0$. If $\mu\sigma_u < \hat{p}^1$, this implies that there exists a solution to (A.36) such that $0 < \bar{p}^1 < \mu\sigma_u$, as asserted in (3.22). If instead $\hat{p}^1 \leq \mu\sigma_u$, then (3.22) follows from the result in the previous paragraph. Hence in either case, the robustly optimal policy satisfies (3.22) for any finite θ.

[29] It is easily seen to be the unique solution, since $f(p)$ is linear in this case. One can also show that this is the optimal policy without restricting attention to linear policies, as is done here; see Clarida *et al.* (1999) or Woodford (2003, chap. 7).

References

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

Gaspar, Vitor, Frank Smets and David Vestin, "Optimal Monetary Policy under Adaptive Learning," unpublished, European Central Bank, August 2005.

Gilboa, Itzhak, and David Schmeidler, "Maxmin Expected Utility with Nonunique Prior," *Journal of Mathematical Economics* 18: 141-153 (1989).

Hansen, Lars Peter and Thomas J. Sargent, "Robust Estimation and Control under Commitment," unpublished, University of Chicago, June 2005a.

—- —- and —- —-, "Recursive Robust Estimation and Control without Commitment," unpublished, University of Chicago, July 2005b.

—- —- and —- —-, *Robustness*, unpublished, University of Chicago, September 2005c.

—- —-, —- —-, Gauhar Turmuhambetova, and Noah Williams, "Robust Control and Misspecification," unpublished, University of Chicago, September 2005.

Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini, "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences," unpublished, Università Bocconi, July 2004.

—- —-, —- —-, and —- —-, "Dynamic Variational Preferences," unpublished, Università Bocconi, July 2005.

Orphanides, Athanasios, and John C. Williams, "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," in B.S. Bernanke and M. Woodford, eds., *The Inflation Targeting Debate,* Chicago: University of Chicago Press, 2005a.

—- —- and —- —-, "Monetary Policy with Imperfect Knowledge," unpublished, Federal Reserve Board, August 2005b.

Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy,* Princeton: Princeton University Press, 2003.