Robustly Optimal Monetary Policy with Near-Rational Expectations *

Michael Woodford

August 4, 2009

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.*

JEL: *D81, D84, E52*

Keywords: *robust control, multiplier preferences, commitment, discretion*

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

*Department of Economics, Columbia University, 420 W. 118th Street, New York, NY 10027. Email: michael.woodford@columbia.edu. I would like to thank Klaus Adam, George Evans, Lars Hansen, Christian Hellwig, Ricardo Reis, Tom Sargent, Chris Sims, Tom Tallarini, and Tack Yun for comments, Mauro Roca for research assistance, and the National Science Foundation for research support through grant number SBR-0422403.

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 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

[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 Athanasios Orphanides and John C. Williams (2005, 2007) and Vitor Gaspar, Frank Smets and David Vestin (2006).

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 CB 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.

NRE are given a precise meaning here by specifying a quantitative measure of the degree of discrepancy between the private-sector beliefs and the those of the central bank; the policy analyst entertains the possibility of any probability beliefs on the part of the private sector that are not too distant from the bank's under this (discounted relative entropy) measure. A robustly optimal policy is then the solution to a min-max problem, in which the policy analyst chooses a policy to minimize the value of her loss function in the case of those distorted beliefs that would *maximize* her expected losses under that policy.

Both this way of specifying the set of contemplated misperceptions and the conception of robust policy choice as a minmax problem follow the work of Lars Peter Hansen and Thomas J. Sargent (2007b). The robust policy problem considered here has some different elements, however, from the type of problems generally considered in the work of Hansen and Sargent. Their primary interest (as in the engineering literature on robust control) has been in the consequences of a policy analyst's uncertainty about the correctness of *her own* model of the economy,[3] rather than about the degree to which the private sector's expectations will agree with its own. Much of the available theory has been developed for cases in which private sector expectations are not an issue at all.

Hansen and Sargent (2003; 2007b, chap. 16) do discuss a class of "Stackelberg problems" in which a "leader" chooses a policy taking into account not only the optimizing response of the "follower" to the policy, but also the fact that the follower

[3]Of course, I do not mean to minimize the relevance of this kind of uncertainty for practical policy analysis, even though I abstract from it here in order to focus on a different issue.

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.[4] 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.*[5] Anastasios G. Karantounias, Hansen and Sargent (2007) consider an optimal dynamic fiscal policy problem, in which private sector expectations of future policy are a determinant of the effects of policy, as here.[6] But again the concern is with possible misspecification of the policy analyst's model, and since in this case the objective of the policy analyst and the representative private household are assumed to be the same, the misspecifications about which both the policy analyst and households are assumed to be most concerned are the same.

One might think that this difference should not matter in practice, if the policy analyst's objective coincides with that of the private sector — 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 monetary 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 with the greatest *subjective losses* from the point of view of the private sector. In the problem consid-

[4]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.

[5]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).

[6]See also Justin Svec (2008) for analysis of a similar problem.

ered here, instead, the NRE beliefs that matter are those that result in an equilibrium with the greatest expected losses under the central bank's probability beliefs. Even if the loss function is identical for the central bank and the private sector, I assume that it is the *policy analyst's* evaluation of expected losses that matters for robust policy analysis.

A number of papers have also considered the consequences of a concern for robustness for optimal monetary policy using a "new Keynesian" model of the effects of monetary policy similar to the one assumed below (*e.g.*, Richard Dennis, 2007; Kai Leitemo and Ulf Söderström, 2008; Carl E. Walsh, 2004).[7] Like Hansen, Sargent, and co-authors, these authors assume that the problem is the policy analyst's doubt about the correctness of her own model, and assume that in the policy analyst's "worst case" analysis, private-sector expectations are expected to be based on the same alternative model as she fears is correct. These papers also model the class of contemplated misspecifications differently than is done here, and (in the case of both Dennis and Leitemo-Soderstrom) assume discretionary optimization on the part of the central bank, rather than analyzing an optimal policy commitment. Nonetheless, it is interesting to observe some qualitative similarities of the conclusions reached by these authors and the ones obtained below on the basis of other considerations.[8]

Section 1 introduces the policy problem that I wish to analyze, defining "near-rational expectations" and a concept of robustly optimal policy. Section 2 then characterizes the robustly optimal policy commitment. Section 3 considers, for comparison, policy in a Markov perfect equilibrium under discretion, in order to investigate the degree to which commitment improves policy in the case of near-rational expectations. Section 4 concludes.

---

[7]An early contribution to this literature, Marc P. Giannoni (2002), considers a problem even less closely related to the problem treated here. Not only is Giannoni concerned with potential misspecification of the cental bank's model, but the policies considered are restricted to parametric families of interest-rate reaction functions.

[8]For example, in three of the four cases considered by Leitemo and Soderstrom (2008), they find that optimizing policy will allow less response of inflation to "cost-push" shocks than would occur in the absence of a concern for robustness, as is also true here, both under a robustly optimal commitment and in a robust Markov perfect equilibrium.

1 Stabilization Policy with Near-Rational Expectations

Here I develop the general idea sketched above in the context of a specific example, that weakens the assumption regarding private-sector expectations in the well-known analysis by Richard Clarida, Jordi Gali and Mark Gertler (1999) of optimal monetary policy in response to "cost-push shocks." This example is chosen because the results under the assumption of rational expectations will already be familiar to many readers.

1.1 The Objective of Policy

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

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

where $\kappa > 0, 0 < \beta < 1$, $\hat{\mathrm{E}}_t[\cdot]$ 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 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[9] of (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

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

where $\lambda > 0, \quad x^* \geq 0$, and the discount factor β is the same as in (1). Here $\mathrm{E}_t[\cdot]$ denotes the conditional expectation of a variable under the CB beliefs, which the policy analyst treats 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. (The condition expectation is taken with respect to the economy's state at date -1, *i.e.,* before the realization of the period zero disturbance.) I do not allow

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

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.[10] The CB believes that the exogenous state u_t is drawn independently each period from a normal distribution; specifically, it believes that

$$u_t = \sigma_u w_t, \tag{3}$$

where w_t is distributed i.i.d. $N(0, 1)$.[11] Note that this property of the joint distribution of the $\{u_t\}$ 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 economy's exogenous state. 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.[12] This is likely to require that someone in the central bank 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);[13] I assume uncertainty about PS expectations only at

[10]Thus I abstract here from the main kind of uncertainty considered by Hansen and Sargent (2007b).

[11]This notation allows us to consider the effects of variation in the volatility of the cost-push shocks, without changing the CB beliefs about the probability of different states identified by histories $\{w_t\}$.

[12]Even so, the assumption that the central bank commits itself to a state-contingent path for 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 need not always lead to the same set of equilibrium allocations being consistent with near-rational expectations. This raises the question of which form of policy commitment is most robust to potential departures from rational expectations, a topic to be addressed in future work. It should not be assumed that the robustly optimal strategy within this class is necessarily also optimal within some broader class of specifications.

[13]In general, implementation of a desired state-contingent inflation rate regardless of the nature of (possibly distorted) PS inflation expectations requires the central bank to directly monitor and

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 (1), and thus a well-defined value for CB expected losses (2).

The analysis is made considerably more tractable if the set of contemplated strategies is further restricted. A *linear policy* is one under which the planned inflation target at each date is a linear function of the history of shocks,

$$\pi_t = \alpha_t + \sum_{j=0}^{t} \phi_{j,t} w_{t-j}, \tag{4}$$

for some coefficients $\{\alpha_t, \phi_{j,t}\}$ that may be time-varying, but evolve *deterministically*, rather than themselves depending on the history of shocks. Restriction of attention to policies in this class has the advantage that a closed-form solution for the worst-case near-rational beliefs is possible, as shown in section 2.[14] And the optimal policy under rational expectations (RE), characterized by Clarida *et al.* (1999), belongs to this family of policies. In the case of a concern for robustness with respect to near-rational expectations, the restriction to linear policies is no longer innocuous. But the characterization of robustly optimal policy within this class of policies is nonetheless of interest. As we shall see, the optimal policy under RE is no longer the optimal choice, even within this restricted class of policies, and the coefficients of the robustly optimal linear policy rule provide a convenient parameterization of the ways in a concern for robustness changes the optimal conduct of policy.[15] Moreover, the (Markov-perfect equilibrium) policy resulting from discretionary optimization under RE is also a linear policy. Thus a consideration of robustly optimal policy within this class also suffices to allow us to determine to what extent allowance for departures from RE may lead optimal policy to resemble discretionary policy under the RE analysis.

respond to those expectations, as in the "expectations-based" approach to implementation proposed by George W. Evans and Seppo Honkapohja (2003).

[14]To be precise, what is needed is that the policy be *conditionally* linear, in the sense defined in (9) below. In the case that $\{u_t\}$ evolves in accordance with a more general linear process, rather than being i.i.d., what is needed is conditional linearity in the period t *innovation*, and not necessarily linearity in the disturbance u_t, as shown in Woodford (2005).

[15]Consideration of the extent to which the robustly optimal policy within a more flexible class of contemplated policies may differ from the robustly optimal linear policy is an important topic for further study.

1.2 Near-Rational Expectations

I turn now to the specification of PS beliefs. These will be described by a probability measure over possible paths for the evolution of the exogenous and endogenous variables, that need not coincide with that of the policy analyst. I do assume that in each of the equilibria contemplated by the policy analyst as possible outcomes under a given policy, the PS is expected to act on the basis of a coherent system of probability beliefs that are maintained over time.[16] (Thus PS conditional probabilities at any date t are determined by Bayesian updating, given the PS prior over possible paths and the paths of exogenous and endogenous variables up to that date.) These probability beliefs need not correspond, however, to any particular theory about how inflation or other variables are determined. The policy analyst neither assumes that the PS believes that inflation is determined by a New Keynesian Phillips curve nor that it believes in some other theory; the assumption of "near-rationality" is instead an assumption about the degree of correspondence between PS probability beliefs (however obtained) and those of the policy analyst herself.[17]

Why I do not require the analyst to assume that PS probability beliefs coincide exactly with her own, I propose that she should not expect them to be completely unlike her own calculation of the probabilities of different outcomes, either. 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 the analyst about which outcomes have zero probability. (More precisely, I shall assume that all contemplated PS beliefs are absolutely continuous *over finite time intervals*, as in Hansen *et al.* (2006).[18]) Thus if policy ensures that something

[16]This differs from the assumption made in analyses of optimal policy with PS "learning" such as those of Orphanides and Williams (2005, 2007) or Gaspar *et al.* (2006).

[17]Under the interpretation taken here, the conventional hypothesis of RE is not an assumption that people "know the true model" and correctly solve its equations, but rather an assumption that they have probability beliefs that *coincide with the analyst's own calculation* of equilibrium outcomes. This coincidence might be thought to arise because people in the economy share the analyst's model; but it might also be expected to result from observation of empirical frequencies, without any understanding of why those probabilities constitute an equilibrium.

[18]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.

always occurs, or that it *never* occurs, the policy analysts expects the PS to notice this, though it may misjudge the probabilities of events that occur with probabilities between zero and one.

The assumption of absolute continuity implies that there must exist a scalar-valued "distortion factor" m_{t+1}, a function of the history h_{t+1} of exogenous states to that point, satisfying

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

and such that

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

for any random variable X_{t+1}.[19] In effect, we may suppose that people correctly understand the equilibrium mapping from states of the world to outcomes — thus, the function $X_{t+1}(h_{t+1})$ — even if they do not also correctly assign probabilities to states of the world, as would be required for an RE equilibrium. I assume this, however, *not* on the ground that people understand and agree with the policy analyst's model of the economy, but simply on the ground that they agree with the policy analyst about zero-probability events; since in equilibrium (according to the calculations of the policy analyst), a history h_{t+1} is *necessarily* associated with a particular value X_{t+1}, the PS is also expected to assign probability one to the value X_{t+1} in the event that history h_{t+1} is realized, though they need not agree with the analyst about the probability of this event.

This representation of the distorted beliefs of the private sector is useful in defining a measure of the distance of the private-sector beliefs from those of the policy analyst. As discussed in Hansen and Sargent (2005, 2007a, b), the *relative entropy*

$$R_t \equiv \mathrm{E}_t[m_{t+1} \ \log m_{t+1}]$$

is a measure of the distance of (one-period-ahead) PS beliefs from the CB beliefs with a number of appealing properties.[20] In particular, PS beliefs that are not too different

[19] The existence of the function $m(h_{t+1})$ is guaranteed by the Radon-Nikodym theorem. In the case of a discrete set of states w that are possible at date $t+1$, given the economy's state at date t, $m(w)$ is simply the ratio $\hat{\pi}(w)/\pi(w)$, where $\pi(w)$ is the probability assigned by the CB to state w and $\hat{\pi}(w)$ is the probability assigned by the PS to that state. This way of describing distorted beliefs is used, for example, by Hansen and Sargent (2005, 2007a) and Hansen *et al.* (2006).

[20] For example, R_t is a positive-valued, convex function of the distorted probability measure, uniquely minimized (with the value zero) when $m_{t+1} = 1$ almost surely (the case of RE).

from those of the policy analyst in the sense that R_t is small are ones 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 any given distorted beliefs as more plausible the smaller is R_t.

The overall degree of distortion of PS probability beliefs about possible histories over the indefinite future can furthermore be measured by a discounted relative entropy criterion

$$\mathrm{E}_{-1} \sum_{t=0}^{\infty} \beta^t m_{t+1} \log m_{t+1},$$

as in Hansen and Sargent (2005). We shall suppose that the policy analyst wishes to guard against the outcomes that can result under any PS beliefs that do not involve too large a value of this criterion. The presence of the discount factor β^t in this expression implies 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.[21]

More precisely, I shall assume that the policy analyst seeks to ensure as small as possible a value for an augmented loss function

$$\mathrm{E}_{-1} \sum_{t=0}^{\infty} \beta^t \frac{1}{2}[\pi_t^2 + \lambda(x_t - x^*)^2] - \theta \mathrm{E}_{-1} \sum_{t=0}^{\infty} \beta^t m_{t+1} \log m_{t+1} \qquad (5)$$

in the case of *any* possible PS beliefs.[22] The presence of the second term indicates that

[21] If we omit the discount factor β^t in our distance measure, the consequences are the same for the objective (5) below as if one were to assume, instead of a constant "cost" θ of departure from CB beliefs, a cost $\theta\beta^{-t}$ that grows the farther into the future one looks. But a large value of θ allows little departure from RE, so such a specification would imply much less allowance for potential PS misunderstanding far in the future, relative to the one adopted here. I show below that under the specification proposed here, the degree of distortion involved in the "worst-case" NRE beliefs contemplated by the policy analyst is time-invariant. See Hansen *et al.* (2006) for discussion of this issue, in the context of a continuous-time analysis.

[22] Technically, this criterion is defined only for PS beliefs that satisfy the absolute continuity condition discussed above. But if we define the relative entropy to equal $+\infty$ in the case of any beliefs that are not absolutely continuous with respect to those of the CB, then (5) can be defined for arbitrary PS beliefs.

the policy analyst is not troubled by the fact that outcomes could be worse (from the point of view of the stabilization objective (2) in the case of distorted PS expectations, as long as the distance of the beliefs in question from those of the CB (as measured by relative entropy) is sufficiently great relative to the increased stabilization losses. Thus the analyst will only worry about 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.[23] The coefficient $\theta > 0$ measures the analyst's degree of concern for possible departures from RE, with a small value of θ implying a great degree of concern for robustness, while a large value of θ implies that only modest departures from RE are considered plausible. In the limit as $\theta \to \infty$, the RE analysis is recovered as a limiting case of the present one.

1.3 Robustly Optimal Commitment

In the case of any policy commitment $\{\pi_t\}$ contemplated by the policy analyst, and any distorted PS beliefs described by a distortion factor $\{m_{t+1}\}$, one can determine the implied value of (5) by solving for the equilibrium process $\{x_t\}$ implied by (1). Let this value be denoted $\mathcal{L}(\pi, m)$. The *robustly optimal* policy is then the policy π that minimizes

$$\bar{\mathcal{L}}(\pi) \equiv \sup_m \mathcal{L}(\pi, m), \tag{6}$$

so as to ensure as low as possible an upper bound for the value of (5) under any equilibrium that may result from the pursuit of the policy.[24]

[23] "Multiplier preferences" of this form are used extensively by Hansen and Sargent (2007b) to model robust decisionmaking. Axiomatic foundations for preferences of this form are provided by Tomasz Strzalecki (2008).

[24] Alternatively, one might suppose that the policy analyst should choose a policy that minimizes the value of (2) under worst-case NRE beliefs, where the latter are defined as the distortion m that solves the inner problem in (6). Apart from the appeal of the axiomatic foundations offered by Strzalecki (2007) for the "multiplier preferences" used here, 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 (2006), and the same qualitative results are obtained in that case, though some of the algebra is different. See Woodford (2005, Appendix A.3) for comparison of the results under the alternative assumptions.

The policy problem (6) can be thought of as a "game" between the CB and a "malevolent agent" that chooses the PS beliefs so as to frustrate the CB's objectives. However, our interest in this problem does not depend on any belief in the existence of such an "agent." Consideration of the min-max problem (6) is simply a way of ensuring that the policy chosen is as robust as possible to possible departures from RE, without sacrificing too much of the CB's stabilization objectives. It is also sometimes supposed that selection of a policy that minimizes (6) requires extreme pessimism on the part of the policy analyst, since only the "worst-case" distorted beliefs are used to evaluate each contemplated policy. But use of this criterion does not require that the policy analyst believe that the equilibrium that would result from worst-case PS beliefs (the distortion m that solves the problem of the "malevolent agent") is the one that *must occur*. The "malevolent agent's" problem is considered only because this is a convenient mathematical approach to determining the upper bound on losses under a given policy.

The policy $\{\pi_t\}$ for periods $t \geq 0$ that minimizes (6), under no constraints beyond the assumption of linearity (4), is in general not time-invariant (the optimal coefficients for the rule for π_t will vary with the date t), and also not time-consistent (re-optimization at some later date would not lead the policy analyst to choose to continue the sequence of inflation commitments chosen at date zero), for reasons that are familiar from the literature on policy analysis under RE.[25] Both of these complications result from the fact that one supposes that the CB can choose an inflation rate for period zero without having to take account of any effects of its choice on inflation expectations prior to date zero, while the inflation rate that it chooses for any date $t \geq 1$ has consequences for PS inflation expectations,[26] and hence for the feasible degree of inflation and output-gap stabilization in earlier periods. We can instead obtain an optimal policy problem with a recursive structure (the solution to which is a time-invariant policy rule) if, instead of supposing that the policy analyst chooses a sequence of (possibly time-varying) inflation commitments $\{\pi_t\}$ for all $t \geq 0$, we consider only the problem of choosing an optimal sequence of inflation commitments for periods $t \geq 1$, taking as given a commitment $\pi_0(w_0)$ that the CB's policy must

[25] The issue is discussed in detail in Woodford (2003, chap. 7).

[26] While I do not assume that PS expectations must *exactly coincide* with the CB's policy intention, as in the RE analysis, nevertheless the CB's state-contingent policy intention affects the plausibility of particular PS inflation expectations, as long as $\theta > 0$.

fulfill.[27]

Let us suppose that the initial inflation commitment is itself linear,

$$\pi_0(w_0) = p^0_{-1} + p^1_{-1} w_0, \tag{7}$$

for some coefficients (p^0_{-1}, p^1_{-1}), and let us suppose that the inflation commitments that are chosen for periods $t \geq 1$ may depend on the value of p^0_{-1}, as well as upon the shocks that occur in periods zero through t. Thus we consider policies that can be written in the form

$$\pi_t = \alpha_t + \gamma_t p^0_{-1} + \sum_{j=0}^{t} \phi_{j,t} w_{t-j}, \tag{8}$$

for some coefficients $\{\alpha_t, \gamma_t, \phi_{j,t}\}$. The separate term $\gamma_t p^0_{-1}$ matters because I shall suppose that the coefficients of the linear rule are chosen before the value of p^0_{-1} is known, and are to apply regardless of that value (that may depend on the economy's state at date -1). I shall let Φ denote the set of linear policies for dates $t \geq 1$ of the form (8).

It will also be useful to discuss the broader set Π of *conditionally linear* policies, under which the state-contingent inflation rate one period in the future can be written

$$\pi_{t+1}(w_{t+1}) = p^0_t + p^1_t w_{t+1} \tag{9}$$

in any period $t \geq 0$, where p^0_t may depend on both the state h_t and the initial condition p^0_{-1}, but p^1_t depends only on the date. Any policy $\phi \in \Phi$ corresponds to a policy $p \in \Pi$, where the coefficients p are given by

$$p^0_t(h_t; p^0_{-1}) = \alpha_{t+1} + \gamma_{t+1} p^0_{-1} + \sum_{j=1}^{t+1} \phi_{j,t+1} w_{t+1-j},$$

$$p^1_t = \phi_{0,t+1}$$

for each $t \geq 0$.

For any given initial commitment (p^0_{-1}, p^1_{-1}) and policy $p \in \Pi$, we can compute an expected value for the augmented loss function $\mathcal{L}(p^0_{-1}, p^1_{-1}, p, m)$ as above, in the case of any contemplated PS beliefs m; and we can correspondingly define the upper bound $\bar{\mathcal{L}}(p^0_{-1}, p^1_{-1}, p)$. Now suppose that the coefficient p^0_{-1} is drawn from a distribution ρ,

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

and that a policy p, specifying the coefficients of (9) for all possible states at dates $t \geq 0$, must be chosen to apply in the case of any realization of p^0_{-1} in the support of this distribution. (The PS beliefs m may depend on the realization of p^0_{-1}.) Then the upper bound on the value of $\mathrm{E}_\rho[\mathcal{L}(p^0_{-1}, p^1_{-1}, p, m)]$ for a given policy p is equal to

$$\hat{\mathcal{L}}(p; p^1_{-1}, \rho) \equiv \mathrm{E}_\rho[\bar{\mathcal{L}}(p^0_{-1}, p^1_{-1}, p)],$$

where E_ρ indicates integration over the distribution ρ of possible values for p^0_{-1}. A *robustly optimal* linear policy commitment is then a set of coefficients $\phi \in \Phi$ that solve the problem

$$\inf_{\phi \in \Phi} \hat{\mathcal{L}}(p(\phi); p^1_{-1}, \rho), \tag{10}$$

where $p(\phi)$ identifies the coefficients $\{p^0_t, p^1_t\}$ corresponding any given linear policy ϕ.

The assumed initial commitment is *self-consistent* if it is a form of commitment that the policy analyst chooses in subsequent periods under the problem just defined.[28] To be precise, initial commitments $(\bar{p}^1, \bar{\rho})$ are self-consistent if when we set $p^1_{-1} = \bar{p}_1, \rho = \bar{\rho}$, the worst-case equilibrium associated with the policy ϕ that solves (10) is such that (i) $p^1_t = \bar{p}_1$ for each $t \geq 0$; and (ii) the unconditional distribution ρ_t of values for the coefficient p^0_t (integrating over the distribution ρ of possible values for p^0_{-1} and over the distribution of possible shocks in each of periods zero through t) is equal to $\bar{\rho}$ for each $t \geq 0$. One can show that a self-consistent specification of the initial commitments is possible, and in this case the robustly optimal linear policy has a time-invariant form, as discussed in the next section.

2 The Robustly Optimal Linear Policy

In this section, I characterize the solution to the optimal policy problem under commitment defined in the previous section, and compare it to the optimal policy under commitment in the RE analysis (as derived for example in Clarida *et al.*, 1999). This means finding the linear policy ϕ of the form (8) that solves (10) in the case that (p^1_{-1}, ρ) are the self-consistent initial commitments $(\bar{p}^1, \bar{\rho})$.

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

2.1 The "Worst-Case" NRE Beliefs

I begin by characterizing the "worst-case" NRE beliefs in the case of any given conditionally linear policy $\{\pi_t\}$. These are characterized by the process $\{m_{t+1}\}$ that solves the "malevolent agent's" problem on the right-hand side of (6). This is the process $\{m_{t+1}\}$ for all $t \geq 0$ that maximizes (5) subject to the constraint that $E_t m_{t+1} = 1$ at all times, where at each date x_t is the solution to the equation

$$\pi_t = \kappa x_t + \beta E_t[m_{t+1}\pi_{t+1}] + u_t. \tag{11}$$

This problem 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

$$\frac{1}{2}[\pi_t^2 + \lambda(x_t - x^*)^2] - \theta E_t[m_{t+1}\log m_{t+1}] \tag{12}$$

subject to the constraint that $E_t m_{t+1} = 1$, where again x_t is given by (11). This single-period problem has a closed-form solution in the case that the commitment $\pi_{t+1}(w_{t+1})$ is of the conditionally linear form (9), where the coefficients (p_t^0, p_t^1) depend only on the history h_t.

One notes that an interior solution to the problem of maximizing (12) exists only if[29]

$$|p_t^1|^2 < \frac{\theta}{\beta^2}\frac{\kappa^2}{\lambda}. \tag{13}$$

Otherwise, the objective (12) 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 (13) at all times. This provides an immediate contrast with optimal policy under RE, where the optimal coefficient p^1 (which is constant over time) is proportional to σ_u, the standard deviation of the cost-push shocks.[30] 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

[29]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 (13) 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.

[30]See, *e.g.*, equation (26) below.

(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 (13), under the worst-case NRE, the CB fears that the PS will expect w_{t+1} to be conditionally distributed as $N(\mu_t, 1)$.[31] If $p_t^1 = 0$, $\mu_t = 0$, while if $p_t^1 \neq 0$,

$$\mu_t = (\bar{\pi}_t - p_t^0)/p_t^1, \tag{14}$$

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

$$\bar{\pi}_t = \Delta_t^{-1} \left[p_t^0 - (\pi_t - u_t - \kappa x^*) \frac{\beta \lambda}{\theta \kappa^2} |p_t^1|^2 \right], \tag{15}$$

$$\Delta_t \equiv 1 - \frac{\beta^2}{\theta} \frac{\lambda}{\kappa^2} |p_t^1|^2 > 0. \tag{16}$$

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 (13) is approached. As a consequence of this possibility, the CB fears an output gap equal to

$$x_t^{pess} - x^* = \frac{(\pi_t - u_t - \kappa x^*) - \beta p_t^0}{\kappa \Delta_t}. \tag{17}$$

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

$$R_t^{pess} \equiv \hat{E}_t[\log m_{t+1}] = \frac{1}{2} \left[\frac{\beta}{\theta} \frac{\lambda}{\kappa} (x_t - x^*) \right]^2 |p_t^1|^2 \geq 0. \tag{18}$$

[31] As shown in Woodford (2005), this result can easily be extended to the case of a vector of innovations upon which π_{t+1} may depend linearly, generalizing the formulas for the scalar case presented here.

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 (the most unfortunate) 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 (17) for the output gap and (18) for the relative entropy term in (5), we obtain a loss function for the CB of the form[32]

$$\hat{\mathcal{L}}(p; p_{-1}^1, \rho) = \mathrm{E} \sum_{t=0}^{\infty} \beta^t L(p_{t-1}; p_t; w_t), \tag{19}$$

defined for any policy $p \in \Pi$, where p_t is shorthand for the pair (p_t^0, p_t^1), and the unconditional expectation of any random variable X_t in a period $t \geq 0$ is defined as

$$\mathrm{E}[X_t] \equiv \mathrm{E}_\rho \mathrm{E}_{-1}[X_t].$$

The robust policy problem (10) can then be described as the choice of a linear policy that maximizes (19).

2.2 Dynamics of Optimal Commitment

Rather than directly considering the problem of finding the linear policy $\phi \in \Phi$ that maximizes (19), it is simpler to consider the problem of finding the conditionally linear $p \in \Pi$ that maximizes this objective, for some specification of the initial commitment (p_{-1}^1, ρ). In fact, the robustly optimal policy within this class is always a fully linear policy, so that we will have also found the robustly optimal element of the more restrictive class of policies Φ.

The reason for this is fairly simple. The optimal conditionally linear policy p must involve a process $\{p_t^0\}$ that is optimal taking as given the sequence $\{p_t^1\}$. But for a given sequence $\{p_t^1\}$ satisfying (13) at all dates, the loss function $L(p_{t-1}; p_t; w_t)$

[32] See Appendix A.1 for the explicit form of the function $L(p_{t-1}; p_t; w_t)$.

is a (convex) quadratic function of (p^0_{t-1}, p^0_t, w_t), with coefficients that vary with t (in the case that the sequence $\{p^1_t\}$ is not constant over time). The problem of choosing an optimal state-contingent commitment $\{p^0_t\}$ given the sequence $\{p^1_t\}$ is therefore a convex linear-quadratic (LQ) optimal control problem, albeit one with (deterministically) time-varying coefficients. Hence the optimal solution is a linear policy of the form

$$p^0_t = \lambda_t + \mu_t p^0_{t-1} + \nu_t w_t, \tag{20}$$

where the coefficients $\{\lambda_t, \mu_t, \nu_t\}$ are deterministic sequences that depend on the sequence $\{p^1_t\}$. But equation (20), which must hold for all $t \geq 0$, together with the fact that the sequence $\{p^1_t\}$ is deterministic, imply that the sequence of conditionally linear inflation commitments (9) constitute a linear policy of the form (8). Hence the optimal conditionally linear policy must be a linear policy, and since all linear policies are conditionally linear, it must be the optimal linear policy.

The linear law of motion (20) also implies that if the unconditional distribution ρ_{t-1} for p^0_{t-1} is a normal distribution $N(\mu_{p,t-1}, \sigma^2_{p,t-1})$, then the unconditional distribution ρ_t for p^0_t will also be a normal distribution, with mean and variance given by a law of motion of the form

$$(\mu_{p,t}, \sigma^2_{p,t}) = \Psi(\mu_{p,t-1}, \sigma^2_{p,t-1}; \psi_t), \tag{21}$$

where $\psi_t \equiv (\lambda_t, \mu_t, \nu_t)$ is the vector of coefficients of the law of motion (20). Because of my interest in choosing a self-consistent initial commitment, I shall suppose that ρ is some normal distribution $N(\mu_{p,-1}, \sigma^2_{p,-1})$, in which case ρ_t will also be normal for all $t \geq 0$ under the optimal linear policy.

Finally, the first-order conditions for the optimal choice of the sequence $\{p^1_t\}$ can be written in the form[33]

$$g(p^1_{t-1}, p^1_t, p^1_{t+1}; \mu_{p,t-1}, \sigma^2_{p,t-1}; \psi_t, \psi_{t+1}) = 0 \tag{22}$$

for each $t \geq 0$. The conditionally linear policy that maximizes (19) for given initial conditions $(p^1_{-1}, \mu_{p,-1}, \sigma^2_{p,-1})$ then corresponds to deterministic sequences $\{p^1_t; \mu_{p,t}, \sigma^2_{p,t}; \psi_t\}$ for $t \geq 0$ that satisfy (21) and (22) for all $t \geq 0$, where the sequence of coefficients $\{\psi_t\}$ describe the solution to the LQ problem defined by the sequence of coefficients $\{p^1_t\}$.

[33] See Appendix A.2 for further discussion.

The self-consistent initial conditions $(\bar{p}^1, \bar{\mu}_p, \bar{\sigma}_p^2)$ are simply the steady-state solution to the above system of difference equations. I show in Appendix A.2 that such a steady state exists. The robustly optimal linear policy, *i.e.*, the policy that maximizes (19) in the case of initial conditions $(\bar{p}^1, \bar{\mu}_p, \bar{\sigma}_p^2)$, is then a policy with time-invariant coefficients, as asserted earlier. Here I compare the properties of this policy with the optimal policy under RE, and also compare equilibrium outcomes under the worst-case equilibrium consistent with this policy to the RE equilibrium outcomes under the RE-optimal policy.

2.3 Characteristics of Optimal Policy

Under the stationary policy corresponding to the steady state of the system (21)–(22), $p_t^1 = \bar{p}_1$ each period, where $\bar{p}_1$ is a positive quantity satisfying the bound (13). It then follows that the LQ problem that we must solve for the optimal state-contingent evolution $\{p_t^0\}$ involves a period loss function with constant coefficients. It follows that the coefficients of the law of motion (20) are time-invariant as well. In fact, one can show[34] that the law of motion takes the form

$$p_t^0 = \mu p_{t-1}^0 + \mu(\bar{p}^1 - \sigma_u) w_t, \tag{23}$$

where $0 < \mu < 1$ is the smaller root of the quadratic equation

$$P(\mu) \equiv \beta\mu^2 - \left(1 + \beta + \frac{\kappa^2 \bar{\Delta}}{\lambda}\right)\mu + 1 = 0. \tag{24}$$

Here $0 < \bar{\Delta} \leq 1$ is the constant value of (16) associated with $\bar{p}^1$. It then follows from (9) that the state-contingent inflation target evolves according to an ARMA(1,1) process

$$\pi_t = \mu\pi_{t-1} + \bar{p}^1 w_t - \mu\sigma_u w_{t-1} \tag{25}$$

for all $t \geq 1$.

Because $0 < \mu < 1$, (23) implies that $\{p_t^0\}$ is a stationary process, with a well-defined unconditional mean and variance $(\bar{\mu}_p, \bar{\sigma}_p^2)$. Moreover, the unconditional mean is zero — so that (25) implies that the inflation rate fluctuates around a long-run average value of zero as well — just as in the optimal policy commitment in the RE case, regardless of the assumed value of θ. Thus the optimal long-run inflation target

[34] This and the other results cited in this section are derived in Appendix A.2.

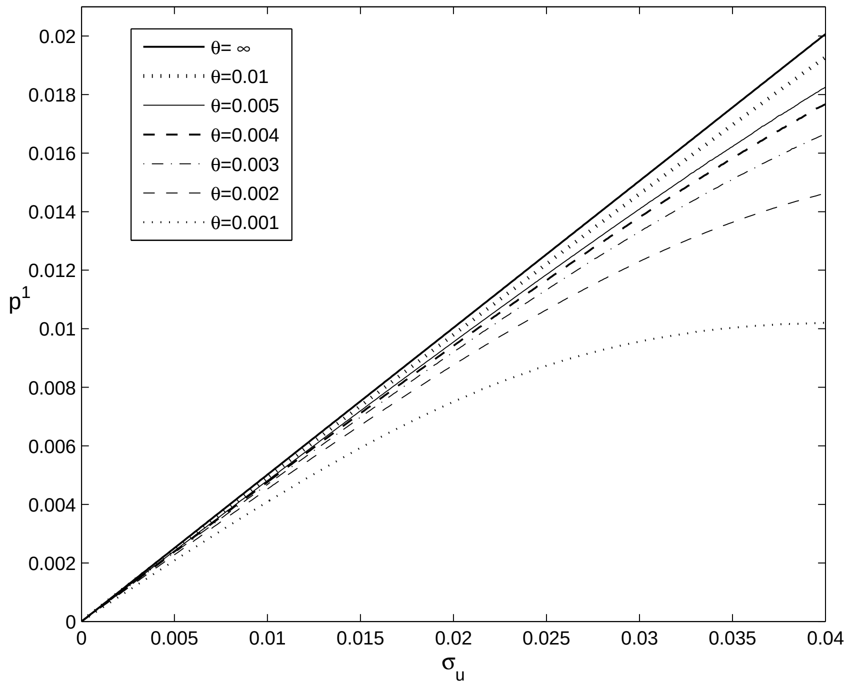

Figure 1: Variation of $\bar{p}^1$ with σ_u, under alternative degrees of concern for robustness.

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.[35]

According to the RE analysis, inflation also evolves according to a stationary ARMA(1,1) process with mean zero. But in the RE case, one can further show that

$$\bar{p}^1 = \mu\sigma_u, \tag{26}$$

so that (25) involves only the first difference of the cost-push shock. (In this case, the law of motion can equivalently be written as a stationary AR(1) process for the log price level.) In the case of a finite value of θ, instead, the optimal response coefficient necessarily satisfies

$$0 < \bar{p}^1 < \mu\sigma_u, \tag{27}$$

so that the price level is no longer stationary.

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

Figure 1 shows how $\bar{p}^1$ varies with σ_u for alternative values of θ.[36] In the RE case, $\bar{p}^1$ increases linearly with σ_u, as indicated by (26) and as required for certainty-equivalence. For any given amplitude of cost-push shocks, lower θ (greater concern 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 (13).

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 (2005) 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.[37] Moreover, (24) implies that μ is closer to 1 in the finite-θ case (where $\bar{\Delta} < 1$) than in the RE case (in which μ is also a root of (24), but with $\bar{\Delta} = 1$). Hence the effect of a past cost-push shock 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.[38] The lower right panel shows the impulse response of the

[36] 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) discussed in Woodford (2003, chap. 6).

[37] This is shown by the negative coefficient multiplying w_t in (23). Note that since $\mu < 1$, (27) implies that $\bar{p}^1 < \sigma_u$.

[38] 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 $\sigma_u = 0.02$. In the upper left panel, the inflation rate is an annualized rate; given that the model periods are interpreted as

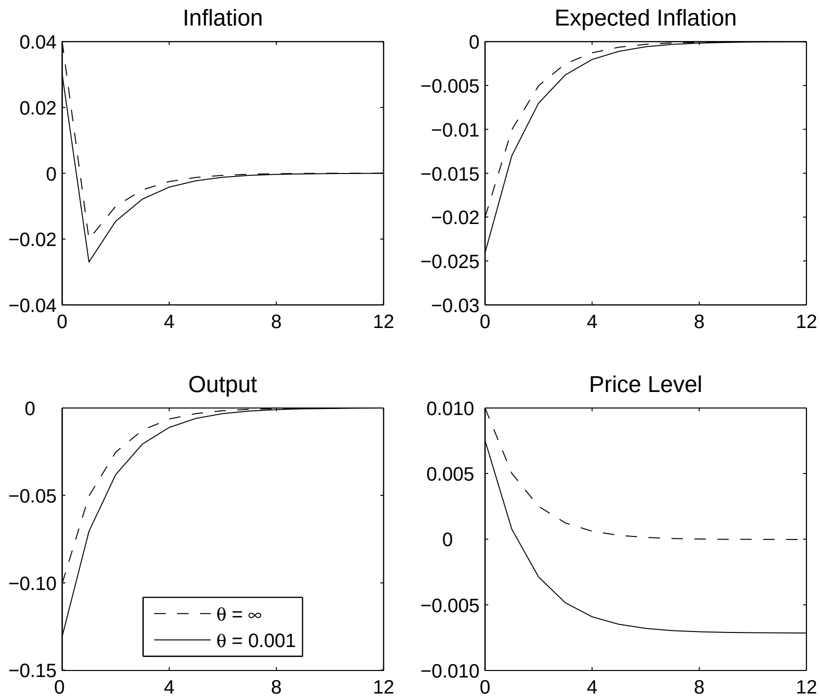

Figure 2: Optimal responses to a positive cost-push shock, with and without concern for robustness.

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, and follows from the fact that lagged MA term in (25) is larger than the contemporaneous term according to (27).

Of course, (25) describes the dynamics of inflation as they are understood by the central bank. PS forecasts of future inflation need not correspond to what this equation for inflation dynamics would imply. In the equilibrium with worst-case NRE expectations, PS inflation expectations evolve in accordance with (15). Substitution

quarters, "inflation" is four times the change in the log price level.

of the law of motion (23) allows this to be rewritten as

$$\bar{\pi}_t = \Lambda p_t^0 + \beta^{-1}(\bar{\Delta}^{-1} - 1)\kappa x^*, \tag{28}$$

where $\Lambda < 1$. This implies that PS inflation expectations are a linear function of CB inflation expectations, but with a bias ($\mathrm{E}[\bar{\pi}] > 0$, since $\bar{\Delta} < 1$ in the case of finite θ), and a derivative less than 1.[39]

The fact that $\Lambda < 1$ 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. 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 way (as shown in Figure 2), in order to ensure that inflation expectations are affected even if expectations are not perfectly model-consistent.

In the limit as $\theta \to 0$ (extreme concern for possible departures from RE), the optimal $\bar{p}^1 \to 0$. (In fact, this can be immediately seen from the bound (13).) In the limit, it is optimal for the CB to prevent cost-push shocks from having any immediate effect on inflation at all. This does not, however, mean that inflation is completely stabilized, for (23) still implies that the planned inflation rate in the *next* period is reduced in the event of a positive cost-push shock. (Note that μ remains bounded away from zero in this limit, since (24) implies that $0 < \mu^{RE} < \mu < 1$ for any value $0 < \bar{\Delta} < 1$.) It remains desirable to reduce intended subsequent inflation, because a reduction in $\hat{\mathrm{E}}_t \pi_{t+1}$ at the same time as an increase in u_t reduces the extent to which the output gap must become more negative due to the cost-push shock; even though PS expectations of inflation cannot be counted on to fall as much as the CB's intended inflation rate does, it is still worthwhile to reduce intended future inflation, in order to ensure *some* moderation of inflation expectations.

[39] In the limit as $\theta \to \infty$, $\bar{\Delta} \to 1$ and $\Lambda \to 1$, so that (15) implies that $\bar{\pi}_t = p_t^0$, the case of RE.

3 Near-Rational Expectations and the Importance of Policy Commitment

I have observed above that robustly optimal policy involves advance commitment, in a similar way as optimal policy under the assumption of rational expectations. But does the degree to which PS expectations may depart from model-consistency affect the *degree* to which commitment matters? In order to address this question, it is necessary to characterize equilibrium policy under discretionary optimization on the part of a CB that understands that private-sector expectations need not be fully model-consistent, and compare this to the robustly optimal policy under commitment.

Suppose that the objective of the central bank is to minimize (5), as above, but that each period the central bank chooses a short-run inflation target π_t after learning the current state w_t, without making any commitment as to the inflation rate that it may choose at any later dates. Because the payoffs and constraints of both the CB and the malevolent agent in the continuation game at date t are independent of the past, in a Markov perfect equilibrium (MPE), π_t will depend only on w_t. I shall assume an equilibrium of this kind;[40] hence there is assumed to exist a time-invariant policy function $\bar{\pi}(\cdot)$ such that in equilibrium $\pi_t = \bar{\pi}(w_t)$ each period. Under discretionary optimization, the CB takes for granted the fact that it will choose to follow the rule $\bar{\pi}(\cdot)$ in all *subsequent* periods, though it is not committed to follow it in the current period. The CB also takes for granted the set of possible NRE beliefs of the PS regarding the economy's future evolution, given that (at least in the view of the CB) the truth is that the exogenous state will be drawn independently each period from a unit normal distribution, monetary policy will follow the rule $\bar{\pi}(\cdot)$, and output will be determined by (1). It then chooses an inflation rate π_t to implement in the current period, given its own model of the economy's subsequent evolution and guarding against the worst-case NRE beliefs given that model. In a MPE, the solution to this problem is precisely the inflation rate $\pi_t = \bar{\pi}(w_t)$.

I shall formally define a robust MPE as follows. Given a policy rule $\bar{\pi}(\cdot)$, let $V(\pi_0; w_0)$ be the value of the objective (5) if the initial state is w_0, the CB chooses an inflation rate π_0 in that initial state and then follows the rule $\bar{\pi}(\cdot)$ in all periods $t \geq 1$,

[40] The restriction to Markov perfect equilibria is commonplace in the literature on discretionary monetary policy under rational expectations; the equilibrium concept proposed here generalizes the one used by Clarida *et al.* (1999) in their RE analysis of this model.

and PS beliefs correspond to the worst-case NRE beliefs given this policy. Then given the inflation rate chosen in any period, the worst-case NRE beliefs $m_{t+1}(\cdot)$ solve the problem

$$\max_{m_{t+1}(\cdot)} \frac{1}{2}[\pi_t^2 + \lambda(x_t - x^*)^2] - \theta E_t \beta^t m_{t+1} \log m_{t+1} + \beta \mathrm{E}_t V(\bar{\pi}(w_{t+1}; w_{t+1}), \tag{29}$$

where x_t is the solution to

$$\pi_t = \kappa x_t + \beta E_t[m_{t+1}\bar{\pi}(w_{t+1})] + u_t.$$

A *robust MPE* is then a pair of functions $\bar{\pi}(\cdot)$ and $V(\cdot;\cdot)$ such that for any pair $(\pi_t; w_t)$, $V(\pi_t; w_t)$ is the maximized value of (29), and for any state w_t, $\bar{\pi}(w_t)$ is the inflation rate that solves the problem

$$\min_{\pi_t} V(\pi_t; w_t). \tag{30}$$

A *robust linear MPE* is a robust MPE in which $\bar{\pi}(\cdot)$ is a linear function of the state,

$$\bar{\pi}(s_t) = \bar{p}^0 + \bar{p}^1 w_t, \tag{31}$$

for some constant coefficients $\bar{p} = (\bar{p}^0, \bar{p}^1)$.[41]

A linear policy (31) is an example of the kind of conditionally linear policy considered in the previous section. Moreover, because the final term in (29) is independent of the choice of $m_{t+1}(\cdot)$, the function $m_{t+1}(\cdot)$ that solves the problem (29) is also the one that maximizes (12), so that the characterization of worst-case NRE beliefs in appendix A.1 again applies. Once again, $|p^1|$ must satisfy the bound (13) in order for there to be well-defined worst-case beliefs;[42] and when this bound is satisfied, the worst-case beliefs are again described by (14) – (15).

[41] Note that it is not necessary, as in our discussion of robustly linear policy under commitment, to suppose that the CB optimizes over a restricted class of policy rules; in fact, in the discretionary policy problem (30), the CB does not choose a rule at all, but only an inflation rate in the particular state that has been realized. Nonetheless, I do not here address the question whether the linear MPE discussed below are the only possible kind of robust MPE.

[42] In the case of discretionary policy, I can no longer argue that the CB will surely choose a policy that satisfies (13), in order to avoid unbounded losses. For now the CB is assumed to choose π_{t+1} without taking into account the effect of the way in which the dependence of π_{t+1} on w_{t+1} affects the worst-case choice of $m_{t+1}(\cdot)$, given that the distorted PS beliefs are a historical fact by the time that π_{t+1} is chosen. Nonetheless, there can be no well-defined equilibrium in which (13) is violated.

Given this characterization of worst-case beliefs, the problem (30) of the discretionary central bank reduces to

$$\min_{\pi_t} \tilde{L}(\pi_t; \bar{p}; w_t), \tag{32}$$

where $\tilde{L}(\pi_t; p_t; w_t)$ is the loss function defined in appendix A.1.[43] Since for any w_t, $\tilde{L}$ is a strictly convex, quadratic function of π_t, the discretionary policy $\bar{\pi}(w_t)$ is just the solution to the first-order condition

$$\tilde{L}_\pi(\pi_t; \bar{p}; w_t) = 0.$$

This linear equation in π_t is easily solved, yielding

$$\bar{\pi}(w_t) = \frac{\lambda}{\kappa^2 \bar{\Delta} + \lambda}[\kappa x^* + u_t + \beta \bar{p}^0]. \tag{33}$$

This in turn implies that $\bar{\pi}(\cdot)$ is indeed a linear function of the form (31), where[44]

$$\bar{p}^0 = \frac{\lambda \kappa x^*}{\kappa^2 \bar{\Delta} + (1-\beta)\lambda} > 0, \tag{34}$$

$$\bar{p}^1 = \frac{\lambda}{\kappa^2 \bar{\Delta} + \lambda} \sigma_u > 0. \tag{35}$$

In both of these expressions, $0 < \bar{\Delta} \leq 1$ is defined as

$$\bar{\Delta} = 1 - \frac{\beta^2}{\theta} \frac{\lambda}{\kappa^2} |\bar{p}^1|^2. \tag{36}$$

Because a MPE solves a fixed-point problem that does not correspond to an optimization problem, depending on parameter values there may be a unique fixed point, multiple fixed points, or none at all; in the latter case, no robust linear MPE exists. Here the fixed-point problem reduces to finding values $(\bar{p}^1, \bar{\Delta})$ that satisfy the two equations (35)–(36) along with the bound (13), so that $0 < \bar{\Delta} \leq 1$. One can show that if $\lambda/\kappa^2 \geq 2$, there is a unique robust linear MPE if $\sigma_u < \hat{p}^1$, while no MPE exist

[43] It is the same as the period loss function in (19), simply written in terms of different variables, because we are now interested in the CB's state-by-state choice of π_t rather its advance choice of the coefficients p_{t-1} of a rule that will determine π_t.

[44] See Woodford (2005, sec. 4) for a generalization of this result to the case of more general linear processes for $\{u_t\}$.

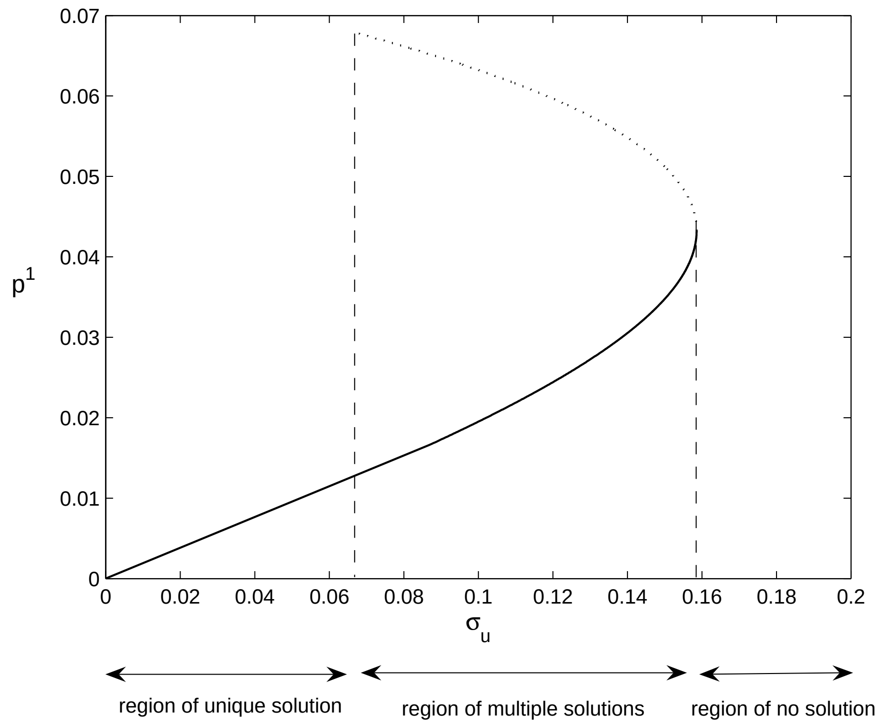

Figure 3: Varying numbers of linear MPE, depending on the size of the cost-push disturbances.

if $\sigma_u \geq \hat{p}^1$, where $\hat{p}^1$ is the upper bound on $|\bar{p}^1|$ defined in (13).[45] If instead $\lambda/\kappa^2 < 2$, then there is a unique MPE if $\sigma_u \leq \hat{p}^1$, but two distinct MPE if $\hat{p}^1 < \sigma_u < \sigma_u^*$, where

$$\sigma_u^* \equiv \frac{2}{3\sqrt{3}} \left[\frac{\theta}{\beta^2} \left(\frac{\kappa^2 + \lambda}{\lambda} \right)^3 \right]^{1/2}. \tag{37}$$

There is again a unique MPE in the special case that $\sigma_u = \sigma_u^*$, but there exist no MPE if $\sigma_u > \sigma_u^*$.[46]

The possibility of multiple solutions is illustrated numerically in Figure 3. Here the parameter values assumed are as in Figure 1, except that now $\kappa = 0.15$,[47] and I

[45]See appendix A.3 for the proof of this result and the ones stated next, and equation (A.21) in the appendix for the definition of $\hat{p}^1$.

[46]Regardless of the value of $\sigma_u > 0$, this bound will be violated in the case of small enough θ, which is to say, in the case of a large enough concern for robustness on the part of the CB.

[47]A larger value of κ is used in this example in order to illustrate the possibility of multiple

graph the locus of solutions only for the case $\theta = 0.001$. A unique solution exists for values of σ_u smaller than 0.068,[48] two solutions exist for values between 0.068 and 0.159, and no solutions exist for larger values of σ_u. In the intermediate range, the second solution (in which inflation is more sensitive to cost-push shocks) is shown by the dotted branch of the locus of fixed points. While these solutions also satisfy the above definition of a robust linear MPE, they are less appealing than the ones on the branch shown as a solid line in the figure, on grounds of what Evans and Honkapohja (2001) refer to as "expectational stability."

One can reduce the system (35) – (36) to the single equation

$$\bar{p}^1 = \Phi(\bar{p}^1), \tag{38}$$

where $\Phi(\tilde{p})$ is the value of $\bar{p}^1$ that satisfies (35), when $\bar{\Delta}$ in this equation is the value obtained by substituting $\bar{p}^1 = \tilde{p}$ in equation (36). Note that $\Phi(p_t^1)$ indicates the degree of sensitivity of inflation to cost-push shocks that would optimally be chosen by a CB choosing under discretion in period t, if it expects the sensitivity of inflation to cost-push shocks in the following period to be given by p_t^1.[49] One can show that the lower branch of solutions corresponds to fixed points at which $0 < \Phi'(\bar{p}^1) < 1$, while the upper branch corresponds to fixed points at which $\Phi'(\bar{p}^1) > 1$. Hence in the former case, an expectation that policy will be near the fixed point far in the future will justify choosing a policy very close to the fixed point now, while in the latter case, even an expectation that policy will be near that fixed point in the distant future will *not* lead the CB to choose policy near that fixed point now — only if future policy is expected to coincide *precisely* with the fixed point will similar behavior be justified now. Hence this fixed point is "unstable" under perturbations of expectations regarding future policy in a way that makes it less plausible that successive central bankers should coordinate on those particular expectations.[50]

solutions, which do not exist under the baseline calibration.

[48] As the graph suggests, there are actually two solutions to the system of equations in this region as well – the dotted branch of the locus of solutions can be extended further to the left. But for values of σ_u this small, the solutions on the dotted branch involve $\bar{\Delta} < 0$, and so do not correspond to MPE.

[49] Thus $\Phi(\cdot)$ is a mapping from the discretionary CB's "perceived law of motion" to the "actual law of motion" resulting from its optimizing decisions, in the terminology of Evans and Honkaphja (2001).

[50] One can also show that the expectationally stable MPE is an asymptotically stable rest point

What happens in the case of an economy in the region where σ_u is too large for any MPE to exist? (Note that this requires that $\sigma_u > \hat{p}^1$.) One observes that $\Phi(0) > 0$, and also that $\Phi(\hat{p}^1) = \sigma_u > \Phi(\hat{p}^1)$. Then, if there are no fixed points in the interval $(0, \hat{p}^1)$, $\Phi(p) > p$ over the entire interval.[51] This means that whatever value of p^1 may be expected to describe monetary policy in the following period, a CB that optimizes under discretion will choose a *larger* value in the current period. There is then no Markov perfect equilibrium; but the situation is clearly one in which (an attempt at) discretionary optimization would be expected to lead to very large responses of inflation to cost-push shocks — there would be no reason for the inflation response to remain within any finite bounds!

In the case of rational expectations (the limit as $\theta \to \infty$), there is always a unique solution, given by

$$\bar{p}^1 = \frac{\lambda}{\kappa^2 + \lambda}\sigma_u > 0. \tag{39}$$

This is the characterization of policy under discretion given by Clarida *et al.* (1999); the linearity in σ_u again indicates that a principle of certainty equivalence applies. Comparison with (26) indicates that under discretionary policy, inflation responds more strongly to a cost-push shock than under the optimal commitment, according to the RE analysis. Moreover, because $\bar{p}^0 > 0$ in the case of discretion, while the long-run average value of p_t^0 is zero under the optimal commitment, discretionary policy is characterized by an "inflationary bias". These discrepancies between what policy would be like in the best possible RE equilibrium and what it is like in the MPE with discretionary policy indicate the importance of advance commitment to an optimal decision procedure for monetary policy.

How are these familiar results affected by allowing for near-rational expectations? We see from (34) that whenever a robust linear MPE exists, it involves a positive average inflation rate $\pi^* = \bar{p}^0$; so again discretionary policy results in an inflationary bias. Moreover, this equation indicates that π^* is a decreasing function of $\bar{\Delta}$; hence the inflationary bias is *increased* by a concern for robustness on the part of the CB

under adaptive learning dynamics, in which a sequence of central bankers seek to forecast the policies of their successors by extrapolating observed policy in the past, while the expectationally unstable MPE will also be unstable under the learning dynamics. On the connection between expectational stability and stability under adaptive learning dynamics, see generally Evans and Honkapohja (2001).

[51] If instead there are two fixed points, the sign of $\Phi(p) - p$ changes between them; this is what makes the lower solution expectationally stable while the upper is unstable.

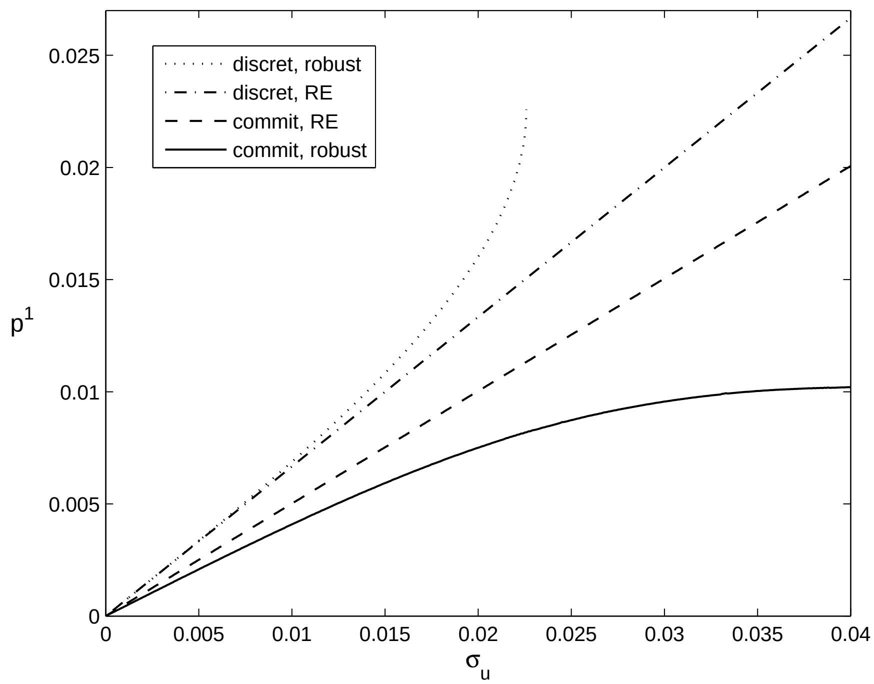

Figure 4: Variation of $\bar{p}^1$ with σ_u, under discretionary policy and under an optimal commitment, with and without allowance for near-rational expectations.

(which makes $\bar{\Delta}$ less than 1). The problem of excessive sensitivity of the inflation rate to cost-push shocks is also increased by a concern for robustness. We observe from (35) that

$$\bar{p}^1 > \frac{\lambda}{\kappa^2 + \lambda}\sigma_u \tag{40}$$

when $\bar{\Delta} < 1$, so that $\bar{p}^1$ is larger than in the RE case, described by (39). One can also show[52] that if we select the lower-sensitivity MPE as "the" prediction of the model when multiple solutions exist, then the solution for $\bar{p}^1$ is monotonically decreasing in θ over the range of values for which a robust linear MPE exists, which means that $\bar{p}^1$ is higher the greater the concern for robustness.

In the RE analysis, a discretionary policymaker allows inflation to respond more to cost-push shocks, because of her inability to commit to a history-dependent policy under which a positive cost-push shock would *reduce* subsequent inflation (as would

[52]Again see appendix A.3 for the proof.

occur under an optimal commitment). In the absence of such a commitment, inflation expectations do not move in a direction that helps to offset the effects of the disturbance on the short-run Phillips-curve tradeoff, and in the absence of such mitigation of the shift in the Phillips-curve tradeoff, it is necessary to allow inflation to respond to a greater extent. When the discretionary policymaker must guard against possible departures from RE, her situation is even more dire. Under the worst-case NRE beliefs, inflation expectations *increase* following a positive cost-push shock, precisely because this moves the Phillips curve in the direction that worsens the policymaker's tradeoff; and so the extent to which the discretionary policymaker finds it necessary to allow inflation to increase is even greater than under the RE analysis (where inflation expectations do not change).

While a concern for robustness increases the sensitivity of inflation to cost-push shocks under discretionary policy, we found in section 2 that it reduces the sensitivity to cost-push shocks under an optimal commitment. This is illustrated numerically in Figure 4, which extends Figure 1 to show how the equilibrium value of $\bar{p}^1$ varies with σ_u under discretionary policy as well as under the optimal commitment from a timeless perspective, both with and without an allowance for near-rational expectations.[53] (The two lower curves correspond to cases also shown in Figure 1.) When RE are assumed, $\bar{p}^1$ is larger under discretionary policy, as just shown; but with a concern for robustness (finite θ), the gap between the values of $\bar{p}^1$ under discretionary policy and under a robustly optimal linear policy is even larger.

Thus the distortions of policy resulting from optimization under discretion are *increased* when the CB allows for the possibility of near-rational expectations, and the lessons of the RE analysis become only more important. When the CB's concern for robustness is sufficiently small (*i.e.*, θ is large) — and when the volatility of fundamentals is sufficiently small (*i.e.*, σ_u is small) — a robust linear MPE exists, but the degree to which it involves both an excessive average rate of inflation and excessive responsiveness of inflation to cost-push shocks, relative to what would occur under the robustly optimal linear policy, is even greater than is true in the RE analysis. In the case of a sufficiently great concern for robustness, *or* a sufficiently unstable environment, a robust linear MPE fails even to exist; in this case, the dangers of discretionary policy are even more severe, and to an extent much greater than would

[53]The parameter values used in the figure are again those used in Figure 1. The RE curves assume that $\theta^{-1} = 0$, while the ones allowing for NRE beliefs assume that $\theta^{-1} = 1000$.

be suggested by the RE analysis.

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. Optimal policy is still *history-dependent* even when rational expectations are not assumed. Indeed, a concern for robustness only increases the optimal degree of history-dependence.

Moreover, just as in the RE analysis, *commitment* is important for optimal policy. The distortions predicted to result from discretionary policymaking become even more severe when the central bank allows for the possibility of near-rational expectations, so that the importance of commitment is increased. 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

Suppose that the policy commitment is of the conditionally linear form

$$\pi_{t+1} = p_t^0 + p_t^1 w_{t+1} \tag{A.1}$$

for some process $\{p_t^0(h_t, p_{-1}^0)\}$ and some deterministic sequence $\{p_t^1\}$. 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

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

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

$$\pi_t = \kappa x_t + \beta E_t[m_{t+1}\pi_{t+1}] + u_t. \tag{A.3}$$

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.2)) on the one hand, and its consequences for PS expected inflation (which affects the constraint (A.3) 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

$$E_t m_{t+1} = 1, \qquad E_t[m_{t+1}\pi_{t+1}] = \bar{\pi}_t, \tag{A.4}$$

whatever degree of distortion the PS inflation expectation $\bar{\pi}_t$ may represent. I first consider this sub-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, in which the first-order condition

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

holds in each state at date $t+1$, where ϕ_{1t}, ϕ_{2t} are Lagrange multipliers associated with the two constraints (A.4). This implies that

$$\log m_{t+1} = c_t + \phi_{2t}\pi_{t+1} \tag{A.5}$$

in each state, for some constant c_t. The two constants c_t and ϕ_{2t} in (A.5) are then the values that satisfy the two constraints (A.4).

Under the assumption of a conditionally linear policy (A.1), π_{t+1} is conditionally normally distributed, so that (A.5) implies that m_{t+1} is conditionally log-normal.[54] It follows that

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

Hence the first constraint (A.4) is satisfied if and only if

$$c_t = -\phi_{2t}p_t^0 - \frac{1}{2}\phi_{2t}^2|p_t^1|^2. \tag{A.6}$$

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 standard normal density. Hence

$$\begin{aligned}\log \tilde{f}(w) &= \log m_{t+1}(w) + \log f(w) \\ &= c_t + \phi_{2t}\pi_{t+1} - \frac{1}{2}log 2\pi - \frac{1}{2}w^2 \\ &= -\frac{1}{2}log 2\pi - \frac{1}{2}[w - \phi_{2t}p_t^1]^2,\end{aligned}$$

using (A.5) to substitute for m_{t+1} in the second line, and (A.1) and (A.6) to substitute for π_{t+1} and c_t respectively in the third line. But this is just the log density function for a variable that is distributed as $N(\mu_t, 1)$, where the bias in the perceived conditional expectation of w_{t+1} is $\mu_t = \phi_{2t}p_t^1$. Hence

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

and the second constraint (A.4) is satisfied if and only if[55]

$$\phi_{2t} = \frac{\bar{\pi}_t - p_t^0}{|p_t^1|^2}. \tag{A.7}$$

[54]This is one of the main reasons for the convenience of restricting our attention to linear policies in this paper.

[55]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.6) suffices to show that $m_{t+1} = 1$ with certainty.

Condition (A.6) then uniquely determines c_t as well, and m_{t+1} is completely described by (A.5), 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 (14), as asserted in the text.

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

$$\begin{aligned} R_t^{pess} = \hat{E}_t[\log m_{t+1}] &= c_t + \phi_{2t}\hat{E}_t\pi_{t+1} \\ &= \frac{1}{2}\frac{(\bar{\pi}_t - p_t^0)^2}{|p_t^1|^2}, \end{aligned} \tag{A.8}$$

using (A.6) and (A.7). 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 $|p_t^1|$.[56]

It remains to determine the worst-case choice of $\bar{\pi}_t$.[57] It follows from (A.3) that

$$(x_t^{pess} - x^*)^2 = \frac{1}{\kappa^2}(\pi_t - u_t - \kappa x^* - \beta\bar{\pi}_t)^2. \tag{A.9}$$

Substituting this for the squared output gap and (A.8) for the relative entropy in (A.2), 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 the inequality

$$|p_t^1|^2 < \frac{\theta}{\beta^2}\frac{\kappa^2}{\lambda}. \tag{A.10}$$

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.2), 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 (A.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

[56] Equation (A.8) 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.

[57] 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.

the objective through an extreme choice of $\bar{\pi}_t$, except in the special case that

$$p_t^0 = \beta^{-1}(\pi_t - u_t - \kappa x^*), \tag{A.11}$$

so that the linear function has a slope of exactly zero. Thus unless p_t^0 satisfies (A.11), p_t^1 must satisfy the bound (A.10) in order for the objective (A.2) to have a finite maximum. Even in the special case that (A.11) holds exactly, p_t^1 must satisfy a variant of (A.10) in which the strict inequality is replaced by a weak inequality.

When (A.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

$$\bar{\pi}_t = \Delta_t^{-1}\left[p_t^0 - (\pi_t - u_t - \kappa x^*)\frac{\beta\lambda}{\theta\kappa^2}|p_t^1|^2\right], \tag{A.12}$$

$$\Delta_t \equiv 1 - \frac{\beta^2}{\theta}\frac{\lambda}{\kappa^2}|p_t^1|^2 > 0, \tag{A.13}$$

as stated in the text. Substituting this solution into (A.8) and (A.9), one obtains the implied output gap (17) and and relative entropy (18) under the worst-case NRE beliefs, as stated in the text. Substituting these expressions into the objective (A.2), one obtains an objective for the CB of the form

$$\hat{\mathcal{L}}(p; p_{-1}^1, \rho) = \mathrm{E}\sum_{t=0}^{\infty}\beta^t L(p_{t-1}; p_t; w_t), \tag{A.14}$$

in which the period loss is given by

$$L(p_{t-1}; p_t; w_t) \equiv \frac{1}{2}\pi_t^2 + \frac{\lambda}{2\kappa^2\Delta_t}[\pi_t - u_t - \kappa x^* - \beta p_t^0]^2, \tag{A.15}$$

where $0 < \Delta_t < 1$ is the function of p_t^1 defined by (A.13), π_t is the function of p_{t-1} and w_t defined by (A.1), and $u_t = \sigma_u w_t$. Note that we can alternatively write

$$L(p_{t-1}; p_t; w_t) = \tilde{L}(\pi_t; p_t; w_t),$$

where the function $\tilde{L}$ is defined by the right-hand side of (A.15), since the coefficients p_{t-1} only enter through their consequences for the value of π_t.[58]

When, instead, (A.10) holds with equality, and (A.11) holds as well, the worst-case value of $\bar{\pi}_t$ is indeterminate, but the maximized value of (A.2) is nonetheless well-defined, and equal to zero. In this case, the period loss function is equal to

$$L(p_{t-1}; p_t; w_t) = \frac{1}{2}\pi_t^2.$$

[58]This alternative expression for the period loss function is convenient in Appendix A.3.

When neither this case nor the one discussed in the previous paragraph applies, we can define $L(p_{t-1}; p_t; s_t)$ 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 $L(p_{t-1}; p_t; s_t)$ 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.2) by choosing $m_{t+1} = 1$ in all states; the maximized value of (A.2) is then necessarily at least this high. It follows that both the conditional expectations and the infinite sum in (A.14) 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 (A.14) is well-defined for an arbitrary conditionally-linear policy $\{p_t\}$ and arbitrary initial conditions (p^1_{-1}, ρ).

A.2 Robustly Optimal Linear Policy

Given the worst-case PS beliefs characterized in the previous appendix, the problem of the CB is to choose a $\{p_t\}$ for all $t \geq 0$ so as to minimize (A.14), for given initial conditions p^1_{-1} and a distribution ρ of possible values for p^0_{-1}. The CB must choose a policy under which p^0_t may depend on both p^0_{-1} and the history of shocks h_t, but p^1_t must be a deterministic function of time.

One can show that the objective (A.14) is a convex function of the sequence $\{p_t\}$. I begin by noting that (A.2) is a convex function of π_t and x_t, for any choice of $m_{t+1}(\cdot)$. Then since (A.3) 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.2) 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.1) is linear, the value of (A.2) implied by any choice of p_t is a convex function of p_{t-1} 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.2) is also a convex function of p_{t-1} and p_t. Thus $L(p_{t-1}; p_t; w_t)$ is a convex function of (p_{t-1}, p_t). Finally, a sum of convex functions is convex; this implies that (A.14) 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,[59] according to which

$$L_3(p_{t-1}; p_t; w_t) + \beta \mathrm{E}_t L_1(p_t; p_{t+1}; w_{t+1}) = 0 \tag{A.16}$$

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

$$\mathrm{E}[L_4(p_{t-1}; p_t; w_t) + \beta L_2(p_t; p_{t+1}; w_{t+1})] = 0 \tag{A.17}$$

for each date $t \geq 0$. Here L_1 through L_4 denote the partial derivatives of $L(p^0_{t-1}, p^1_{t-1}; p^0_t, p^1_t; w_t)$ with respect to its first through fourth arguments, respectively. Condition (A.16) is the first-order condition for the optimal choice of $p^0_t(h_t, p^0_{-1}$, and (A.17) is the corresponding condition for the optimal choice of p^1_t (which must take the same value in all states of the world at date t).

Note that it follows from the characterization in the previous appendix that for any plan satisfying (A.10), the partial derivatives just referred to are well-defined, and equal to

$$L_1(p_{t-1}; p_t; w_t) = \pi_t + \frac{\lambda}{\kappa^2} \frac{\pi_t - u_t - \kappa x^* - \beta p^0_t}{\Delta_t},$$

$$L_2(p_{t-1}; p_t; w_t) = L_1(p_{t-1}; p_t; w_t) w_t,$$

$$L_3(p_{t-1}; p_t; w_t) = -\beta \frac{\lambda}{\kappa^2} \frac{\pi_t - u_t - \kappa x^* - \beta p^0_t}{\Delta_t},$$

$$L_4(p_{t-1}; p_t; w_t) = \frac{\beta^2}{\theta} \left(\frac{\lambda}{\kappa^2}\right)^2 \left(\frac{\pi_t - u_t - \kappa x^* - \beta p^0_t}{\Delta_t}\right)^2 p^1_t.$$

Substituting (A.1) for π_t and (A.13) for Δ_t in these expressions, one can express the first-order conditions (A.16) – (A.17) as restrictions upon the sequence $\{p_t\}$.

Taking as given the deterministic sequence $\{p^1_t\}$, one observes that (A.16) is a linear stochastic difference equation for the evolution of the process $\{p^0_t\}$, with coefficients that are time-varying insofar as they involve the coefficients $\{p^1_t\}$. One can show that these linear equations must have a linear solution of the form (20). Here there is no need to give a general expression for the coefficients of this solution, as

[59] This sequence of first-order conditions by itself is necessary but not sufficient for an optimum; in order to prove that a solution to the FOCs represents an optimum, one must also verify a transversality condition. Here, however, we are interested only in the steady-state solution to the FOCs, which necessarily satisfies the transversality condition. Hence the steady-state solution characterized below does represent the policy that minimizes (A.14) under the self-consistent specification of the initial conditions.

we are interested only in the existence of a steady state. In such a steady state, if it exists, p_t^1 is equal to some constant value $\bar{p}^1$ for all t; so it suffices to consider the solution to (A.16) in this case.

Under the assumption that $p_t^1 = \bar{p}^1$ for all $t \geq -1$, (A.16) is a stochastic linear difference equation for the process $\{p_t^0\}$ of the form

$$E_t[A(L)p_{t+1}^0] = (\sigma_u - \bar{p}^1)w_t, \qquad \text{(A.18)}$$

where

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

(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.18), one can easily show that (A.18) has a unique stationary solution,[60] given by

$$p_t^0 = \mu p_{t-1}^0 - \mu(\sigma_u - \bar{p}^1)w_t, \qquad \text{(A.19)}$$

where $0 < \mu < 1$ is the smaller root of the characteristic equation (24) given in the text. Note that a stationary solution exists regardless of the value assumed for $\bar{p}^1$, as long as it satisfies (A.10), for the quadratic equation is easily seen to have a root in that interval in the case of any $\bar{\Delta} > 0$. In fact, since $0 < \bar{\Delta} < 1$, one can show that $\mu^{RE} < \mu < 1$, where μ^{RE} is the root in the RE case (corresponding to $\bar{\Delta} = 1$).

The law of motion (A.19) implies that if the unconditional distribution for p_{t-1}^0 is $N(\mu_{p,t-1}, \sigma_{p,t-1}^2)$, then (given the assumption that w_t is i.i.d. $N(0,1)$) the unconditional distribution for p_t^0 is also normal, with mean and variance

$$\mu_{p,t} = \mu \mu_{p,t-1}, \qquad \sigma_{p,t}^2 = \mu^2[\sigma_{p,t-1}^2 + (\sigma_u - \bar{p}^1)^2].$$

These difference equations have a unique fixed point, corresponding to the stationary or ergodic distribution implied by the law of motion (A.19), namely,

$$\bar{\mu}_p = 0, \qquad \bar{\sigma}_p^2 = \frac{\mu^2(\sigma_u - \bar{p}^1)^2}{1 - \mu^2}.$$

I turn next to the implications of conditions (A.17). Note that for each period $t \geq 0$, the left-hand side of this equation involves the values of the three quantities $(p_{t-1}^1, p_t^1, p_{t+1}^1)$ and the unconditional joint distribution of $(p_{t-1}^0, p_t^0, p_{t+1}^0; w_t, w_{t+1})$.

[60] This is the solution to the FOCs that satisfies the transversality condition and hence that corresponds to the process that minimizes (A.14). For a generalization of this characterization to the case in which the process $\{u_t\}$ follows a more general linear process, see Woodford (2005, Appendix A.2).

Given the assumption of a normal distribution $N(\mu_{p,t-1}, \sigma^2_{p,t-1})$ for p^0_{t-1} and the law of motion (20) for $\{p^0_t\}$ under optimal policy, we can write this joint distribution as a function of the parameters $(\mu_{p,t-1}, \sigma^2_{p,t-1})$ of the marginal distribution for p^0_{t-1} and the parameters (ψ_t, ψ_{t+1}) of the conditional distribution $(p^0_t, p^0_{t+1}; w_t, w_{t+1}|p^0_{t-1})$. (Recall that ψ_t denotes the vector of coefficients of the law of motion (20).) Hence the left-hand side of (A.17) is a function of the form

$$g(p^1_{t-1}, p^1_t, p^1_{t+1}; \mu_{p,t-1}, \sigma^2_{p,t-1}; \psi_t, \psi_{t+1}),$$

as asserted in (22). Once again, we need not further discuss the form of this equation except in the case of a steady-state solution.

Using the solution above for the unconditional joint distribution of $(p^0_{t-1}, p^0_t, p^0_{t+1}; w_t, w_{t+1})$ in the case of self-consistent initial conditions, condition (A.17) then becomes a second-order nonlinear difference equation in p^1_t (the coefficients of which depend, however, on the assumed value of $\bar{p}^1$). One observes that

$$\begin{aligned} \mathrm{E}[L_4(p_{t-1}; p_t; w_t)] &= \frac{\beta^2}{\theta}\left(\frac{\lambda}{\kappa^2}\right)^2 \frac{\bar{p}^1}{\bar{\Delta}^2}\mathrm{E}[(\pi_t - u_t - \kappa x^* - \beta p^0_t)^2] \\ &= \frac{\beta^2}{\theta}\left(\frac{\lambda}{\kappa^2}\right)^2 \frac{\bar{p}^1}{\bar{\Delta}^2}[a + 2b\bar{p}^1 + (\bar{p}^1)^2], \end{aligned}$$

where

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

Similarly, one can show that

$$\begin{aligned} \mathrm{E}[L_2(p_t; p_{t+1}; w_{t+1})] &= \mathrm{E}[\pi_{t+1}w_{t+1}] + \frac{\lambda}{\kappa^2\bar{\Delta}}\mathrm{E}[(\pi_{t+1} - u_{t+1} - \kappa x^* - \beta p^0_{t+1})w_{t+1}] \\ &= \bar{p}^1 + \frac{\lambda}{\kappa^2\bar{\Delta}}[\bar{p}^1 + b]. \end{aligned}$$

Hence condition (A.17) is equivalent to

$$f(\bar{p}^1) \equiv \frac{\beta^2}{\theta}\left(\frac{\lambda}{\kappa^2}\right)^2 \frac{c}{\bar{\Delta}^2}\bar{p}^1 + \bar{p}^1 + \frac{\lambda}{\kappa^2\bar{\Delta}}[\bar{p}^1 + b] = 0, \tag{A.20}$$

where

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

A robustly optimal linear policy then exists if and only if (A.20) has a solution $\bar{p}^1$ that satisfies the bound (A.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.19), the above definitions imply that

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

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

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 (A.13) for $\bar{\Delta}$ in (A.20) 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 (A.10) can alternatively be written in the form

$$|\bar{p}^1| < \hat{p}^1 \equiv \frac{\kappa}{\lambda^{1/2}}\frac{\theta^{1/2}}{\beta}. \tag{A.21}$$

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 $p^1 \to \hat{p}^1$, the first term in the expression (A.20) 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 (A.10) and (A.20), and so describes a robustly optimal linear policy.

One can further establish that

$$0 < \bar{p}^1 < \mu\sigma_u, \tag{A.22}$$

as asserted in the text. When evaluated at the value $p^1 = \mu\sigma_u$, the second two terms in (A.20) are equal to

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

where $P(\mu)$ is the polynomial defined in (24). Moreover, in the limiting case in which $\theta \to \infty$ (the RE case), the first term in condition (A.20) is identically zero, so that $f(\mu\sigma_u) = 0$, and $\bar{p}^1 = \mu\sigma_u$ is a solution.[61] 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.17) such that (A.22) holds. If instead $\hat{p}^1 \leq \mu\sigma_u$, then (A.22) follows from the result in the previous paragraph. Hence in either case, the robustly optimal policy satisfies (A.22) for any finite θ, while the upper bound holds with equality in the limiting case of infinite θ.

Substitution of the law of motion (A.19) for p_t^0 in (A.12) leads to the solution

$$\bar{\pi}_t = \Lambda p_t^0 + \beta^{-1}(\bar{\Delta}^{-1} - 1)\kappa x^*,$$

where

$$\Lambda \equiv \bar{\Delta}^{-1} - \beta^{-1}\mu^{-1}(\bar{\Delta}^{-1} - 1).$$

Note that

$$\Lambda - 1 = (1 - \beta^{-1}\mu^{-1})(\bar{\Delta}^{-1} - 1) < 0,$$

since $0 < \beta, \mu, \bar{\Delta} < 1$, from which it follows that $\Lambda < 1$.

A.3 Existence and Stability of Robust Linear MPE

A robust linear MPE corresponds to a pair $(\bar{p}^1, \bar{\Delta})$ that satisfy equations

$$\bar{p}^1 = \frac{\lambda}{\kappa^2\bar{\Delta} + \lambda}\sigma_u > 0, \tag{A.23}$$

$$\bar{\Delta} = 1 - \frac{\beta^2}{\theta}\frac{\lambda}{\kappa^2}|\bar{p}^1|^2, \tag{A.24}$$

[61] 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).

with $\bar{\Delta} > 0$ so that (A.10) is satisfied. Equivalently, we are looking for solutions to the two equations in the interval $0 < \bar{p}^1 < \hat{p}^1$, where $\hat{p}^1$ is defined by (A.21).

If we write these equations as $\bar{\Delta} = \Delta_1(\bar{p}^1)$ and $\bar{\Delta} = \Delta_2(\bar{p}^1)$ respectively, we observe that $\Delta_1(p)$ is a decreasing, strictly concave function for all $p > 0$, while $\Delta_2(p)$ is a decreasing, strictly convex function over the same domain. Moreover, $\Delta_1(p) < \Delta_2(p)$ for all small enough $p > 0$ (as $\Delta_2(p) \to +\infty$ as $p \to 0$), and also for all large enough p (as $\Delta_1(p) \to -\infty$ as $p \to +\infty$). Hence there are either *no* intersections of the two curves with $\bar{p}^1 > 0$, or *two* intersections, or a single intersection at a point of tangency between the two curves.

The slopes of the two curves are furthermore given by

$$\Delta_1'(p) = -2\frac{\beta^2}{\theta}\frac{\lambda}{\kappa^2}p,$$

$$\Delta_2'(p) = -\frac{\lambda}{\kappa^2}\frac{\sigma_u}{p^2}.$$

From these expressions one observes that $\Delta_2'(p)$ is less than, equal to, or greater than $\Delta_1'(p)$ according to whether p is less than, equal to, or greater than $\tilde{p}^1$, where

$$\tilde{p}^1 \equiv \left(\frac{\theta}{\beta^2}\frac{\sigma_u}{2}\right)^{1/3} > 0.$$

From this it follows that there are two intersections if and only if $\Delta_2(\tilde{p}^1) < \Delta_1(\tilde{p}^1)$, which holds if and only if $\sigma_u < \sigma_u^*$, where σ_u^* is defined as in (37).[62] Similarly, the two curves are tangent to each other if and only if $\sigma_u = \sigma_u^*$; in this case, the unique intersection is at $\bar{p}^1 = \tilde{p}^1$. And finally, the two curves fail to intersect if and only if $\sigma_u > \sigma_u^*$.

It remains to consider how many of these intersections occur in the interval $0 < \bar{p}^1 < \hat{p}^1$. One notes that there is exactly one solution in that interval (and hence a unique robust linear MPE) if and only if $\Delta_2(\hat{p}^1) < 0$, which holds if and only if $\sigma_u < \hat{p}^1$. When $\sigma_u = \hat{p}^1$ exactly, $\Delta_2(\hat{p}^1) = \Delta_1(\hat{p}^1) = 0$, and the curves intersect at $\bar{p}^1 = \hat{p}^1$. This is the larger of two solutions for $\bar{p}^1$ if and only if

$$\Delta_1'(\hat{p}^1) < \Delta_2'(\hat{p}^1), \tag{A.25}$$

which holds if and only if $\lambda\kappa^2 < 2$. In this case, as σ_u is increased further, the larger of the two solutions for $\bar{p}^1$ decreases with σ_u, so that there are two solutions in the

[62]It is useful to note that this definition implies that $\sigma_u^* \geq \hat{p}^1$, with equality only if $\kappa^2/\lambda = 1/2$.

interval $(0, \hat{p}^1)$, until $\sigma_u = \sigma^*$, and the two solutions collapse into one, as the curves are tangent. (Note that $\sigma_u^* > \hat{p}^1$.) For still larger values of σ_u, there is no intersection, as explained in the previous paragraph.

If instead, $\lambda/\kappa^2 = 2$ exactly, then the curves are tangent when $\sigma_u = \hat{p}^1$ (which in this case is also equal to σ_u^*). At this point the only intersection occurs at $\bar{p}_1 = \hat{p}^1$ (which fails to satisfy condition (A.10)), and for larger values of σ_u there are no intersections. Finally, if $\lambda/\kappa^2 > 2$, then the inequality in (A.25) is reversed, and when $\sigma_u = \hat{p}^1$, the intersection at $\bar{p}^1 = \hat{p}^1$ is the smaller of the two solutions. (The smaller solution approaches $\hat{p}^1$ from below as σ_u increases to $\hat{p}^1$.) In this case, there are no solutions $\bar{p}^1 < \hat{p}^1$ when $\sigma_u = \hat{p}^1$. As σ_u increases further, the smaller solution continues to increase with σ_u, so that even for values of σ_u that continue to be less than or equal to σ_u^* (so that the curves continue to intersect), there are no solutions with $\bar{p}^1 < \hat{p}^1$. And for still larger values of σ_u, there are again no solutions at all. Hence in each case, the number of solutions is as described in the text.

The "expectational stability" analysis proposed in the text involves the properties of the map

$$\Phi(p) \equiv \Delta_2^{-1}(\Delta_1(p)).$$

Formally, a fixed point $\bar{p}^1$ of Φ (which corresponds to an intersection of the two curves studied above) is expectationally stable if and only if there exists a neighborhood P of $\bar{p}^1$ such that

$$\lim_{n\to\infty} \Phi^n(p) = \bar{p}^1$$

for any $p \in P$. Our observations above about the functions $\Delta_1(\cdot), \Delta_2(\cdot)$ imply that $\Phi(\cdot)$ is a monotonically increasing function. Hence a fixed point $\bar{p}^1$ is stable if and only if $\Phi'(\bar{p}^1) < 1$.

The above definition implies that

$$\Phi'(p) = \frac{\Delta_1'(p)}{\Delta_2'(\Delta_2^{-1}(\Delta_1(p)))} > 0.$$

Evaluated at a fixed point of Φ, this reduces to

$$\Phi'(\bar{p}^1) = \frac{\Delta_1'(\bar{p}^1)}{\Delta_2'(\bar{p}^1)}.$$

Hence the stability condition is satisfied if and only if

$$\Delta_2'(\bar{p}^1) < \Delta_1'(\bar{p}^1) < 0. \tag{A.26}$$

Because of the concavity of $\Delta_1(\cdot)$ and the convexity of $\Delta_2(\cdot)$, this condition necessarily holds at the fixed point with the smaller value of $\bar{p}^1$, and not at the higher value. Hence in Figure 3, it is the upper (dashed) branch of solutions that is expectationally unstable, while the lower (solid) branch of solutions is stable. We therefore conclude that regardless of the other parameter values, there is exactly one expectationally stable robust linear MPE for all values of σ_u below some positive critical value, and no robust linear MPE for values of σ_u greater than or equal to that value.

Finally, let us consider the way in which $\bar{p}^1$ changes as θ is reduced (indicating that a broader range of NRE beliefs are considered possible). Letting $\bar{p}^1$ be implicitly defined by the equation

$$\Delta_1(\bar{p}^1) = \Delta_2(\bar{p}^1),$$

the implicit function theorem implies that

$$\frac{d\bar{p}^1}{d\theta} = -\frac{\partial \Delta_1 / \partial \theta}{\Delta_1' - \Delta_2'}. \tag{A.27}$$

It follows from (A.26) that in the case of an expectationally stable MPE, the denominator of the fraction in (A.27) is positive. We also observe that

$$\frac{\partial \Delta_1}{\partial \theta} = \frac{\beta^2}{\theta^2} \frac{\lambda}{\kappa^2} (\bar{p}^1)^2 > 0,$$

so that the numerator is positive as well, and hence $\bar{p}^1$ decreases as θ increases. This means that $\bar{p}^1$ increases as the CB's concern for robustness increases (corresponding to a lower value of θ, up until the point where there ceases to any longer be a robust linear MPE at all. In that case, as discussed in the text, we can think of the equilibrium sensitivity of inflation to cost-push shocks as being *unbounded;* so the conclusion that greater concern for robustness leads to greater sensitivity of inflation to cost-push shocks extends, in a looser sense, to that case as well.

References

Clarida, Richard, Jordi Gali and Mark Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37(2): 1661-1707.

Dennis, Richard. 2007. "Model Uncertainty and Monetary Policy." Federal Reserve Bank of San Francisco working paper no. 2007-09, November.

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

Evans, George W., and Seppo Honkapohja. 2003. "Expectations and the Stability Problem for Optimal Monetary Policies." *Review of Economic Studies,* 70(4): 715-742.

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

Giannoni, Marc P. 2002. "Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model." *Macroeconomic Dynamics,* 6(1): 111-144.

Hansen, Lars Peter and Thomas J. Sargent. 2003. "Robust Control of Forward-Looking Models." *Journal of Monetary Economics,* 50(3): 581-604.

Hansen, Lars Peter and Thomas J. Sargent. 2005. "Robust Estimation and Control under Commitment." *Journal of Economic Theory,* 124(2): 248-301.

Hansen, Lars Peter and Thomas J. Sargent. 2007a. "Recursive Robust Estimation and Control without Commitment." *Journal of Economic Theory,* 136(1): 1-27.

Hansen, Lars Peter and Thomas J. Sargent. 2007b. *Robustness.* Princeton, NJ: Princeton University Press.

Hansen, Lars Peter, Thomas J. Sargent, Gauhar Turmuhambetova, and Noah Williams. 2006. "Robust Control and Model Misspecification." *Journal of Economic Theory,* 128(1): 45-90.

Karantounias, Anastasios G., with Lars Peter Hansen and Thomas J. Sargent.

2007. “Ramsey Taxation and Fear of Misspecification.” http://www.atl-res.com/ Tasos/job market paper.pdf.

Leitemo, Kai, and Ulf Söderström. 2008. “Robust Monetary Policy in the New-Keynesian Framework.” *Macroeconomic Dynamics,* 12(S1): 126-135.

Orphanides, Athanasios, and John C. Williams. 2005. “Imperfect Knowledge, Inflation Expectations, and Monetary Policy.” In *The Inflation Targeting Debate,* ed. Ben S. Bernanke and Michael Woodford. Chicago: University of Chicago Press.

Orphanides, Athanasios, and John C. Williams. 2007. “Robust Monetary Policy with Imperfect Knowledge.” *Journal of Monetary Economics*, 54(5): 1406-1435.

Strzalecki, Tomasz. 2008. “Axiomatic Foundations of Multiplier Preferences.” http://strzalecki.com/pdf/strzalecki-mult.pdf.

Svec, Justin. 2008. “Optimal Fiscal Policy with Robust Control.” http://www.columbia.edu/ jcs2104/research/paper.pdf.

Walsh, Carl E. 2004. “Robustly Optimal Instrument Rules and Robust Control: An Equivalence Result.” *Journal of Money, Credit and Banking,* 36(6): 1105-1113.

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

Woodford, Michael. 2005. “Robustly Optimal Monetary Policy with Near-Rational Expectations.” NBER working paper no. 11896, December.

Woodford, Michael. 2006. “An Example of Robustly Optimal Monetary Policy with Near-Rational Expectations.” *Journal of the European Economics Association,* 4(2-3): 386-395.