

Model Uncertainty, Robust Policies, and the Value of Commitment

Kenneth Kasa
Research Department
Federal Reserve Bank of San Francisco¹
e-mail: Kenneth.Kasa@sf.frb.org

November, 1998

Abstract. Using results from the literature on H^∞ -control, this paper incorporates model uncertainty into Whiteman's (1986) frequency domain approach to stabilization policy. The derived policies guarantee a minimum performance level even in the worst of (a bounded set of) circumstances.

For a given level of model uncertainty, robust H^∞ policies are shown to be more 'activist' than Whiteman's H^2 policies in the sense that their impulse responses are larger. Robust policies also tend to be more autocorrelated. Consequently, the premium associated with being able to commit is greater under model uncertainty. Without commitment, the policymaker isn't able to (credibly) smooth his response to the degree that he would like.

From a technical standpoint, a contribution of this paper is its analysis of robust control in a model featuring a forward-looking state transition equation, which arises from the fact that the private sector bases its decisions on expectations of future government policy. Existing applications of H^∞ -control in economics follow the engineering literature, and only consider backward-looking state transition equations. It is the forward-looking nature of the state transition equation that makes a frequency domain approach attractive.

JEL Classification #'s: C61, E61.

¹Please address correspondence to: Kenneth Kasa, Research Department, Federal Reserve Bank of San Francisco, P.O. Box 7702, San Francisco, CA 94120.

Model Uncertainty, Robust Policies, and the Value of Commitment

Kenneth Kasa

1 Introduction

In a pioneering analysis, Whiteman (1986) used frequency domain optimization methods to derive explicit, closed-form expressions for optimal government stabilization policies under alternative assumptions about the government's ability to precommit to its policies. He showed that when the government can precommit, policy tends to be smoother and more persistent than when it cannot precommit. Whiteman's results can also be used to show that the welfare gain from being able to commit tends to increase as the underlying shocks hitting the economy become more persistent.

Whiteman's analysis, along with the related work of Hansen, Epple, and Roberds (1985), Miller and Salmon (1985), Cohen and Michel (1988), and Kasa (1998), is a response to the famous 'Lucas Critique'. Initially, this critique engendered a deep skepticism about the applicability of optimal control methods to the formulation of government policy. (See, e.g., Prescott (1977)). After all, governments are not controlling an inanimate mechanical system, they are playing a game against forward-looking individuals, who are well aware of the constraints and motivations of the government. As a result, the government's current payoffs are a function of its anticipated *future* actions, not just its past actions, as is the case in the control literature.

Over time, however, it came to be realized that there is nothing in principle preventing the use of control methods in the design of (time-consistent) government policy. The key is to modify the government's objective function and constraints in order to remove the government's incentive to make (noncredible) promises. This can be done by eliminating the dependence of current payoffs on future actions. Hansen, Epple, and Roberds (1985) interpret this modification of the government's optimization problem by thinking of the government as consisting of a sequence of identical administrations, each of which takes the actions of subsequent administrations as given. Although analysis of this kind of model is

somewhat more complicated (e.g., standard assumptions no longer guarantee the existence of an equilibrium), the papers cited in the previous paragraph show that it can be done, and that doing so leads to some interesting comparisons between precommitment and time-consistent policies.

Besides these straightforward control-theoretic responses, the challenges of the Lucas Critique have been answered in several other ways. One of the more radical responses has been offered by McCallum (1995). He essentially denies the logic of backward induction by arguing that the government's incentive constraints by themselves will never be the source of Pareto inefficient outcomes. McCallum simply asserts that patient and farsighted policymakers are smart enough to recognize the futility of any attempt to exploit the private sector's expectations. Absent other distortions, his advice is to compute and analyze precommitment/Ramsey policies, and to ignore time-consistent policies.

Although McCallum's predictions about policy may be accurate from a purely positive standpoint, his normative admonishment to "just do it" is unconvincing. It simply discards by fiat much of modern game theory, which if given the same assumptions about policymakers' foresight and patience, is able to rationalize the same efficient policy outcomes, but in a way that respects widely accepted ground rules.

Another response to the Lucas Critique has been the pragmatic and defensive one of arguing that Lucas' advice to search for 'deep parameters' and policy-invariant econometric specifications has in practice produced models that are just as unstable as traditional reduced form Keynesian models. (See, e.g., Oliner, Rudebusch, and Sichel (1996).) While this response may be accurate given our current state of knowledge, and perhaps is satisfactory for current policymakers, it is not very satisfactory from a scientific standpoint.

To date, the most convincing critique of the Lucas Critique has been offered by Sims (1982) and Sargent (1984). They argue that Lucas' analysis is schizophrenic, or at best asymmetric. The central thought experiment in the Lucas Critique is that of a 'regime change'. In Lucas' analysis these regime changes are exogenous. Sims and Sargent argue that if the government's actions are made endogenous, and a fully symmetric game between

the private sector and the government is analyzed, then the whole issue of regime changes evaporates, and the Lucas Critique becomes irrelevant. The logic of this argument leads directly to a game-theoretic approach to government policy formulation, in which the ‘commitment technology’ becomes fully endogenous. (See, e.g., Chari and Kehoe (1990) and Stokey (1991).)

Although quite convincing, it is interesting that neither Sims nor Sargent subsequently pursued the logic of their own arguments. Fully endogenizing government policy destroys the ability of economists to offer normative advice to policymakers, and neither author seemed willing to go that far. Perhaps this was in anticipation of future difficulties encountered by a game-theoretic approach (e.g., multiple equilibria). Alternatively, it might have been that both recognized that even the most sophisticated dynamic game-theoretic models of policy miss essential aspects of real world policymaking.

Recently, an emerging literature has produced yet another response to the Lucas Critique. This literature focuses on ‘model uncertainty’ rather than commitment. It builds on recent developments in the engineering literature, which during the 1980s, made dramatic strides in analyzing control problems featuring model uncertainty. Fascinating linkages are currently being discovered between the engineer’s concept of ‘unstructured uncertainty’ (see, e.g., Zhou, Doyle, and Glover (1996)) and the economist’s concept of ‘Knightian Uncertainty’ (see, e.g., Gilboa and Schmeidler (1989) and Hansen, Sargent, and Tallarini (1997)). The goal of this literature is to devise policies that perform *adequately* (i.e., achieve a certain threshold performance level) under a wide range of circumstances. It was the contribution of Zames (1981) to recognize that the goal of guaranteeing adequate performance in the presence of model uncertainty could be formalized and made tractable simply by switching norms. His idea of analyzing traditional control problems in the H^∞ (supremum) norm rather than the standard H^2 (sum-of-squares) norm sparked a revolution in control theory.

Marcellino and Salmon (1997) argue that these recent developments in robust control theory have implications for the Lucas Critique. In their analysis individuals are uncertain about the structure of the economy, which includes government policy. In response, they

formulate decision rules that guarantee a minimum performance level, even in the worst of circumstances. Marcellino and Salmon point out that such a rule is insensitive to disturbances that lie within a predefined set of potential disturbances. To guarantee a minimum performance level, individuals base their decisions on a shock sequence that lies on the *boundary* of the feasible set, so that any sequence within the interior produces the same decision.¹

The relevance of these ideas to the Lucas Critique is immediate; to the extent that individuals are uncertain about the policy formation process, a robust decision rule becomes insensitive to a class of policy interventions. This softens the blow of the Lucas Critique.²

In a sense, this paper is the flip-side of Marcellino and Salmon (1997). Like them, I apply results from robust control theory to revisit standard issues in the analysis and design of government policy. However, here the policymaker’s perspective is adopted, not the private sector’s. Specifically, I consider a policymaker who is uncertain about the structure of the economy (in a way that cannot be captured adequately by additive disturbances with known statistical properties), only now it is the private sector that is the source of uncertainty. As in Marcellino and Salmon, the decision-maker formulates a policy that guarantees a minimum performance level. Given my reverse perspective, however, here the focus is on how model uncertainty affects the nature of government policy and the gains from precommitment, as opposed to the Lucas Critique.³

Basically, my strategy is to take an “off-the-shelf” model of dynamic policy formation, which is invariably analyzed using the H^2 -norm, and to simply re-do everything using the H^∞ -norm. I use Whiteman’s frequency domain approach as a springboard because it provides a convenient way to handle a forward-looking state transition equation. Forward-

¹Note the contrast with the usual practice of tacking on an additive disturbance with known statistical properties, and calculating optimal decision rules based on these properties. One interpretation of robust control is based on the distinction between risk (Savage expected utility is applicable) and uncertainty (Savage expected utility is inapplicable). See Gilboa and Schmeidler (1989) and Hansen, Sargent, and Tallarini (1997).

²There is an evident similarity here between robust insensitivity and the earlier literature on learning about ‘regime changes’. (See, e.g., Taylor (1975).) An advantage of a robust control perspective is that it doesn’t rely on the troublesome concept of a regime change.

³Sargent (1998a) also studies government policy formation under model uncertainty. He incorporates robustness considerations by using adaptive, constant-gain learning algorithms.

looking state transition equations are the hallmark of the time-consistency literature. The H^∞ analysis of a model with a forward-looking state transition equation is a key technical contribution of this paper, given the engineering literature’s exclusive focus on backward-looking state transition equations.

The remainder of the paper is organized as follows. The next section reviews Whiteman’s (1986) frequency domain approach to optimal policy design. The analysis takes place in H^2 . Then, with the H^2 results as a benchmark, I go on in section 3 to consider optimal stabilization policy under model uncertainty, using the H^∞ -norm. The first step in doing this is to reformulate Whiteman’s problem as a minimum norm problem. This is done in Lemma 3.1.1. The robust control problem can then be approached via the classical Nehari approximation theorem (see, e.g., Young (1988, chpt. 15)), and solved using well known interpolation methods. This is done in Lemma 3.2.1 and Theorem 3.2.1. From the ‘small gain theorem’ (Basar and Bernhard (1995, p. 16)), the inverse of the resulting H^∞ -norm can be interpreted as a measure of the range of uncertainty within which the policy is robust. As Hansen and Sargent (1998) note, time-consistency in the H^∞ case requires updating of the bound on the ℓ^2 -norm of the unstructured shocks.⁴ Thus, policy function invariance is associated with a time-varying degree of model uncertainty.

Sections 4 and 5 turn to comparisons between H^2 and H^∞ policy and value functions. As in Sargent (1998b), I find that for a given level of uncertainty, robust stabilization policy is more ‘activist’. Because of the greater demand for activism, it follows that the gains from precommitment are greater under model uncertainty and robust control. For example, in a benchmark specification of the model it turns out that time-consistent losses are 75% greater than precommitment losses in the traditional H^2 case. With model uncertainty this premium increases to 90%. Moreover, the gain differential widens as the underlying shocks become more persistent.

Section 6 concludes the paper, and offers some suggestions for future research.

⁴In the H^∞ case, this unstructured shock sequence turns out to be degenerate, with all its spectral power concentrated on a single frequency. (See, e.g., Hansen and Sargent (1998, Appendix A)). What is updated then is a scale factor measuring the ‘variance’ of this process.

2 H^2 -Optimal Stabilization Policy

This section reviews the model and the results of Whiteman (1986). To facilitate comparison between H^2 and H^∞ stabilization policy, I follow exactly Whiteman's assumptions and notation. Because of the close parallel in this section to Whiteman's analysis, the presentation will be brief. The reader should consult Whiteman's paper for full details and proofs of the following results.

The model begins with a policymaker who attempts to minimize the expected present discounted value of a loss function that trades-off variation in his instrument variable, x_t , with variation in a target variable, y_t , which is determined by the private sector:

$$\mathcal{L} = \min_{\{x_{t+j}\}} E_t \sum_{j=0}^{\infty} \beta^j [y_{t+1+j}^2 + \lambda x_{t+1+j}^2] \quad (1)$$

where λ is a positive scalar measuring the cost of instrument instability relative to target instability. Although this loss function is undeniably ad hoc, something closely resembling it can be derived in a variety of settings featuring a well-defined welfare criterion.

The analysis presumes that the economy starts at time $t = 0$ with $x_0 = 0$ and $y_0 = 0$. This is an inessential simplification when solving the model under precommitment. By assumption, the policymaker will never be able to revise his policy rule as a function of future shock realizations and initial conditions. However, in the time-consistent case, when the policymaker is allowed to re-optimize, we have to make sure that each period's initial conditions do not trigger a change in the policymaker's decision rule. Assuming that optimal policy has been pursued in the past, this can easily be done following the methods of Hansen, Epple, and Roberds (1985) and Whiteman (1986).⁵

When minimizing his loss function, the policymaker faces a constraint relating his choices of x_t to realizations of the target, y_t . Whiteman writes this constraint as follows:

$$E_t y_{t+1} = \rho y_t + x_t + e_t \quad |\rho| > 1 \quad (2)$$

⁵Hence, policies will be time-consistent, but not necessarily subgame perfect, since we do not permit 'off-equilibrium-path' deviations. Alternatively, in the language of Basar and Bernhard (1995), policies will be weakly, but not strongly, time-consistent.

where e_t is an exogenous forcing process, reflecting perhaps shifts in demand or technology. The interpretation of this constraint is that it reflects optimizing behavior by the private sector, e.g., it could be an Euler equation. It implies that agents in the private sector base their decisions on expectations of the future. This becomes clearer by iterating equation (2) forward to get:

$$y_t = -\rho^{-1} E_t \sum_{j=0}^{\infty} \left(\frac{1}{\rho}\right)^j (x_{t+j} + e_{t+j}) \quad (3)$$

so that ρ^{-1} has the interpretation of a discount rate. Because choices of y_t must be based on expectations of future policy, there is an incentive for the government to make promises, which it may not want to keep ex post.⁶

Although private sector agents do not know the future values of x_t and e_t , it is assumed (in contrast to Marcellino and Salmon (1997)) that they do know the *rules*, or stochastic processes, generating these variables. In particular, x_t and e_t are known to have the following Wold representations,⁷

$$e_t = \sum_{j=0}^{\infty} A_j u_{t-j} = A(L)u_t \quad (4)$$

$$x_t = \sum_{j=0}^{\infty} F_j u_{t-j} = F(L)u_t \quad (5)$$

The u_t sequence can be interpreted as the ‘fundamental innovations’ to the agents’ information sets. It is i.i.d and is normalized to have unit variance.

Note that while the policy function in (5) can be expressed as a closed-loop feedback from the exogenous variable (i.e., $x_t = F(L)A^{-1}(L)e_t$), it is open-loop with respect to the decisions of the private sector. On the face of it, this would seem to be suboptimal given the government’s desire to stabilize y_t . However, following Whiteman and most of the dynamic policy literature, it is assumed here that the government is a Stackelberg leader, meaning that it recognizes its influence on the private sector, while individual agents in the private

⁶Note, as Whiteman (1986, Appendix C) demonstrates, it is not necessarily the case that the government will want to renege. However, it would take a very special sequence of shocks (of measure zero) for the government not to renege.

⁷As in Whiteman, all sequences in this section are assumed to be ‘ β -summable’, so that, e.g., $\sum_{j=0}^{\infty} \beta^j A_j^2 < \infty$. Consequently, all z -transforms belong to the Hardy space of square-integrable analytic functions inside a disk of radius $\sqrt{\beta}$ centered at the origin. This space is denoted $H^2(\sqrt{\beta})$.

sector do not think of themselves as having any influence over government policy. Given that the private sector believes government policy is exogenous, the government obtains no strategic advantage from reacting directly to y_t . Of course, if in contrast agents actually believed their own individual decisions influenced government policy, then indeed there would be an advantage to specifying a rule that reacted directly to y_t . If credible, such a policy could keep agents ‘in line’.

If (4) and (5) are now plugged into (3), the Hansen-Sargent (1980) prediction formula can be applied to get the following convenient expression for y_t ,

$$y_t = \frac{L[A(L) + F(L)] - \rho^{-1}[A(\rho^{-1}) + F(\rho^{-1})]}{1 - \rho L} u_t \equiv C(L)u_t \quad (6)$$

Then, using (5) and (6) along with Parseval’s formula we get the following frequency domain representation of the government’s objective function in (1):

$$V = \min_{F(z)} \frac{(1 - \beta)^{-1}}{2\pi i} \oint [C(z)C(\beta z^{-1}) + \lambda F(z)F(\beta z^{-1})] \frac{dz}{z} \quad (7)$$

where \oint denotes contour integration around a disk of radius $\sqrt{\beta}$ centered at the origin. The policymaker’s goal is to find an analytic function, $F(z) \in H^2(\sqrt{\beta})$, which minimizes V subject to (6).

The following two subsections contain Whiteman’s solutions of this problem under the polar assumptions of perfect precommitment and no precommitment. The latter will produce time-consistent policies by construction, while the former will in general produce time-inconsistent policies.⁸

2.1 The Precommitment Case

With precommitment, the policymaker is imagined to be in business for one day. At time $t = 0$ he formulates a contingency plan for minimizing (7), with all initial conditions set to zero. Then, having devised this policy function, he programs it into his computer and retires.

⁸See Kasa (1998) for an analysis of intermediate cases, where the government has an arbitrary, but fixed, precommitment horizon of n -periods. See Roberds (1987) for an analysis of random precommitment.

The solution to this once-in-a-lifetime policy design problem is given by Theorem 1 in Whiteman (1986),

Proposition 2.1.1 (Whiteman (1986), Theorem 1): *When the government can precommit, the z -transform of the optimal policy function, obtained by minimizing (7) subject to (6), is given by:*

$$F^p(z) = -\frac{\beta}{\gamma(1-\theta z)} \left[\frac{A(z)}{(1-\beta\theta z^{-1})} \right]_+ \quad (8)$$

where $[\cdot]_+$ is an annihilation operator, meaning “ignore negative powers in z ”, and where γ and θ are determined by the spectral factorization equation:

$$\gamma(1-\theta z)(1-\beta\theta z^{-1}) = \beta + \lambda(1-\rho z)(1-\beta\rho z^{-1}) \quad |\theta| < \beta^{-1/2} \quad (9)$$

Plugging (8) into (6) and simplifying using the relationship in (9) between (θ, γ) and (λ, ρ) yields the following feedback policy rule,

$$x_t = (\beta\rho)^{-1}x_{t-1} + (\lambda\rho)^{-1}y_t \quad (10)$$

There are three noteworthy features of this feedback policy. First, in general it is time-inconsistent. If allowed to re-optimize in the future, the policymaker would almost certainly want to revise his policy in response to intervening shocks. Second, it is independent of $A(L)$. Although the univariate time-series properties of x_t and y_t depend on the stochastic properties of the exogenous shocks, the equilibrium relationship between x_t and y_t is independent of the stochastic properties of e_t . Third, as Whiteman notes, the autocorrelated response of x_t to y_t reflects a sort of ‘intertemporal substitution’ due to costs of instrument instability. (Note that if $\lambda = 0$, then optimal policy is clearly $x_t = -e_t$. This eliminates variability in y_t , and thus attains the bliss point in every period.) As we shall see in section 2.2, this kind of smoothing and intertemporal substitution is only feasible when the government can commit to its policy.

Plugging (8) and (10) into (7), and then simplifying using (9), delivers the following minimized loss function,

Proposition 2.1.2 (Whiteman (1986), Theorem 3): *When the government can precommit to its policies, the z -transform of the minimized loss function is given by:*

$$\mathcal{L}_{\min}^p = \frac{1}{1-\beta} \frac{\theta\beta/\rho}{2\pi i} \oint \left(\frac{zA(z) - \theta\beta A(\theta\beta)}{z - \theta\beta} \right) \left(\frac{\beta z^{-1} A(\beta z^{-1}) - \theta\beta A(\theta\beta)}{\beta z^{-1} - \theta\beta} \right) \frac{dz}{z} \quad (11)$$

2.2 The Time-Consistent Case

Following Hansen, Epple, and Roberds (1985), now assume the government consists of a sequence of ‘administrations’. Each administration is in office for a single period, and has no ability to compel future administrations to adhere to its policies. Although in office for only a single period, each administration still cares about the entire present discounted value of future losses, so that it considers the future consequences of its current actions. When looking to the future, each administration believes that subsequent administrations will be exactly like itself, so that despite the process of administration turnover, the government’s objectives remain constant over time.

To compute a time-consistent policy, the terms $E_t x_{t+n}$ for $n > 0$ that appear in the constraint in (3) must be held constant when deriving the policymaker’s first-order condition. This will remove the incentive to make promises. In equilibrium, of course, the elements of $F(L)$ that are held constant must match the corresponding elements that are optimized, but this equilibrium condition isn’t imposed until *after* the first-order condition has been derived. A time-consistent policy is invariant to re-optimization because, by construction, initial conditions always satisfy the first-order condition. (See Hansen, Epple, and Roberds (1985) or Whiteman (1986) for details.)

Applying this procedure generates the following time-consistent policy function,

Proposition 2.2.1 (Whiteman (1986), Theorem 2): *When the government cannot precommit to its policies, the z -transform of the optimal policy function is:*

$$F^c(z) = -\frac{1}{(1 + \lambda\rho^2)} \left[\frac{A(z)}{(1 - \phi z^{-1})} \right]_+ \quad (12)$$

where $\phi = \lambda\rho/(1 + \lambda\rho^2)$.

Plugging (12) into (6) now yields the following time-consistent feedback policy rule,

$$x_t = (\lambda\rho)^{-1}y_t \quad (13)$$

Compared to the (time-inconsistent) precommitment policy in (10), it is clear that the ability to commit leads to a more gradual response to shocks. Finally, the minimized loss is:

$$\mathcal{L}_{\min}^c = \frac{1}{1-\beta} \frac{\phi/\rho}{2\pi i} \oint \left(\frac{zA(z) - \phi A(\phi)}{z - \phi} \right) \left(\frac{\beta z^{-1}A(\beta z^{-1}) - \phi A(\phi)}{\beta z^{-1} - \phi} \right) \frac{dz}{z} \quad (14)$$

This completes our whirlwind tour of Whiteman's H^2 analysis of optimal stabilization policy. In the next section, the policy and value functions in equations (8), (11), (12), and (14) will serve as benchmarks for our H^∞ analysis of robust stabilization policy.

3 H^∞ -Optimal Stabilization Policy

The best way to think about robust control is that it just uses a different norm to compute policy functions. Traditional H^2 -control evaluates policies using a sum-of-squares metric. In the frequency domain this translates into minimizing the area under a spectral density. The problem with this strategy is that it exposes the decision-maker to potentially unbounded losses should his policy function be applied to the 'wrong' model.

To avoid potentially huge losses in the presence of model uncertainty, H^∞ -control adopts a minmax perspective, which is implemented via the supremum norm. In the frequency domain this translates into minimizing the maximum value of a spectral density. This puts a cap on potential losses.

This section develops these ideas in some detail. It proceeds in three steps. First, I reformulate Whiteman's stabilization problem as a minimum norm, or 'model-matching', problem. This involves using the policy and value functions in section 2 to back-solve for a two-sided L^2 function. This function serves as a target, which the one-sided H^2 policy function tries to approximate. The target function is selected so that the solution generates the same policy and value functions as Whiteman's. There is a separate, though closely related, target function for each of the precommitment and time-consistent cases. The

second step is to re-solve these two model-matching problems using the H^∞ supremum norm. There are many ways to do this, state-space methods being the most general and powerful (see, e.g., Zhou, Doyle, and Glover (1996)). However, state-space methods are designed to achieve only an approximate solution to the minimum norm problem, and in the present univariate context, there is a convenient and revealing method for actually obtaining closed-form expressions for the exact solution. To do this we must formulate the problem as one of minimum norm interpolation. This is done in section 3.2. Finally, the last step is to implement this solution methodology for the simple case of $AR(1)$ shocks. This solution will be used in sections 4 and 5 to compare H^2 and H^∞ stabilization policy.

3.1 H^2 -Control as a Minimum-Norm Problem

The following lemma summarizes the results from this back-solving exercise,

Lemma 3.1.1: *The stabilization problems analyzed in Section 2 can be viewed equivalently as the following minimum norm problem in $H^2(\sqrt{\beta})$:*

$$\mathcal{L}_2 = \min_{F_2(z)} \|T(z) - g(z)F_2(z)\|_2^2 = \min_{F_2(z)} \frac{1}{2\pi i} \oint |T(z) - g(z)F_2(z)|^2 \frac{dz}{z} \quad (15)$$

with $T(z)$ given by:

$$T(z) = \sqrt{\frac{\xi}{\rho(1-\beta)}} \left\{ \frac{zA(z) - \xi A(\xi)}{z - \xi} + \frac{\beta z^{-1}A(\beta z^{-1}) - \xi A(\xi)}{z - \xi} \right\} \quad (16)$$

The analytic function, $g(z)$, and the scalar, $|\xi| < \beta^{-1/2}$, depend on whether or not the government can precommit, and $g(z)$ contains no zeros inside the unit circle.

Proof: First, from standard Hilbert space theory (see, e.g., Young (1988, p. 188)), the solution to problem (15) is given by $\hat{F}_2(z) = g^{-1}(z)[T(z)]_+$. Second, from (16) we have,

$$[T(z)]_+ = \sqrt{\frac{\xi}{\rho(1-\beta)}} \left(\frac{zA(z) - \xi A(\xi)}{z - \xi} \right)$$

Third, from Hansen and Sargent (1980, Appendix A), we know for example that $[A(z)/(1 - \theta\beta z^{-1})]_+ = [zA(z) - \theta\beta A(\theta\beta)]/[z - \theta\beta]$. Thus, $\hat{F}_2(z)$ will equal the function $F^p(z)$ given by equation (8) if $g(z) = -\gamma(1 - \theta z)\sqrt{\rho(1-\beta)}/\xi$ and $\xi = \theta\beta$. Alternatively, if we set

$g(z) = -(1 + \lambda\rho^2)\sqrt{\rho(1 - \beta)}/\xi$ and $\xi = \phi$, then we obtain the function $F^c(z)$ given by equation (12).

Next, plugging $\hat{F}_2(z)$ back into (15) yields the minimized loss function,

$$\mathcal{L}_{\min} = \frac{1}{2\pi i} \oint |T(z) - [T(z)]_+|^2 \frac{dz}{z} = \frac{1}{2\pi i} \oint |[T(z)]_-|^2 \frac{dz}{z} \quad (17)$$

Again from (16), note that

$$[T(z)]_- = \sqrt{\frac{\xi}{\rho(1 - \beta)}} \left(\frac{\beta z^{-1} A(\beta z^{-1}) - \xi A(\xi)}{z - \xi} \right)$$

Thus, setting $\xi = \theta\beta$ makes (17) equivalent to the precommitment value function in (11), while setting $\xi = \phi$ makes (17) equivalent to the time-consistent value function in (14).⁹

Thus, depending on $g(z)$ and ξ , the minimum norm problem in (15) produces the same policy functions ($F^p(z), F^c(z)$) and the same value functions ($\mathcal{L}_{\min}^p, \mathcal{L}_{\min}^c$) as the stabilization problems in section 2. \square

The construction here is not the only, or perhaps even the most obvious, method of nesting the H^2 and H^∞ control problems. An alternative strategy would be to apply the spectral factorization theorem to (7) and write the objective function as,

$$\min_F \frac{\sigma^2}{2\pi i} \oint G(z)G(\beta z^{-1}) \frac{dz}{z}$$

where $G(z)$ and σ^2 solve $\sigma^2 G(z)G(\beta z^{-1}) = C(z)C(\beta z^{-1}) + \lambda F(z)F(\beta z^{-1})$. Model uncertainty could then be incorporated via unstructured perturbations of the transfer function $G(L)$, and the robust H^∞ -control problem would call for the minimization of the operator norm of G with respect to F . This problem could in turn be formulated as a ‘minimum entropy’ problem, which from Mustafa and Glover (1990), nests both the H^2 and H^∞ problems.

From a computational standpoint, however, it turns out that in a univariate context a model-matching approach is more convenient. From this perspective, model uncertainty is associated with (two-sided) additive perturbations of $T(z)$ in equation (15). It should be noted, however, that the two-sidedness of $T(z)$ makes a model-matching approach inherently unsuited to the analysis of the effects of initial conditions and re-optimization. That’s why a separate back-solving problem must be solved for the precommitment and time-consistent

⁹Remember that due to discounting the squared modulus of a complex-variable function, $h(z)$, is defined to be $|h(z)|^2 = h(z)h(\beta z^{-1})$.

cases. Verifying time-consistency in the H^∞ case is done indirectly, and will exploit the previously noted nesting properties of the minimum entropy controller.

3.2 H^∞ -Control as a Minimum-Norm Interpolation Problem

Having formulated the model-matching problem in (15), the analysis of the robust control problem becomes straightforward, at least conceptually. All we do now is solve the minimum norm problem in (15) using a different norm, i.e., the H^∞ supremum norm. Specifically, we want to solve the following problem,

$$\mathcal{L}_\infty = \inf_{F_\infty(z)} \|T(z) - g(z)F_\infty(z)\|_\infty = \inf_{F_\infty(z)} \sup_{|z|=1} |T(z) - g(z)F_\infty(z)|^2 \quad (18)$$

Again, the idea here is that even an H^2 -minimizer might be interested in solving this problem if there is uncertainty about $T(z)$, which could reflect uncertainty about $A(z)$ or ρ . For example, suppose the baseline or *nominal* model $T^n(z)$ is replaced by the *actual* model, $T^a(z) = T^n(z) + \Delta(z)$, where $\Delta(z)$ is some bounded analytic function in an annulus around the unit circle. Applying the original policy function, $[T^n]_+$, now produces the actual minimized loss function,

$$\mathcal{L}_2^a = \mathcal{L}_2^n + \|\Delta\|_2^2 + \frac{2}{2\pi i} \oint \left(\frac{zA(z) - \xi A(\xi)}{\beta z^{-1} - \xi} \right) \Delta(z) \frac{dz}{z} \quad (19)$$

where \mathcal{L}_2^n denotes the original loss function in (17) associated with the nominal $T^n(z)$. Notice that the last term in (19) could be quite large, depending on how the unstable poles of $\Delta(z)$ interact with ξ and the poles of $A(z)$.

In contrast, suppose we construct a policy $\hat{F}_\infty(z)$ that solves the problem in (18), and (as it will turn out) that $\max |T^n(z) - g(z)\hat{F}_\infty(z)|^2 = k^2$. Then it can be shown that $(k^2 + \|\Delta\|_\infty)$ provides an upper bound on \mathcal{L}_2^a . Although $\mathcal{L}_2^n \leq k^2$ and $\|\Delta\|_2^2 \leq \|\Delta\|_\infty$, it may well turn out that the last term in (19) dominates, so even by the H^2 criterion the policymaker is better off solving (18).

The following lemma is the key to obtaining an analytical solution to this problem.

Lemma 3.2.1: *Denote the unstable poles of the function $T(z)$ defined in (16) by $p_i, i =$*

1, 2, \dots, P. Using these poles, construct the Blaschke product,

$$B(z) = \prod_{i=1}^P \frac{|p_i|}{p_i} \frac{z - p_i}{1 - \bar{p}_i z}$$

where $|p_i| < 1$ by definition. Then the H^∞ minimum norm problem in (18) can be formulated equivalently as the following minimum norm interpolation problem – Find an analytic function $\varphi(z) \in H^\infty$ of minimum H^∞ norm that satisfies the P interpolation constraints

$$\varphi(p_i) = \tilde{T}(p_i)$$

where,

$$\tilde{T}(z) \equiv T(z)B(z) \tag{20}$$

Proof: First, define the function $\tilde{F}(z) \equiv g(z)F_\infty(z)$, and restate the problem in (18) as,

$$\mathcal{L}_\infty = \min_{\tilde{F}(z) \in H^\infty} \|T(z) - \tilde{F}(z)\|_\infty \tag{21}$$

Given a solution for $\tilde{F}(z)$ we can obtain the solution to the original problem by $F_\infty(z) = g^{-1}(z)\tilde{F}(z)$, since $g(z)$ has no zeros inside the unit circle.

Notice that (21) calls for minimizing the H^∞ distance between a two-sided L^∞ function and a one-sided H^∞ function. A general existence and uniqueness proof, which relates the solution to the Hankel norm of T , is provided by Nehari’s Theorem (See, e.g., Young (1988, p. 190).) Given our univariate set-up, however, an easier route is to reformulate it as an interpolation problem as follows.

Since Blaschke products have a modulus of unity on the unit circle, (21) can in turn be restated as follows,

$$\mathcal{L}_\infty = \min_{\tilde{F}(z) \in H^\infty} \|\tilde{T}(z) - B(z)\tilde{F}(z)\|_\infty \tag{22}$$

where $\tilde{T}(z)$ is defined in (20). Now let,

$$\tilde{F}(z) = \frac{\tilde{T}(z) - \varphi(z)}{B(z)} \tag{23}$$

Notice that $\tilde{F}(z) \in H^\infty$ since the interpolation conditions make the numerator in (23) vanish at the zeros of $B(z)$. Finally, plugging (23) into (22) gives us the problem,

$$\mathcal{L}_\infty = \min_{\varphi \in H^\infty} \|\varphi(z)\|_\infty \quad \text{s.t.} \quad \varphi(p_i) = \tilde{T}(p_i) \tag{24}$$

An optimal $\varphi(z)$ can then be transformed into an optimal $F_\infty(z)$ using (23) and the definition of $\tilde{F}(z)$. \square

As it turns out, the problem in (24) is an example of a general class of problems, the analysis of which is a well developed branch of the theory of H^p spaces (see, e.g., Duren (1970, chpt. 8)). Its solution is given by the following theorem,

Theorem 3.2.1: *The solution of the minimum norm interpolation problem in (24) is given by,*

$$\hat{\varphi}(z) = k \prod_{i=1}^{P-1} \frac{z - \psi_i}{1 - \bar{\psi}_i z} \quad (25)$$

where the scalars k and ψ_i are determined by the simultaneous nonlinear equations,

$$k \prod_{i=1}^{P-1} \frac{p_i - \psi_i}{1 - \bar{\psi}_i p_i} = \tilde{T}(p_i) \quad i = 1, 2, \dots, P \quad (26)$$

Proof: See, e.g., Chui and Chen (1997, p. 103).

Corollary 3.2.1: *The minimum norm problem in (18) has the constant value k^2 , where k depends on commitment ability and is derived from the solution of the equations in (26).*

Proof: First, from Lemma 3.2.1 and Theorem 3.2.1, the solution of (18) is implicit in the solutions of (25) and (26), with an optimized value of $\|\hat{\varphi}(z)\|_\infty$. Second, by the definition of the H^∞ norm, and the fact that an analytic function attains its maximum on the boundary of its domain, $\|\hat{\varphi}(z)\|_\infty = k^2$, since the terms $(z - \psi_i)/(1 - \bar{\psi}_i z)$ all have a modulus of one on the unit circle.

The fact that k^2 equals the H^∞ norm implies that in the nonlinear equations characterizing ψ_i , we must select the roots that produce the smallest (modulus of) k . \square

Thus, the H^∞ robust controller produces an ‘all-pass’ transfer function, i.e., the spectral density of the model-matching error is flat, equaling k^2 at all frequencies. This is intuitive. A minmax decision-maker would always be willing to accept a little more variance at frequencies where the spectral density is relatively low in exchange for a reduction in variance at frequencies where it is relatively high.

Time-Consistency of H^∞ -Control Policies

Remember, in deriving these policy functions, the initial shock, u_0 , has been normalized to zero. At time $t = 1$, however, it is almost certainly the case that a nonzero u_1 will

have been realized. How can we be sure that this realization will not trigger a change in policy? In the H^2 case, section 2.2 outlined a methodology that will deliver this invariance. Unfortunately, this method is specific to an H^2 objective, and we cannot expect it to apply to the H^∞ case.

Following Hansen and Sargent (1998), the key to guaranteeing time-consistency in H^∞ -control problems is to incorporate initial conditions into the model's unstructured uncertainty. The idea is as follows.

Associated with every policy function is a worst-case sequence of disturbances. A robust policy minimizes the potential damage caused by this worst-case shock sequence. To produce a sensible (i.e., bounded) result, these shocks must clearly be bounded in some way. With precommitment, this bound turns out to be just a constant scale factor, which does not influence any decisions. However, if re-optimization is allowed, then since the realized initial conditions will in general *change* the worst-case shock sequence, the policymaker may want to reconsider his policy.

Hansen and Sargent (1998) and Hansen, Sargent, and Tallarini (1997) show how to incorporate the realized initial conditions into the bound on the unstructured shocks in a way that preserves the original worst-case shocks, and hence, preserves the original robust policy function. One way to interpret this recursive updating of the bound on the shocks is that it represents an evolving degree of model uncertainty. This makes sense. As time unfolds, one would expect the degree of model uncertainty to change.¹⁰

Although a formal proof of time-consistency is not offered here, the main ingredients of such a proof are as follows. First, as noted earlier, the minimum entropy approach of Mustafa and Glover (1990) nests the H^2 and H^∞ -control problems.¹¹ This approach introduces a free parameter into the decision-maker's objective function, which can be interpreted either as a Lagrange Multiplier on the constraint bounding the unstructured shocks, or as an upper bound on the problem's H^∞ -norm. As this parameter goes to infinity, the minimum

¹⁰Note, however, that we are explicitly ruling out purposeful learning and adaptive control. See below for further discussion.

¹¹Hansen and Sargent (1998) extend their results to incorporate discounting.

entropy problem converges to the H^2 problem. Alternatively, the minimum value consistent with the existence of a solution replicates the solution of the H^∞ -control problem. Second, we know that by construction the H^2 policy (derived without precommitment, of course) is time-consistent. Hence, the time-consistency of any minimum entropy control policy can be maintained by exploiting the trade-off between the constraint on the unstructured shocks and its associated Lagrange Multiplier. The third ingredient is to then drive this entropy parameter down to its H^∞ limit, and verify that continuity is maintained. This can be tricky in multivariate problems, since it depends on whether a positive-definiteness or a rank condition is violated first, but in a univariate setting it is relatively straightforward. (See Zhou, Doyle, and Glover (1996, p. 439).)

It is interesting to relate this discussion of time-consistency back to Marcellino and Salmon's (1997) analysis of the Lucas Critique. Remember, they attribute model uncertainty to the private sector, and point out that robust private sector decision rules mitigate the Lucas Critique, since they remain invariant to a class of policy interventions. In a sense, this discussion of time-consistency is the flip-side of this insight, i.e., viewed from the government's perspective, model uncertainty reduces the importance of precommitment. To the extent that realized shocks lie within the domain of model uncertainty, re-optimization becomes less of an issue. However, one must be careful to not over-generalize this intuition, since there are other factors at play influencing the gains from precommitment. In fact, as we shall see, it turns out that once we account for differences in optimal policies, the (relative) gain from precommitment actually increases with model uncertainty.

3.3 An Example: AR(1) Shocks

An advantage of a frequency domain approach to policy design is its ability to handle general specifications of the shock dynamics. However, if we are to continue in the realm of pencil-and-paper analysis, we must adopt relatively simple specifications for the underlying shocks, e_t . For example, notice from (16) that if e_t is an $AR(s)$ process then $T(z)$ has $s + 1$ unstable poles. Then, from (25) and (26), deriving the policy function requires the solution

of s simultaneous polynomials of order s . Clearly, with only a pencil and paper at hand, this becomes intractable for $s > 2$. Accordingly, to keep things as simple as possible, while retaining the ability to say something about how shock persistence affects the analysis, in the remainder of the paper I assume e_t follows an $AR(1)$ process, so that $A(L) = (1 - \alpha L)^{-1}$.

The following proposition and corollary summarize the results in the $AR(1)$ case.

Proposition 3.3.1: *Assume the exogenous shocks follow the $AR(1)$ process, $e_t = \alpha e_{t-1} + u_t$.*

Then the z -transform of the robust precommitment policy function is:

$$F_\infty^p(z) = -\frac{\beta}{\gamma(1 - \alpha\theta\beta)} \left[\frac{(1 - \psi^p z) + (1 - \alpha z)(1 - r^p z)}{(1 - \alpha z)(1 - \theta z)(1 - \psi^p z)} \right] \quad (27)$$

and the the z -transform of the robust time-consistent policy function is:

$$F_\infty^c(z) = -\frac{1}{\gamma(1 - \alpha\phi)} \left[\frac{(1 - \psi^c z) + (1 - \alpha z)(1 - r^c z)}{(1 - \alpha z)(1 - \psi^c z)} \right] \quad (28)$$

where r^p and r^c are determined by the equations:

$$1 - r^p \psi^p = \frac{(1 - (\psi^p)^2)(\beta - \theta\beta\psi^p)}{(\psi^p - \alpha\beta)(\psi^p - \theta\beta)} \quad 1 - r^c \psi^c = \frac{(1 - (\psi^c)^2)(\beta - \phi\psi^c)}{(\psi^c - \alpha\beta)(\psi^c - \phi)}$$

and where the scalars ψ^p and ψ^c are determined by the precommitment and time-consistent solutions for the $AR(1)$ version of the equations in (26), so that $P = 2$.

Proof: From (16), when $A(z) = (1 - \alpha z)^{-1}$ then $T(z)$ has two unstable poles. One is at $z = \alpha\beta$. The other is at $z = \theta\beta$ with precommitment, and at $z = \phi$ with time-consistency.

From (26), the interpolation constraints become:

$$k \frac{\alpha\beta - \psi_i}{1 - \psi_i\alpha\beta} = \tilde{T}(\alpha\beta) \quad (29)$$

$$k \frac{\xi - \psi_i}{1 - \psi_i\xi} = \tilde{T}(\xi) \quad (30)$$

where ξ (and therefore the resulting ψ_i) depends on commitment as before, and \tilde{T} is defined by (16) and (20). Dividing (29) by (30) yields the the following quadratic equation for the ψ_i :

$$\left(\frac{\alpha\beta - \psi_i}{1 - \psi_i\alpha\beta} \right) \left(\frac{1 - \psi_i\xi}{\xi - \psi_i} \right) = \beta \frac{(1 - \alpha\xi)(1 - \xi^2)}{(1 - \alpha^2\beta^2)(\beta - \xi^2)}$$

If $|\alpha| < 1$ then k increases with ψ , so the smaller of the two roots must be selected. Next, solving for k in terms of ψ_i yields,

$$k = \tilde{T}(\alpha\beta) \frac{1 - \psi_i\alpha\beta}{\alpha\beta - \psi_i} = \sqrt{\frac{\xi}{\rho(1 - \beta)}} \left(\frac{\beta}{1 - \alpha\xi\beta} \right) \left(\frac{1}{1 - \alpha^2\beta^2} \right) \left(\frac{1 - \psi_i\alpha\beta}{\alpha\beta - \psi_i} \right) \quad (31)$$

Finally, from (23) and (26),

$$\tilde{F}(z) = \left[\tilde{T}(z) - k \frac{z - \psi_i}{1 - \psi_i z} \right] \left[\frac{(1 - \alpha\beta z)(1 - \xi z)}{(z - \alpha\beta)(z - \xi)} \right]$$

Upon plugging in for k and \tilde{T} it is seen the $(z - \alpha\beta)$ and $(z - \xi)$ terms can be factored out. Collecting terms and using the definitions of \tilde{F} and $g(z)$ given in lemma 3.1.1 yields equations (27) and (28). \square .

Corollary 3.3.1: *Assume the exogenous shocks follow the AR(1) process, $e_t = \alpha e_{t-1} + u_t$. Then in the precommitment case the H^∞ value function is given by:*

$$(k^p)^2 = \frac{\theta\beta}{\rho(1-\beta)} \left[\left(\frac{\beta}{1-\alpha\theta\beta^2} \right) \left(\frac{1}{1-\alpha^2\beta^2} \right) \left(\frac{1-\psi^p\alpha\beta}{\alpha\beta-\psi^p} \right) \right]^2 \quad (32)$$

and in the time-consistent case it is given by:

$$(k^c)^2 = \frac{\phi}{\rho(1-\beta)} \left[\left(\frac{\beta}{1-\alpha\phi\beta} \right) \left(\frac{1}{1-\alpha^2\beta^2} \right) \left(\frac{1-\psi^c\alpha\beta}{\alpha\beta-\psi^c} \right) \right]^2 \quad (33)$$

Proof: Follows directly from equation (31), with appropriate choice of ξ and ψ_i . \square

From inspection of equations (27) and (28) we see that the robust precommitment policy implies x_t is governed by an $ARMA(3,3)$ process, while the robust time-consistent policy implies x_t follows an $ARMA(2,3)$ process. In the next section these policy functions are compared to Whiteman's H^2 policy functions.

Before doing this, however, it is useful to visualize what's going on by plotting out frequency decompositions of the model-matching errors for both the H^2 and H^∞ cases. In the H^2 case we get:

$$|T(z) - g(z)F_2(z)|^2 = \frac{\xi}{\rho(1-\beta)} \left(\frac{1}{1-\alpha\xi} \right)^2 \frac{1}{1 + \alpha^2\beta - 2\alpha\sqrt{\beta}\cos\omega}$$

where ω denotes frequency measured in radians. Of course, in the H^∞ case the transfer function of the model-matching error is 'all-pass' by design, i.e., the error is the same at all frequencies. Its magnitude is given by (32) in the precommitment case, and by (33) in the time-consistent case.

Figures 1a and 1b compare these decompositions in the precommitment and time-consistent cases, respectively. The parameters are set at $\beta = .95$, $\alpha = 0.7$, $\lambda = 1.0$, and $\rho = 1.1$. There

are three points to notice. First, not surprisingly, since $\phi > \theta\beta$ and $\psi^c > \psi^p$, time-consistent losses exceed precommitment losses in both the H^2 and H^∞ cases. Second, it is clear from the figures that the area under the H^2 curves is less than the area under the H^∞ curves. Without model uncertainty, you are obviously better off with the H^2 policy. However, the third point to notice is that the H^2 policy is sensitive to low frequency misspecifications, which interact with the low frequency of the exogenous shocks. In contrast, the H^∞ policy immunizes the policymaker against these misspecifications, but at the cost of introducing high frequency noise.

4 H^2 vs. H^∞ Comparisons: Policy Functions

Referring back to equations (8) and (12), we see that when $A(L) = (1 - \alpha L)^{-1}$ the H^2 policy functions become:

$$F_2^p(z) = -\frac{\beta}{\gamma(1 - \alpha\theta\beta)} \left[\frac{1}{(1 - \theta z)(1 - \alpha z)} \right]$$

$$F_2^c(z) = -\frac{1}{\gamma(1 - \alpha\phi)} \left[\frac{1}{1 - \alpha z} \right]$$

Thus, with precommitment x_t is $ARMA(2,0)$, and with time-consistency x_t is $ARMA(1,0)$. The interesting point to notice is that these are the same as the first terms of the H^∞ policy functions given in (27) and (28), since the $(1 - \psi_i z)$ factors cancel out. This allows us to relate the H^2 and H^∞ policy functions as follows:

$$F_\infty^p(z) = F_2^p(z) - \frac{\beta}{\gamma(1 - \alpha\theta\beta)} \left[\frac{1 - r^p z}{(1 - \theta z)(1 - \psi^p z)} \right]$$

$$F_\infty^c(z) = F_2^c(z) - \frac{1}{\gamma(1 - \alpha\phi)} \left[\frac{1 - r^c z}{1 - \psi^c z} \right]$$

Two differences are clear. First, evaluating at $z = 0$ shows that the H^∞ policies have larger initial impulse responses. In this sense, a robust policy is more ‘aggressive’, as it attempts to countervail potential low frequency misspecifications. Second, as long as ψ^p and ψ^c are positive, it turns out that robust policies are more ‘persistent’ as well. It is this latter characteristic that influences the relative gains from precommitment, to which we now turn.

5 H^2 vs. H^∞ Comparisons: Gains to Commitment

Referring back to equations (11) and (14), we see that when $A(L) = (1 - \alpha L)^{-1}$ the percentage ‘gains to commitment’ in the H^2 case are:

$$\frac{\mathcal{L}_{\min}^c}{\mathcal{L}_{\min}^p} - 1 = \frac{\phi}{\theta\beta} \left(\frac{1 - \alpha\theta\beta}{1 - \alpha\phi} \right)^2 - 1 \quad (34)$$

and in the H^∞ case they are:

$$\left(\frac{k^c}{k^p} \right)^2 - 1 = \frac{\phi}{\theta\beta} \left[\left(\frac{1 - \alpha\theta\beta^2}{1 - \alpha\phi\beta} \right) \left(\frac{1 - \psi^c\alpha\beta}{1 - \psi^p\alpha\beta} \right) \left(\frac{\alpha\beta - \psi^p}{\alpha\beta - \psi^c} \right) \right]^2 - 1 \quad (35)$$

Figures 2 and 3 plot equations (34) and (35) for alternative values of α and λ , respectively. The remaining parameters are set as before, i.e., $\beta = .95$ and $\rho = 1.1$.

The first thing to notice is that in both the H^2 and H^∞ cases the gains to commitment increase with both α and λ . All else equal, higher values of these parameters make optimal policy smoother and more autocorrelated. In the case of α this occurs directly, since the shocks themselves are becoming more autocorrelated. In the case of λ this occurs because the cost of changing the instrument increases. Thus, since without commitment the policymaker cannot (credibly) deliver the optimal degree of smoothness, the gains from being able to commit increase as α and λ increase.

The second thing to notice in these figures, which is more interesting for the purposes of this paper, is that the gains to commitment are greater in the H^∞ case, for all values of α and λ . For example, when $\lambda = 1.0$ and α is small, time-consistent losses are about 35% greater than precommitment losses, in both the H^2 case and the H^∞ case. However, as α rises a wedge begins to develop. By the time α reaches 0.7 the gains to commitment have reached 75% in the H^2 case and 90% in the H^∞ case. When $\alpha = 0.9$, time-consistent losses are nearly double precommitment losses for H^2 , and more than double for H^∞ .

Finally, a note of caution should be sounded regarding these figures. The nesting, or model-matching, procedure used in this paper makes it straightforward to compare H^2 and H^∞ policies, *given an assumption about commitment*. This was done in section 4 and in Figures 1a and 1b. As discussed in section 3, however, examining the effects of *different*

commitment assumptions for the H^∞ -control problem is tricky, since dynamic consistency involves updating the degree of model uncertainty. As a result, there may be a sense in which figures 2 and 3 compare apples and oranges, since the time-consistent loss functions are associated with a time-varying degree of model uncertainty while the precommitment loss functions are not.

6 Conclusion

Economists usually assume their agents know the model (even if the econometrician does not). In those instances where some uncertainty *is* allowed, it is invariably highly structured. If agents don't know the model exactly, they at least know it up to additive i.i.d shocks, and maybe a handful of parameters they learn about via Bayes Rule.

This paper has pursued an alternative approach to model uncertainty, based on recent results from the H^∞ -control literature. In this approach, model uncertainty is unstructured in the extreme. The only assumption is that it is bounded in some norm. In fact, so little is assumed that the traditional marginal approach to optimization becomes unworkable. Instead, decision-makers are presumed to operate on the basis of 'robustness', meaning they use policies that perform adequately even in the worst of circumstances. Operationally, this just involves a switch in the norm used to calculate an optimal policy. Loosely stated, rather than minimizing the sum-of-squared deviations, we assume the decision-maker minimizes the maximum deviation.

Is this reasonable? After all, minimax decision rules fell from grace during the 1950s as researchers became dissatisfied with their (lack of) decision-theoretic foundations. Is H^∞ -control simply rediscovering the errors of the past? I would argue that it isn't. As noted by Basar and Bernhard (1995) and Hansen and Sargent (1998), H^∞ -control can be related to problems that *do* have plausible decision-theoretic foundations, e.g., risk-sensitive (LEQG) control, or Gilboa and Schmeidler's (1989) axiomatization of Knightian Uncertainty. Moreover, as was discussed earlier, the work of Mustafa and Glover (1990) shows clearly the sense in which H^∞ -control is just a limiting case of traditional H^2 -control.

The goal of this paper was to apply this recent literature to a standard policy design problem (albeit one with a forward-looking state transition equation), and to see how robustness considerations influence the nature of optimal stabilization policy and the gains to commitment. I found that robust policies are more ‘activist’, and as a result, the gains to commitment turn out to be larger under model uncertainty.

The fact that the gains from commitment are larger with model uncertainty is surprising for two reasons. First, in H^∞ -control dynamic consistency is enforced by incorporating initial conditions into the domain of the model’s unstructured uncertainty. By itself, this suggests time-consistency becomes less of an issue when there’s model uncertainty. In fact, this intuition can be thought of as the government policy analog of Marcellino and Salmon’s (1997) critique of the Lucas Critique. However, this ignores the fact that the nature of optimal policy also changes when there is model uncertainty. I find that an increased demand for smoothness in the presence of model uncertainty leads to a net increase in the cost of being unable to deliver this smoothness due to lack of credibility.

The second reason this result is surprising is more subversive, and suggests an important topic for future research. In this paper a policymaker formulates a robust decision rule in an effort to minimize the damage he could incur due to his lack of knowledge about the ‘true’ model. In the real world, the best way to deal with uncertainty is to reduce it by learning and experimentation. Such actions have been explicitly ruled out here. My guess is that if learning and model uncertainty were combined, so that a fully state-contingent experimentation strategy could not be easily formulated, then there might be gains to ‘flexibility’, which could offset the traditional gains to commitment. A useful start might be to employ the dual of the H^∞ -control problem, which involves robust filtering of an unknown hidden state when the ‘measurement error model’ is uncertain. Not surprisingly, Hansen and Sargent have already begun to work on this.

REFERENCES

- Basar, Tamer, and Pierre Bernhard, 1995, *H[∞]-Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*, 2nd Edition. Boston-Basel-Berlin: Birkhauser.
- Chari, V.V., and Patrick J. Kehoe, 1990, Sustainable Plans, *Journal of Political Economy* 98, 783-802.
- Chui, C.K., and G. Chen, 1997, *Discrete H[∞]-Optimization*, 2nd Edition, Springer
- Cohen, Daniel, and Philippe Michel, 1988, How Should Control Theory Be Used to Calculate a Time-Consistent Government Policy?, *Review of Economic Studies* 55, 263-74.
- Duren, Peter L., 1970, *Theory of H^p Spaces*, Academic Press.
- Gilboa, Itzhak, and David Schmeidler, 1989, Maximin Expected Utility with Non-unique Prior, *Journal of Mathematical Economics* 18, 141-53.
- Hansen, Lars P., Dennis Epple, and William Roberds, 1985, Linear-Quadratic Duopoly Models of Resource Depletion, in *Energy, Foresight, and Strategy*, edited by Thomas J. Sargent, Johns Hopkins Univ. Press.
- Hansen, Lars P., and Thomas J. Sargent, 1980, Formulating and Estimating Dynamic Linear Rational Expectations Models, *Journal of Economic Dynamics and Control* 2, 7-46.
- _____, 1998, Discounted Robust Filtering and Control in the Frequency Domain, Unpublished working paper, Univ. of Chicago.
- Hansen, Lars P., Thomas J. Sargent, and Thomas D. Tallarini, 1997, Robust Permanent Income and Pricing, mimeo, Univ. of Chicago.
- Kasa, Kenneth, 1998, Optimal Policy With Limited Commitment, *Journal of Economic Dynamics and Control* 22, 887-910.
- Marcellino, Massimiliano, and Mark Salmon, 1997, Robust Decision Theory and the Lucas Critique, Unpublished working paper, London Business School.
- McCallum, Bennett T., 1995, Two Fallacies Concerning Central-Bank Independence, *American Economic Review Papers and Proceedings* 85, 207-11.
- Miller, Marcus, and Mark Salmon, 1985, Policy Coordination and the Time Inconsistency of Optimal Policy in Open Economies, *Economic Journal* 95, 124-37.
- Mustafa, Denis, and Keith Glover, 1990, *Minimum Entropy H[∞]-Control*, Springer-Verlag.
- Oliner, Stephen D., Glenn D. Rudebusch, and Daniel Sichel, 1996, The Lucas Critique Revisited: Assessing the Stability of Empirical Euler Equations for Investment, *Journal of Econometrics* 70, 291-316.
- Prescott, Edward C., 1977, Should Control Theory be Used for Economic Stabilization?, *Carnegie-Rochester Conference Series on Public Policy* 7, 13-38.
- Roberds, William, 1987, Models of Policy Under Stochastic Replanning, *International Economic Review* 28, 731-55.

- Sargent, Thomas J., Autoregressions, Expectations, and Advice, 1984, *American Economic Review Papers and Proceedings* 74, 408-15.
- _____, 1998a, *The Conquest of American Inflation*, forthcoming monograph.
- _____, 1998b, Discussion of paper by Lawrence Ball, forthcoming in *Monetary Policy Rules*, edited by John Taylor, Univ. of Chicago Press.
- Sims, Christopher A., 1982, Policy Analysis With Econometric Models, *Brookings Papers on Economic Activity* 1, 107-64.
- Stokey, Nancy L., 1991, Credible Public Policy, *Journal of Economic Dynamics and Control* 15, 627-56.
- Taylor, John B., 1975, Monetary Policy During a Transition to Rational Expectations, *Journal of Political Economy* 83, 1009-21.
- Whiteman, Charles H., 1986, Analytical Policy Design Under Rational Expectations, *Econometrica* 54, 1387-405.
- Young, Nicholas, 1988, *An Introduction to Hilbert Space*, Cambridge Univ. Press.
- Zames, George, 1981, Feedback and Optimal Sensitivity: Model Reference Transformation, Multiplicative Seminorms, and Approximate Inverses, *IEEE Transactions on Automatic Control* 26, 301-20.
- Zhou, Kemin with John C. Doyle and Keith Glover, 1996, *Robust and Optimal Control*, Prentice-Hall.