Four Types of Ignorance

Lars Peter Hansen and Thomas J. Sargent*

April 4, 2015

Abstract

This paper studies alternative ways of representing uncertainty about a law of motion in a version of a classic macroeconomic targetting problem of Milton Friedman (1953). We study both "unstructured uncertainty" – ignorance of the conditional distribution of the target next period as a function of states and controls – and more "structured uncertainty" – ignorance of the probability distribution of a response coefficient in an otherwise fully trusted specification of the conditional distribution of next period's target. We study whether and how different uncertainties affect Friedman's advice to be cautious in using a quantitative model to fine tune macroeconomic outcomes.

1 Introduction

> "As Josh Billings wrote many years ago, "The trouble with most folks isn't so much their ignorance, as knowing so many things that ain't so." Pertinent as this remark is to economics in general, it is especially so in monetary economics."
> Milton Friedman (1965) [1]

Josh Billings may never have said that. Some credit Mark Twain. Despite, or maybe *because* of the ambiguity about who said them, those words convey the sense of calculations

*We thank Harald Uhlig for thoughtful comments and David Evans and Paul Ho for excellent research assistance.

[1]From the forward to Phillip Cagan's *Determinants and Effects of Changes in the Stock of Money, 1875-1960*, Columbia University Press, New York and London, 1965. page xxiii.

that Milton Friedman (1953) used to advise against using quantitative macroeconomic models to "fine tune" an economy. Ignorance about details of an economic structure prompted Friedman to recommend caution.

We use a dynamic version of Friedman's model as a laboratory within which we study the consequences of four ways that a policy maker might confess ignorance. One of these corresponds to Friedman's, while the other three go beyond Friedman's. Our model states that a macroeconomic authority takes an observable state variable X_t as given and chooses a control variable U_t that produces a random outcome for X_{t+1}:

$$X_{t+1} = \kappa X_t + \beta U_t + \alpha W_{t+1}. \tag{1}$$

The shock process W is an iid sequence of standard normally distributed random variables. We interpret the state variable X_{t+1} as a deviation from a target, so ideally the policy maker wants to set $X_{t+1} = 0$, but the shock W_{t+1} prevents this.

Friemdan framed the choice between "doing more" and "doing less" in terms of the *slope* of the response of a policy maker's decision U_t to its information X_t about the state of the economy. Friedman's purpose was to convince policy makers to lower the slope. He did this by comparing optimal policies for situations in which the policy maker knows β and in which it doesn't know β.

For working purposes, it is useful tentatively to classify types of ignorance into not knowing (i) response coefficients (β), and (ii) conditional probability distributions of random shocks (W_{t+1}). Both categories of unknowns potentially reside in our model, and we'll study the consequences of both types of ignorance. As we'll see, confining ignorance to not knowing coefficients puts substantial structure on the source of ignorance by trusting significant parts of a specification. Not knowing the shock distribution translates into not knowing the conditional distribution of X_{t+1} given time t information and so admits a potentially large and less structured class of misspecifications.

After describing a baseline case in which a policy maker completely trusts specification (1), we study the consequences of four ways of expressing how a policy maker might distrust that model.[2]

I. A "Bayesian decision maker" doesn't know the coefficient β but trusts a prior probability distribution over β. (This was Friedman's way of proclaiming model uncer-

[2]Our preoccupation within enumerating types of ignorance and ambiguity here and in Hansen and Sargent (2012) is inspired by Epson (1947).

tainty.)

II. A "robust Bayesian decision maker" uses operators of Hansen and Sargent (2007) to express distrust of a prior distribution for the response coefficient β. The operators tell the decision maker how to make cautious decisions by twisting the prior distribution in a direction that increases probabilities of β's yielding lower values.

III. A "robust decision maker" uses either the multiplier or the constraint preferences of Hansen and Sargent (2001) to express his doubts about the probability distribution of W_{t+1} conditional on X_t and a decision U_t implied by model (1). Here an operator of Hansen and Sargent (2007) twists the conditional distribution of X_{t+1} to increase probabilities of X_{t+1} values that yield low continutation utilities.

IV. A robust decision maker asserts ignorance about the same conditional distribution mentioned in item (III) by adjusting an entropy penalty in a way that Petersen et al. (2000) used to express a decision maker's desire for a decision rule that is robust at least to particular alternative probability models.

Approaches (I) and (II) are ways of 'not knowing coefficients' while approaches (III) and (IV) are ways of 'not knowing a shock distribution.' We compare how these types of ignorance affect Friedman's conclusion that ignorance should induce caution in policy making.[3]

2 Baseline model without uncertainty

Following Friedman, we begin with a decision maker who trusts model (1). The decision maker's objective function at date zero is

$$\begin{aligned}-\frac{1}{2}\sum_{t=0}^{\infty}\exp(-\delta t)E\left[(X_t)^2|X_0=x\right] = &-\frac{1}{2}\sum_{t=0}^{\infty}\exp[-\delta(t+1)]E\left[(\kappa X_t+\beta U_t)^2|X_0=x\right]\\ &-\frac{1}{2}x^2-\frac{\alpha^2\exp(-\delta)}{2[1-\exp(-\delta)]},\end{aligned} \tag{2}$$

[3]Approaches (I), (II), and (III) have been applied in macroeconomics and finance, but with the exception of Hansen and Sargent (2015), approach (IV) has not.

where $\delta > 0$ is a discount rate. The decision maker chooses U_t as a function of X_t to maximize (2) subject to the sequence of constraints (1). The optimal decision rule

$$U_t = -\frac{\kappa}{\beta} X_t \tag{3}$$

attains the following value of the objective function (2):

$$-\frac{1}{2}x^2 - \frac{\alpha^2 \exp(-\delta)}{2[1-\exp(-\delta)]}.$$

Under decision rule (3) and model (1), $X_{t+1} = \alpha W_{t+1}$.

In subsequent sections, we study how two types of ignorance change the decision rule for U_t relative to (3):

- Ignorance about β.
- Ignorance about the probability distribution of W_{t+1} conditional on information available at time t.

3 Friedman's Bayesian expression of caution

This section sets Friedman's analysis within a perturbation of model (1). We study how the decision maker adjusts U_t to offset adverse effects of X_t on X_{t+1} when he does not know the response coefficient β. Does he do a lot or a little? Friedman's purpose was to advocate doing less relative to the benchmark rule (3) for setting U_t.

We replace (1) with

$$X_{t+1} = \kappa X_t + \beta_t U_t + \alpha W_{t+1}, \tag{4}$$

where the decision maker believes that $\{\beta_t\}$ is iid normal with mean μ and variance σ^2, which we write

$$\beta_t \sim \mathcal{N}(\mu, \sigma^2). \tag{5}$$

An iid prior asserts that information about β_t is useless for making inferences about β_{t+j} for $j > 0$, thus ruling out the possibility of learning.[4]

[4] de Finetti (1937) introduced exchangeability of a stochastic process as a way to make rooom for learning. Kreps (1988, ch. 11) described de Finetti's work as a foundation of statistical learning. Kreps offered a compelling explanation of why an iid prior excludes learning. Prescott (1972) departed from iid risk to create a setting in which a decision maker designs experiments because he wants to learn. **?** study

Let x^* denote a next period value.[5] Model (4)–(5) implies that the conditional density for x^* is normal with mean $\kappa x + \mu u$ and variance $\sigma^2 u^2 + \alpha^2$. Conjecture a value function

$$V(x) = -\frac{1}{2}\nu_1 x^2 - \frac{1}{2}\nu_0$$

that satisfies the Bellman equation:

$$V(x) = \max_u - \left\{ \frac{\exp(-\delta)}{2}\nu_1(\kappa x + \mu u)^2 - \frac{\exp(-\delta)}{2}\nu_1 \sigma^2 u^2 - \frac{\exp(-\delta)}{2}\nu_1 \alpha^2 - \frac{\exp(-\delta)}{2}\nu_0 - \frac{1}{2}x^2 \right\}. \tag{6}$$

First-order conditions with respect to u imply that

$$-\exp \nu_1 \left[\mu(\kappa x + \mu u) + \sigma^2 u\right] = 0$$

so that the optimal decision rule is

$$u = -\left(\frac{\mu\kappa}{\mu^2 + \sigma^2}\right) x, \tag{7}$$

which has the same form as the rule that Friedman derived for his static model. Because

$$\frac{\mu\kappa}{\mu^2 + \sigma^2} < \frac{\kappa}{\mu},$$

caution shows up in decision rule (7) when $\sigma^2 > 0$, indicating how Friedman's Bayesian decision rule recommends "doing less" than the known-β rule (3).[6]

4 Robust Bayesians

A robust Bayesian is unsure of his prior and therefore pursues a systematic analysis of the sensitivity of expected utility to perturbations of his prior. Aversion to prior uncertainty induces a robust Bayesian decision maker to calculate a lower bound on expected utility over a set of priors, and to investigate how decisions affect that lower bound. We describe two recursive implementations.

---

a setting with a robust decision maker who experiments and learns.

[5] From now on, upper case letters denote random vectors (subscripted by t) and lower case letters are realized values.

[6] See Tetlow and von zur Muehlen (2001) and Barlevy (2009) for discussions of whether adjusting for model uncertainty renders policy more or less aggressive.

In this section, we modify the Bayesian model of section 3. The decision maker still believes that β_t is an iid scalar process. But now he doubts his prior distribution (5) for β_t. He expresses his doubts about (5) by using one of two operators of Hansen and Sargent (2007) to replace the conditional expectation operator in Bellman equation (6). We assume that the decision maker confronts ambiguity about the baseline normal distribution for $\{\beta_t\}$ that is not diminished by looking at past history. The decision maker's uncertainty about the distribution of β_t is not tied to his uncertainty about the distribution of β_{t+1} except through the specification of the baseline normal distribution. In this paper, we deliberately close down learning for simplicity. In other papers (e.g., see Hansen and Sargent (2007, 2010)), we have applied the methods used here to environments in which a decision maker learns about hidden Markov states.

4.1 T^2 and C^2 operators as indirect utility functions

We define two operators, T^2, expressing "multiplier preferences", and C^2, expressing a type of "constraint preferences" in the language of Hansen and Sargent (2001). A common parameter θ appears in both operators. But in T^2, θ is a fixed penalty parameter, while in C^2, θ is a Lagrange multipler, an endogenous object whose value varies with both the state x and the decision u.

4.1.1 Relative entropy of a distortion to the density of β_t

Let $\phi(\beta; \mu, \sigma^2)$ denote the Gaussian density for β assumed in Friedman's model. The relative entropy of an alternative density $f(\beta; x, u)$ with respect to the benchmark density $\phi(\beta; \mu, \sigma^2)$ is

$$\text{ent}(f, \phi) = \int m(\beta; x, u) \log(m(\beta; x, u))\phi(\beta; \mu, \sigma^2) d\,\beta$$

where $m(\beta; x, u)$ is the likelihood ratio

$$m(\beta; x, u) = \left(\frac{f(\beta; x, u)}{\phi(\beta; \mu, \sigma^2)} \right).$$

Relative entropy is thus an expected log-likelihood ratio evaluated with respect to the distorted density $f = m\phi$ for β.

4.1.2 The T^2 operator

Let $\theta \in [\underline{\theta}, +\infty]$ be a fixed penalty parameter where $\underline{\theta} > 0$. Where $\widetilde{V}(\beta; x, u)$ is a value function, the operator $[\mathsf{T}^2\widetilde{V}](x,u)$ is the indirect utility function for the following problem:

$$
\begin{aligned}
[\mathsf{T}^2\widetilde{V}](x,u) \quad &= \min_{m\geq 0, \int_\beta m\phi=1} \left[\int \exp(-\delta)\widetilde{V}(x,\beta,u)m(\beta;x,u)\phi(\beta;\mu,\sigma^2)d\beta + \theta \text{ ent}(m\phi,\phi)\right] \\
&= E_\beta\big(\check{m}(\beta;x,u)\exp(-\delta)\widetilde{V}(x,\beta,u)\big) + \theta \text{ ent}(\check{m}\phi,\phi)
\end{aligned} \tag{8}
$$

where $\phi(\beta;\mu,\sigma^2)$ is a Gaussian density and the minimizer is $\check{m}$. The minimizing likelihood ratio $\check{m}$ satisfies the exponential twisting formula

$$
\check{m}(\beta;x,u) \propto \exp\left[-\frac{\exp(-\delta)}{\theta}\widetilde{V}(x,\beta,u)\right], \tag{9}
$$

with the factor of proportionality being the mathematical expectation of the object on the right, so that $\check{m}(\beta;x,u)$ integrates against ϕ to unity. The "worst-case" density for β associated with $[\mathsf{T}^2\widetilde{V}](x,u)$ is[7]

$$
\hat{f}(\beta;x,u) = \check{m}(\beta;x,u)\phi(\beta;\mu,\sigma^2). \tag{10}
$$

Substituting the minimizing m into the right side of (8) and rearranging shows that the T^2 operator can be expressed as the indirect utility function

$$
[\mathsf{T}^2\widetilde{V}](x,u) = -\theta \log E\left(\exp\left[-\frac{\exp(-\delta)}{\theta}\widetilde{V}(X_t,\beta_t,U_t)\right]\Big| X_t = x, U_t = u\right), \tag{11}
$$

where according to (5), $\beta_t \sim \mathcal{N}(\mu,\sigma^2)$. Equation (8) represents $[\mathsf{T}^2\widetilde{V}](x,u)$ as the sum of the expectation of $\widetilde{V}$ evaluated at the minimizing distorted density $\check{f} = \check{m}\phi$ plus θ times entropy of the distorted density $\check{f}$. Associated with a fixed θ is an entropy

$$
\text{ent}(\check{m}\phi,\phi)(x,u) = \int \check{m}(\beta;x,u)\log(\check{m}(\beta;x,u))\phi(\beta;\mu,\sigma^2)d\,\beta \tag{12}
$$

that for a fixed θ evidently varies with x and u.

Remark 4.1. *(Relationship to smooth ambiguity models) The operator* (11) *is an indirect*

[7] Notice that the minimizing density (10) depends on both the control law u and the endogenous state x. In subsection 4.5, we describe how to construct another representation of the worst-case density in which this dependence is replaced by dependence on an additional state variable that is exogenous to the decision maker's choice of control u.

utility function that emerges from solving a penalized minimization problem over a continuum of probability distortions. The smooth ambiguity models of Klibanoff et al. (2005, 2009) posit a specification resembling and generalizing (11)*, but without any reference to such a minimization problem. In the context of a random coefficient model like* (4)*, Klibanoff et al. would use one concave function to express aversion to the risk associated with randomness in the shock W_{t+1} but conditioned on β_t, on the one hand, and another concave function to express aversion to ambiguity about the distribution of β_t, on the other hand. The functional form* (11) *emerges from the formulation of Klibanoff et al. (2005) if we use a quadratic function for aversion to risk and an exponential function for aversion to ambiguity about the probability distribution of β_t.*

4.1.3 The C^2 operator

For the T^2 operator, θ is a fixed parameter that penalizes the m-choosing agent in problem (8) for increasing entropy. The minimizing decision maker in problem (8) can change entropy but at a marginal cost θ.

For the C^2 operator, instead of a fixed θ, there is a fixed level $\eta \geq 0$ of entropy that the probability-distorting minimizer cannot exceed. The C^2 operator is the indirect utility function defined by the following problem:

$$[\mathsf{C}^2\widetilde{V}](x,u) = \min_{m(\beta;x,u)\geq 0} \int \exp(-\delta)\widetilde{V}(x,\beta,u)m(\beta;x,u)\phi(\beta;\mu,\sigma^2)d\beta \tag{13}$$

where the minimization is subject to $E_\beta m(\beta;x,u) = 1$ and

$$\text{ent}(m\phi,\phi) \leq \eta. \tag{14}$$

The minimizing choice of m again obeys equation (10), except that θ is not a fixed penalty parameter, but instead a Lagrange multiplier on constraint (13). Now the Lagrange multiplier θ varies with x and u, not entropy.

Next we shall use T^2 or C^2 to design a decision rule that is robust to a decision maker's doubts about the probability distribution of β_t.

4.2 Multiplier preference adjustment for doubts about distribution of β_t

Here the decision maker's (robust) value function V satisfies a functional equation (15) cast in terms of intermediate objects that we construct in the following steps:

- Average over the probability distribution of the random shock W_{t+1} to compute

  $$\widetilde{V}(x,\beta,u) = E\big[V(\kappa X_t + \beta_t U_t + \alpha W_{t+1})\big| X_t = x, \beta_t = \beta, U_t = u\big],$$

  where $W \sim \mathcal{N}(0,1)$.

- Next, use definition (11) to average $\widetilde{V}$ over the distribution of β_t and thereby compute the indirect utility function $[\mathsf{T}^2\widetilde{V}](x,u)$.

- Choose u to attain the right side of

  $$V(x) = \max_u -\frac{1}{2}x^2 + [\mathsf{T}^2]\left(\tilde{V}(x,u)\right). \tag{15}$$

- Iterate to convergence to construct a fixed point $V(\cdot)$.

4.3 Constraint preference adjustment for doubts about distribution of β_t

Here we use the same iterative procedure described in subsection 4.2 except that we replace the steps described in the second and third bullet points with

- Use definition (13) of C^2 to average $\widetilde{V}$ over random β_t and thereby compute $[\mathsf{C}^2]\left(\tilde{V}(x,u)\right)$.

- Choose u to attain the right side of

  $$V(x) = \max_u -\frac{1}{2}x^2 + [\mathsf{C}^2]\left(\tilde{V}(x,u)\right). \tag{16}$$

Appendix B describes how we approximate V and $\tilde{V}$ in Bellman equations (15) and (16).

4.4 Examples

Figures 1–3 show robust decision rules that attain the right sides of (15) and (16) for an infinite horizon model with parameters set at $\mu = 1, \sigma = 1, \delta = .1, \alpha = .6, \kappa = .8$. We analyze a two period version in appendix A. Figures 1, 2, and 4 show robust decision rules that attain the right side of (15) while figure 3 describes decision rules that attain the right side of (16).

Figure 1 shows decision rules for u as a function of x for θ values of .5 and 1, as well as for the $\theta = +\infty$ value associated with Friedman's Bayesian rule (7) with $\sigma^2 > 0$ and the benchmark no-ignorance rule (3) associated with $\sigma^2 = 0$. With $\theta < +\infty$, the robust decision rules are non-monotonic. They show caution relative to Friedman's rule, and increasing caution for larger values of $|x|$. For low values of $|x|$, the robust rules are close to Friedman's Bayesian rule, but for large values of $|x|$, the robust rules are much more cautious. For the two-period model analyzed in appendix A, the worst-case distribution of β_t remains Gaussian. For the infinite horizon model, the worst-case distributions seem very close to Gaussian. Figure 2 shows means and variances of worst-case probability densities (in the top two panels) and relative entropies of the worst-case distributions with respect to the density $\mathcal{N}(\mu, \sigma^2)$ under the approximating model, all as functions of x. As x increases in absolute value, the mean of the worst-case distribution of β decreases, indicating that the policy is less effective. As $|x|$ increases, the worst-case variance increases at first, but eventually it decreases. The shape of entropy as a function of x, shown in the bottom panel of figure 2, sheds light on why the worst-case variance ultimately starts decreasing. As $|x|$ increases, at first entropy increases too, reaching a maximum and then slowly decreasing. Notice how the "wings" in the two robust decision rules in figure 1 and the worst-case variances in figure 2 mirror those in the graph of entropy as a function of x in the bottom panel of figure 2. Figure 3 plots two robust decision rules and stationary distributions of X_t under those decision rules and the baseline model (1). The graphs are for two levels of σ: $\sigma = .75$ and $\sigma = 1.5$. The graphs reveal that the 'wings' in the decision rules do not occur too far out in the tails of the stationary distribution.

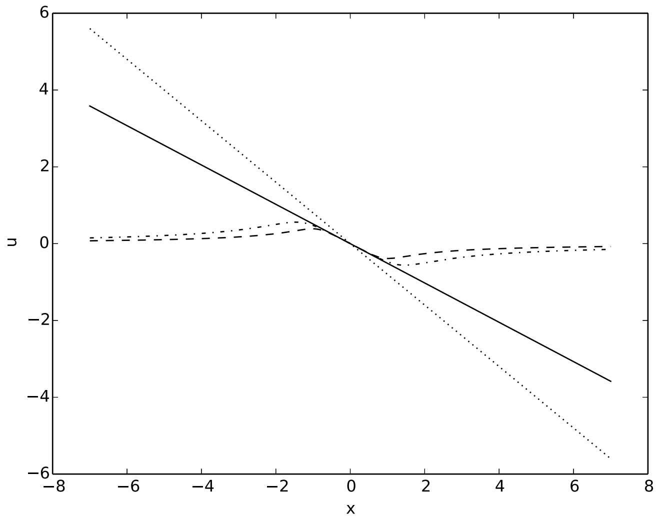

Figure 1: The β known and Friedman's Baysian decision rules (the two linear rules), and also robust decision rules that solve $\max_u -\frac{1}{2}x^2 + [\mathsf{T}^2]\left(\tilde{V}(x,u)\right)$ for two values of θ.

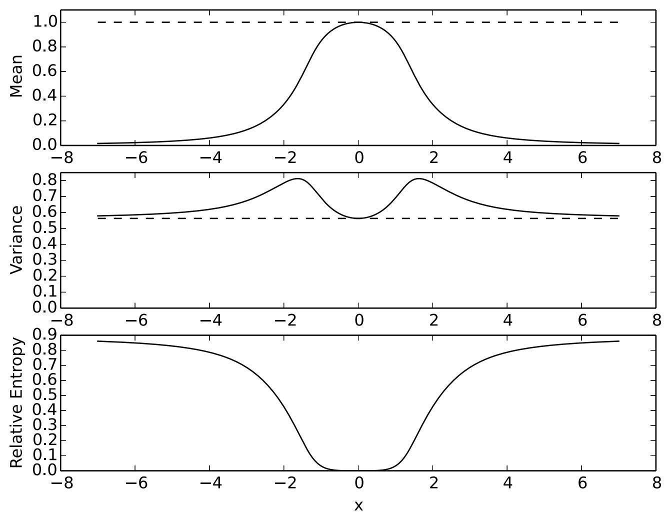

Figure 2: Top two panels: worst-case means and variances associated with decision rules that solve $\max_u -\frac{1}{2}x^2 + [\mathsf{T}^2]\left(\tilde{V}(x,u)\right)$ for various θ's; lower panel: entropy as a function of x.

.

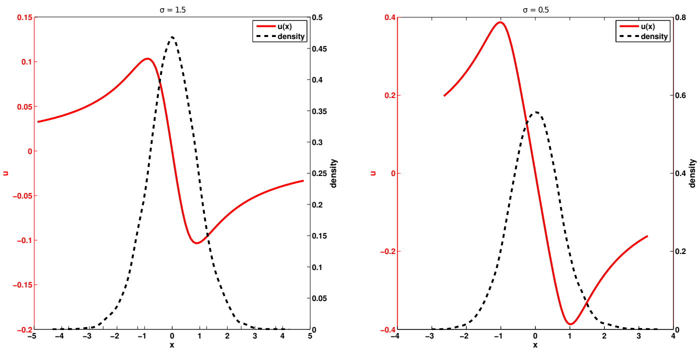

Figure 3: Robust decision rules and stationary distribution under approximating model under those rules.

The robust decision rules that attain the right side of (16) are reported in the top panel figure 4 for values of η of 0.1 and 0.4, as well as the non-robust $\eta = 0$ rule. The middle and bottom panels show the associated worst-case means and variances, which for a given entropy are constant functions of x. Here, with fixed entropy, the decision rules seem to be linear.[8] They are flatter (more cautious), the higher is entropy. The minimizing player responds to an increase in entropy by increasing the variance and decreasing the mean of β_t, making policy both less effective and more uncertain. The robust (maximizing) decision maker responds by making u less responsive to x. The constaint-preference robust Bayesian decision rule looks like Friedman's, except that it is even more cautious because of how the robust Bayesian decision maker acts as if the mean of β_t is lower and the variance of the β_t is larger than does Friedman's Bayesian decision maker.[9,10]

Figure 5 shows that the Lagrange multiplier θ for the constraint preferences in (16) increases apparently quadratically with $|x|$. For a given entropy η, the increase in the "price of robustness" θ as $|x|$ increases causes the "wings" that in the robust (15) decision

[8]In the two-period version of the model analyzed in appendix A we can prove that they are linear.

[9]This behavior is familiar from what we have seen in other robust decision problems that we have studied: in the sense that variance (noise about β) gets translated into a mean distortion.

[10]Recall formula (7) for Friedman's Bayesian decision rule and notice that

$$\frac{\partial}{\partial \mu}\left(\frac{\mu\kappa}{\mu^2+\sigma^2}\right) = \frac{\kappa(\sigma^2-\mu^2)}{(\mu^2+\sigma^2)^2},$$

which says that depending on the sign of $\sigma^2 - \mu^2$, the absolute value of the response coefficient may increase or decrease with an increase in μ. Notice that we have set $\mu = \sigma$ in our numerical, which makes this effect vanish. When $\sigma^2 = 0$, an increase in μ causes the absolute value of the slope to decrease – the policy maker does *less* because the policy instrument is *more* effective.

rules displayed in the top panel of figure 1 to disappear in the (16) decision rules displayed in figure 4.

Remark 4.2. *In applying the* C^2 *operator in recursion* (16)*, we have imposed a fixed entropy constraint period-by-period. That specification caused the "wings" associated with recursion* (15) *to vanish. If instead we had imposed a single intertemporal constraint on discounted entropy at time* 0 *and solved that entropy constraint for an associated Lagrange multiplier* θ*, it would give rise to the fixed-*θ *decision rule associated with the* T^2 *recursion* (15) *having the "wings" discussed above. Our discussion of how entropy varies with* x *in the fixed* θ *rule would then pertain to how a fixed discounted entropy would be allocated across time and states conditioned on an initial realized state* X_0*.*

4.5 Another representation of worst-case densities

In subsection 4.1.2, we reported formula (10), which expresses a worst case m as a function of the control u and the state x. This worst-case density emerges from a recursive version of the max-min problem that fist maximizes over u, then minimizes over m.[11,12] In this section, we briefly describe some analytical dividends that accrue when it is possible to form an alternative representation of a (sequence of) worst-case densities, an outcome that requires that it be legitimate to replace max-min with min-max. We presume that current and past values of $\{(W_t, \beta_{t-1})\}$ along with an initial condition for X_0 are observed when the date t control U_t is decided.

The issue under consideration is whether it is possible to construct a Bayesian interpretation of a robust decision rule. This requires a different representation of a worst-case distribution than provided by formula (10), one that takes the form of a sequence of probability densities for $\beta_t, t \geq 0$. When enough additional assumptions are satisfied, this is possible. Then it can be shown easily that the robust decision rule is an optimal decision rule for an ordinary Bayesian who takes this sequence of distributions as given.[13] In the applications described here, this sequence of worst-case distributions of β_t has some interesting features. In particular, although the baseline model asserts that the sequence

[11]The minimization is implicit when we use an indirect utility function expressed by applying a T^2 or C^2 operator, as we do in recursions (15) and (16),respectively.

[12]The type of analysis sketched in this section is presented in more detail in the context of infinite horizon linear-quadratic models in Hansen and Sargent (2008, ch. 7).

[13]Establishing that a robust decision rule can be interpreted as an optimal decision rule for a Bayesian with *some* prior distribution is needed to establish that a robust decision rule is "admissible" in the sense of Bayesian statistical decision theory.

$\beta_t, t \geq 0$ is iid, they are in general *not* iid under a worst-case model. Instead the worst-case distribution for β_{t+1} that depends on current and past information.

4.5.1 Replacing max-min with min-max

A key first step in constructing such a worst-case model that supports such a Bayesian interpretation of a robust rule involves verifying that exchanging the order of the minimization and maximization in problem (8) or (13) leaves the associated value function unaltered. If this is true, a Bellman-Isaacs condition for the associated zero-sum, two-player game is said to be satisfied. That makes it possible to construct a representation of the worst-case m of the form

$$m^*(\beta, x) = \check{m}[\beta, \check{u}(x), x], \tag{17}$$

where $\check{u}(x)$ is the robust control law. Given m^*, the robust decision rule $\check{u}$ solves a dynamic recursive version of Friedman's Bayesian problem, except that $m^*(\beta_t, x)\phi(\beta_t; \mu, \sigma^2)$ replaces the baseline model $\phi(\beta_t; \mu, \sigma^2)$.

4.5.2 Application of "Big K, little k"

For our purpose here, it is troublesome that the state variable X_t, which appears as x in (17) is endogenous in the sense that its future values are partly under the control of the maximizing decision maker. This can be avoided by using a version of the macroeconomist's "Big K, little k" trick to represent the worst-case density in a way that puts it beyond influence of the maximizing player.[14] The idea is to start by recalling that ordinary dynamic programming is a recursive method for solving a date zero problem originally cast in terms of choices of sequences. When the Bellman-Isaacs condition is satisfied, we can use an analogous mapping between date 0 problems cast in terms of sequences and problems expressed recursively a la dynamic programming applies to two-person zero sum games like the ones we are studying. That reasoning lets us also exchange orders between maximization minimization at date 0 for a robust control problem cast in terms of choices of sequences. We accomplish by inventing a new exogenous state variable process $\{Z_t\}$ whose evolution is the same as that of X_{t+1} "along an equilibrium path", meaning "when the robust control law is imposed". But now $\{Z_t\}$ is beyond the control of the maximizing agent. This lets us construct a sequence of densities $\{m^*(\beta, Z_t)\}$ relative to the normal baseline for β_{t+1}. The associated sequence of densities *ex post* justifies the robust decision rule as solving a

[14] This approach is used in Hansen and Sargent (2008, chs. 7,10,12,13).

version of Friedman' Bayesian problem. Here *ex post* means "after the outer minimizing agent has chosen its sequence $\{m^*\}$ of worst-case probability distortions."

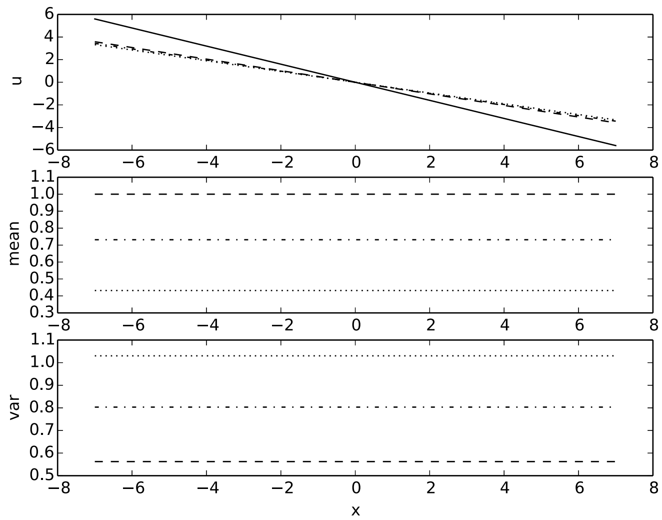

Figure 4: Top panel reports the known-β decision rule, Friedman's Bayesian decision rule, and robust decision rules that attain $\max_u -\frac{1}{2}x^2 + [\mathsf{C}^2]\left(\tilde{V}(x,u)\right)$ for different entropies η; bottom two panels report worst-case means and variances of β_t as functions of x. For bigger η, decision rules are flatter, worst-case means lower, and worst-case variances larger.

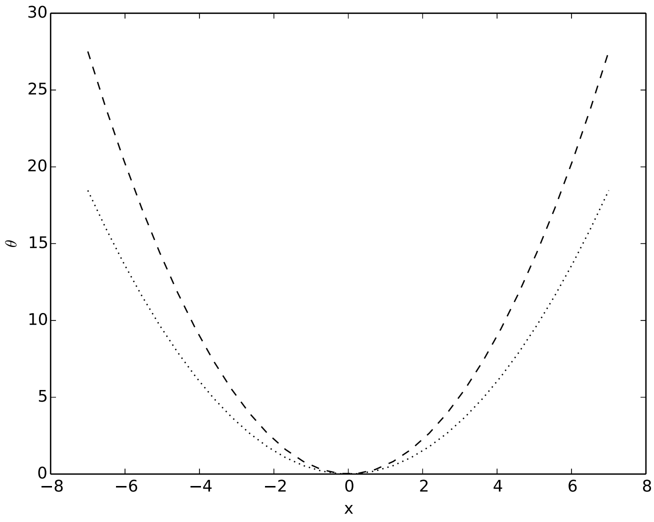

Figure 5: Lagrange multiplier θ as function of $|x|$ associated with entropy constaint for decision rules that solve $\max_u -\frac{1}{2}x^2 + [\mathsf{C}^2]\left(\tilde{V}(x,u)\right)$ for two values of entropy.

5 Uncertainty about the shock

In this section, we again apply a min-max approach in the spirit of Wald (1939), but now we use it to describe a decision maker's way of representing and coping with his doubts about the specification of the shock distribution. The decision maker regards (1) as his baseline model. But because he does not trust the implied conditional distribution for X_{t+1}, the decision maker now takes the iid normal model for W only as a benchmark model that is surrounded by other distributions that he suspects might prevail. Perturbed shock distributions can represent many alternative conditional distributions for X_{t+1}. For example, a shift in the mean of W that is a function of U_t and X_t effectively changes the statistical model (1) in ways that can include nonlinearities and history dependencies.

5.1 Multiplier preferences

We suppose that the decision maker expresses his distrust of the model (1) by behaving as if he has what Hansen and Sargent (2001) call multiplier preferences. Let $m_{t+1} = \frac{M_{t+1}}{M_t}$ be the ratio of another *conditional* probability density for W_{t+1} to the density $W_{t+1} \sim \mathcal{N}(0,1)$ in the benchmark model (1). (We apologize for recycling the $m(\cdot)$ notation previously used in section 4 but trust that it won't cause confusion.) Let $\text{ent}_t \equiv E\left[\frac{M_{t+1}}{M_t}\left(\log M_{t+1} - \log M_t\right)|\mathcal{F}_t\right]$ be the entropy of the distorted distribution relative to the benchmark distribution associated with model (1).

Let θ be a penalty parameter obeying $\theta > \alpha^2$. Define

$$
\begin{aligned}
[\mathsf{T}^1 V](x,u) &= \min_{m \geq 0, \int m(w)\psi(w)dw=1} \int m(w)(\kappa x + \beta u + \alpha w) + \theta m(w)\log m(w)\psi(w)dw \\
&= -\theta \log \int \exp\left[-\frac{1}{\theta}V(\kappa x + \beta u + \alpha w)\right]\psi(w)dw. \qquad (18)
\end{aligned}
$$

where m depends explicitly on w and implicitly (x,u) and ψ is the standard normal density. Think of $m(W_{t+1})$ as a candidate for $\frac{M_{t+1}}{M_t}$. The minimizing distortion m^* that attains the right side of the first line of (18) exponentially tilts the W_{t+1} distribution towards lower continuation values by multiplying the normal density for W_{t+1} by the likelihood ratio

$$
m^*(w) = \frac{\exp\left[-\frac{1}{\theta}V(\kappa x + \beta u + \alpha w)\right]}{E\left(\exp\left[-\frac{1}{\theta}V(\kappa x + \beta u + \alpha W_{t+1})\right]|x,u\right)}. \qquad (19)
$$

A Bellman equation for constructing a robust decision rule for U is

$$
V(x) = \max_u -\frac{1}{2}x^2 + \exp(-\delta)\mathsf{T}^1[V](x,u). \qquad (20)
$$

To solve this Bellman equation, guess a quadratic value function

$$
V(x) = -\frac{1}{2}\nu_2 x^2 - \frac{1}{2}\nu_0.
$$

Then

$$
\begin{aligned}
[\mathsf{T}^1 V](x,u) =& -\theta \log \int \exp\left[\frac{\nu_2}{2\theta}(\kappa x + \beta u + \alpha w)^2\right]\psi(w)dw - \frac{1}{2}\nu_0 \\
=& -\frac{1}{2}\nu_2(\kappa x + \beta u)^2 - \frac{1}{2}\nu_0 \\
& -\theta \log \int \exp\left[\frac{\alpha \nu_2}{\theta}(\kappa x + \beta u)w + \frac{\nu_2}{2\theta}(\alpha w)^2\right]\psi(w)dw.
\end{aligned}
$$

We can compute the worst-case m^*-distorted distribution for W_{t+1} by completing the square. The worst-case distribution is normal with precision:

$$1 - \frac{\nu_2 \alpha^2}{\theta} = \frac{\theta - \nu_2 \alpha^2}{\theta}$$

where we assume that $\theta > \nu_2 \alpha^2$. Notice that the altered precision does not depend on u. The mean of the worst-case distribution for W_{t+1} is

$$\left(\frac{\theta}{\theta - \nu_2 \alpha^2} \right) \left(\frac{\alpha \nu_2}{\theta} \right) (\kappa x + \beta u) = \left(\frac{\alpha \nu_2}{\theta - \nu_2 \alpha^2} \right) (\kappa x + \beta u),$$

which depends on (x, u) via $(\kappa x + \beta u)$. A simple calculation shows that

$$\begin{aligned}
[\mathsf{T}^1 V](x, u) =& -\frac{1}{2} \nu_2 (\kappa x + \beta u)^2 - \frac{1}{2} \nu_0 \\
& - \frac{1}{2} \left[\frac{\alpha^2 (\nu_2)^2}{\theta - \nu_2 \alpha^2} \right] (\kappa x + \beta u)^2 \\
& + \frac{\theta}{2} \left[\log \left(\theta - \nu_2 \alpha^2 \right) - \log \theta \right] \\
=& -\frac{1}{2} \left(\frac{\theta \nu_2}{\theta - \nu_2 \alpha^2} \right) (\kappa x + \beta u)^2 \\
& + \frac{\theta}{2} \left[\log \left(\theta - \nu_2 \alpha^2 \right) - \log \theta \right] - \frac{1}{2} \nu_0 .
\end{aligned}$$

The objective function on the right side of (20) is quadratic in $(\kappa x + \beta u)$ and thus the maximizing solution for u is

$$u = -\frac{\kappa}{\beta} x, \tag{21}$$

which is same control law (3) that prevails with β known and no model uncertainty. Notice also that since the log function is concave

$$\log \left(\theta - \nu_2 \alpha^2 \right) - \log \theta \le -\frac{\nu_2 \alpha^2}{\theta},$$

and thus

$$\frac{\theta}{2} \left[\log \left(\theta - \nu_2 \alpha^2 \right) - \log \theta \right] \le -\frac{1}{2} \nu_2 \alpha^2 .$$

Under control law (21), the implied worst-case mean of W_{t+1} is zero, but its variance is larger than its value 1 under the benchmark model (1). The value function satisfies:

$$-\frac{1}{2} \nu_2 x^2 - \frac{1}{2} \nu_0 = -\frac{1}{2} x^2 + \frac{\exp(-\delta)\theta}{2} \left[\log \left(\theta - \nu_2 \alpha^2 \right) - \log \theta \right] - \frac{\exp(-\delta)}{2} \nu_0 .$$

Thus, $\nu_2 = 1$ and

$$\nu_0 = -\left(\frac{\exp(-\delta)\theta}{[1-\exp(-\delta)]}\right)\left[\log\left(\theta - \nu_2\alpha^2\right) - \log\theta\right] \geq \left(\frac{\exp(-\delta)}{[1-\exp(-\delta)]}\right)\nu_2\alpha^2$$

for $\alpha^2 < \theta < \infty$. The expression on the right side is the constant term in the value function without a concern for robustness. Relative to the baseline model, the worst-case model decreases the shock precision but does nothing to the mean. The contribution of discounted entropy to ν_0 is

$$-\left(\frac{\exp(-\delta)\theta}{1-\exp(-\delta)}\right)\left[\log\left(\theta-\alpha^2\right)-\log\theta\right] - \left[\frac{\exp(-\delta)\theta\alpha^2}{1-\exp(-\delta)}\right].$$

The remainder term

$$\left[\frac{\exp(-\delta)\theta\alpha^2}{1-\exp(-\delta)}\right]$$

is the constant term for the discounted objective under the worst-case variance specification.

5.2 Explanation for no increase in caution

The outcome of the preceding analysis is that the section 2 decision rule (3) is robust to concerns about misspecifications of the conditional distribution of W_{t+1} as represented by the T^1 operator. An important ingredient of this outcome is that the worst-case model does not distort the conditional mean of W. Consider the extremization problem

$$\max_u \min_v -\frac{1}{2}x^2 - \exp(-\delta)\frac{\nu_2}{2}(\kappa x + \beta u + \alpha v)^2 - \exp(-\delta)\frac{\nu_0}{2} + \exp(-\delta)\theta\frac{v^2}{2}, \tag{22}$$

where v is a mean shift in the standard normal distribution for w. The minimizing choice of v potentially feeds back on both u and x. Why is the minimizing mean shift zero? When $\theta > \alpha^2$, the minimizing v solves solves the first-order condition:

$$-\alpha(\kappa x + \beta u + \alpha v) + \theta v = 0, \tag{23}$$

while simultaneously u satisfies the first-order condition

$$-\beta(\kappa x + \beta u + \alpha v) = 0. \tag{24}$$

In writing the first-order conditions for u, we can legitimately ignore terms in $\frac{dv}{du}$ because v itself satisfies first-order condition (23). Thus, under a policy that satisfies (24), a nonzero choice of v would leave the following key term in the objective on the right side of (22) unaffected

$$-\exp(-\delta)\frac{\nu_2}{2}(\kappa x + \beta u + \alpha v)^2 = 0,$$

but it would add an entropy penalty, so there is no gain to setting v to anything other than zero.

In summary, our "type III" specification doubts about the conditional distribution of W_{t+1} do not promote caution in the sense of 'doing less' with u in response to x. The minimizing agent gains nothing by distorting the mean of the conditional distribution of the mean of W_{t+1} and therefore chooses to distort the variance, which harms the maximizing agent but does not affect his decision rule. In section 6, we alter the entropy penalty to induce the minimizing agent to alter the worst-case mean of W in a way that makes the u-setting agent be cautious. But first, it is helpful to study in greater depth how the decision maker in the present setting best sets u to respond to possibly different mean distortions v that might have been chosen by the minimizing agent, but that were not.

5.3 Prologenomenon to structured uncertainty

To prepare the way for the analysis in section 6, we study best responses of u to "off equilibrium" choices of v in the the two-player zero-sum game (22). Thus, consider a v that instead of satisfying the v first-order necessary condition (23) takes the arbitrary form:

$$v = \xi_x x + \xi_u u + \xi_0, \tag{25}$$

where the coefficients ξ_x, ξ_u, ξ_0 need not equal those implied by (23). Such an arbitrary choice of v implies the altered state evolution equation:

$$X_{t+1} = \alpha\xi_0 + (\kappa + \alpha\xi_x)X_t + (\beta + \alpha\xi_u)U_t + \alpha\widetilde{W}_{t+1}, \tag{26}$$

where $\widetilde{W}_t$ is normally distributed with mean zero under an altered distribution for W_{t+1} having mean v instead of its value of 0 under the baseline model. It is possible for the u-setting decision maker to offset the effects of both v and x on x^* by setting u to satisfy

$$\kappa x + \beta u + \alpha v = 0, \tag{27}$$

provided that $\xi_u \neq -\frac{\beta}{\alpha}$. This response by the u-setting agent means that the v setting agent achieves nothing by using the arbitrary choice (25), but that choice incurs a time t contribution to an entropy penalty of the form

$$\exp(-\delta)\frac{1}{2\theta}(\xi_x x + \xi_u u + \xi_0)^2, \tag{28}$$

harming the v-setting agent. This deters the minimizing agent from setting (25).[15] The term $\frac{1}{\theta}(\xi_x x + \xi_u u + \xi_0)^2$ contributing to the entropy penalty for the arbitrary choice of v distortion (25) persuades the v-setting minimizing player to prefer *not* to choose that arbitrary v and instead to set $v = 0$ in the zero-sum two-player game (22).

But what if the decision maker wants to perform a sensitivity analysis of perturbations to his baseline model (1) that he insists include *particular* alternative models having the form (26). He wants some way to tell the minimizing agent this, an ability that he lacks in the framework of the present section. For that reason, in section 6, we consider situations in which the u-setting decision maker wants his decision rule to be robust to perturbations of the baseline model (1) that among others include ones that can be represented as mean distortions having the form of (25). We achieve that goal by adjusting the entropy penalty with a term involving $\frac{1}{\theta}(\xi_x x + \xi_u u + \xi_0)^2$.

6 Structured Uncertainty

In this section, we induce Friedman-like caution by making it cheaper for a minimizing player to choose a conditional mean v of W_{t+1} of the arbitrary form (25). Like Petersen et al. (2000), we do this by deducting the discounted relative entropy of the model perturbation $v = \xi_x x + \xi_u u + \xi_0$ from the right side of the entropy constraint. We can then convert θ into a Lagrange multiplier by setting it to satisfy the adjusted entropy constraint at some initial state.

Although this adjustment to the entropy constraint makes it feasible for the minimizing player to set the worst-case mean v according to (25), the minimizing player will usually make some other choice. But making the choice (25) feasible ends up altering the robust

[15] Briefly consider the case $\xi_u = -\frac{\beta}{\alpha}$ which implies that

$$\kappa x + \beta u + \alpha v = (\kappa + \alpha \xi_x)x,$$

which means that now u shows up only in the penalty term (28). The best response for the u-setting player is to set u to be arbitrarily large, making (25) a very unattractive choice for the v-setting player.

decision rule in a way that produces caution.

We now assume that the decision maker wants to explore, among other things, the consequences of parameter changes that reside within a restricted class with the form:

$$(\xi_x x + \xi_u u + \xi_0)^2 \leq \begin{bmatrix} 1 & x & u \end{bmatrix} H \begin{bmatrix} 1 \\ x \\ u \end{bmatrix}. \tag{29}$$

Parameters that satisfy this restriction can be time varying in ways about which the decision maker is unsure. To accommodate this type of specification concern, we adjust the entropy penalty by using a new operator S^1 defined as

$$[\mathsf{S}^1 V](x, u, \theta) = [\mathsf{T}^1](x, u, \theta) - \frac{\theta}{2} \begin{bmatrix} 1 & x & u \end{bmatrix} H \begin{bmatrix} 1 \\ x \\ u \end{bmatrix}$$

where H is a positive semidefinite matrix and where we now denote the explicitly the dependence on θ. To find a robust decision rule we now solve the Bellman equation

$$V(x, \theta) = \max_u -\frac{1}{2}x^2 + \exp(-\delta)[\mathsf{S}^1 V](x, u, \theta). \tag{30}$$

We now guess a value function of the form:

$$V(x, \theta) = -\frac{\nu_2(\theta)}{2}x^2 - \nu_1(\theta)x - \frac{\nu_0(\theta)}{2}$$

The first-order condition for u implies that

$$-\exp(-\delta)\left[\frac{\beta\theta\nu_2}{(\theta - \nu_2\alpha^2)}\right](\kappa x + \beta u) - \exp(-\delta)\theta \begin{bmatrix} 0 & 0 & 1 \end{bmatrix} H \begin{bmatrix} 1 \\ x \\ u \end{bmatrix} = 0.$$

which implies a robust decision rule

$$u = F(x, \theta) = -\left[\frac{\nu_2(\theta)\beta\kappa + [\theta - \nu_2(\theta)\alpha^2]\, h_{32}}{\nu_2(\theta)\beta^2 + [\theta - \nu_2(\theta)\alpha^2]\, h_{33}}\right] x - \left[\frac{\nu_1(\theta)\beta + [\theta - \nu_2(\theta)\alpha^2]\, h_{31}}{\nu_2(\theta)\beta^2 + (\theta - \nu_2(\theta)\alpha^2) h_{33}}\right], \tag{31}$$

where h_{ij} is entry (i,j) of the matrix H. The value function satisfies

$$V(x,\theta) = -\frac{1}{2}x^2 + \exp(-\delta)[\mathsf{T}^1 V][x, F(x,\theta), \theta] - \frac{\theta}{2}\begin{bmatrix} 1 & x & F(x,\theta) \end{bmatrix} H \begin{bmatrix} 1 \\ x \\ F(X,\theta) \end{bmatrix}.$$

6.1 Robustness expressed by value function bounds

To describe the sense in which decision rule (31) is robust, it is useful to construct two value function bounds.

6.1.1 Two preliminary bounds

Given a decision rule $u = F(x)$ and a positive probability distortion $m(w|x,u)$ with $\int m(w|x,u)\psi(w)dw = 1$, we construct fixed points of two recursions. The fixed point $[\mathsf{U}^1(m,F)](x)$ of a first recursion

$$[\mathsf{U}^1(m,F)](x) = -\frac{1}{2}x^2 + \exp(-\delta)\int m(w|x,F(x))\mathsf{U}^1[m,F][\kappa x + \beta F(x) + \alpha w]\psi(w)dw$$

equals discounted expected utility for a given decision rule F and under a given probability distortion m. The fixed point $[\mathsf{U}^2(m,F)](x)$ of a second recursion

$$\begin{aligned}[\mathsf{U}^2(m,F)](x) = &\int m[w|x,F(x)]\log m[w|x,F(x)]\psi(w)dw - \frac{1}{2}\begin{bmatrix} 1 & x & u \end{bmatrix} H \begin{bmatrix} 1 \\ x \\ u \end{bmatrix} \\ &+ \exp(-\delta)\int m[w|x,F(x)]\mathsf{U}^2[m,F][\kappa x + \beta F(x) + \alpha w]\psi(w)dw, \qquad (32)\end{aligned}$$

equals discounted expected relative entropy net of the adjustment for concern about the particular misspecifications m of the type (29). The probability distortion m implies an altered state evolution: given a current state x and a control u, the density for $X_{t+1} = x^*$ is

$$\frac{1}{\alpha}m\left[\frac{x^* - \kappa x - \beta u}{\alpha}\right]\psi\left[\frac{x^* - \kappa x + \beta u}{\alpha}\right].$$

Thus, m alters how the control u influences next period's state. As a particular example, the density for X_{t+1} could be normal with mean

$$(\kappa + \alpha\xi_x)x + (\beta + \alpha\xi_u)u$$

and variance α^2. Whenever inequality (29) is satisfied

$$\mathsf{U}^2(m, F)](x) \leq 0 \tag{33}$$

for all x.

Of course, in general we allow for a much richer collection of m's than the one in this example. Inequality (29) holds for each (x, u), but there will other m's that are statistically close to these deviations for which inequality (33) is also satisfied. Inequality (33) makes comparisons in terms of conditional expectations of discounted sums or relative entropies, which are weaker than the term-by-term comparisons featured in inequality (29):

$$\int m[w|x, F(x)] \log m[w|x, F(x)]\psi(w)dw - \frac{1}{2}\begin{bmatrix} 1 & x & u \end{bmatrix} H \begin{bmatrix} 1 \\ x \\ u \end{bmatrix} \leq 0.$$

6.1.2 Value function bound

We now study the max-min problem used to construct V. Using the operators U^1 and U^2,

$$[\mathsf{U}^1(m, F^\theta)](x) \geq V(x, \theta)$$

for all m and x such that

$$[\mathsf{U}^2(m, F^\theta)](x) \leq 0$$

where $F^\theta(x) = F(x, \theta)$ is a robust control law associated with θ. To make the lower bound $V(x, \theta)$ as large as possible, we solve

$$\bar{V}(x) = \max_\theta V(x, \theta).$$

For a given initial condition $\bar{x}$, we let $\bar{F}$ denote the robust control law associated with the maximizing θ. Then

$$[\mathsf{U}^1(m, \bar{F})](x) \geq \bar{V}(x)$$

for all m and x such that

$$[\mathsf{U}^2(m,\bar{F})](x) \le 0.$$

By the Lagrange multiplier theorem

$$[\mathsf{U}^1(m,\bar{F})]\,(\bar{x}) = \bar{V}\,(\bar{x})\,,$$

which indicates that the bound is sharp at $x = \bar{x}$. In this way, we have produced a robustness bound for the decision rule $\bar{F}$.[16]

Our max-min construction means that decision rules other than $\bar{F}$ cannot provide superior performance in terms of the robustness inequalities. For instance, consider some other control law $\hat{F}$. Solve the fixed point problem:

$$\hat{V}(x,\theta) = -\frac{1}{2}x^2 + \exp(-\delta)[\mathsf{S}^1\hat{V}]\left[x,\hat{F}(x),\theta\right].$$

Let $\hat{\theta}$ solve

$$\max_\theta \hat{V}\,(\bar{x},\theta)\,.$$

Then by the Lagrange multiplier theorem, $\hat{V}\left(\bar{x},\hat{\theta}\right)$ is the greatest lower bound on $\left[\mathsf{U}^1\left(m,\hat{F}\right)\right](\bar{x})$ over m's for which

$$\left[\mathsf{U}^2\left(m,\hat{F}\right)\right](\bar{x}) \le 0.$$

Since

$$\hat{V}\left(\bar{x},\hat{\theta}\right) \le V\left(\bar{x},\hat{\theta}\right) \le \bar{V}\,(\bar{x})\,,$$

the worst-case utility bound for $\hat{F}$ is lower than that for the robust control law $\bar{F}$.

In constructing U^2, we considered only time invariant choices of m. It is straightforward

[16] In this robustness analysis, we have chosen to feature one initial condition $X_0 = \bar{x}$. Alternatively, we could use an average over an initial state or include an additional date zero term in the maximization over θ. For instance we could compute

$$\min_{q\ge 0,\int q=1}\int q((x)V(x,\theta) + \theta\int[\log q(x) - \log\bar{q}(x)]\,dx$$

for a give baseline density $\bar{q}$ and use $\int[\mathsf{U}^1(m,F^\theta)](x)q(x)dx$ in place $[\mathsf{U}^1(m,F^\theta)](x)$ and

$$\int[\mathsf{U}^2(m,F^\theta)](x)q(x) + \int[\log q(x) - \log\bar{q}(x)]\,dx$$

in place of $[\mathsf{U}^2(m,F^\theta)](x)$ in the robustness inequalities.

to show that time varying parameter changes are also allowed provided that

$$\int m_t[w|x, F(x)] \log m_t[w|x, u] \psi(w) dw - \frac{1}{2} \begin{bmatrix} 1 & x & u \end{bmatrix} H \begin{bmatrix} 1 \\ x \\ u \end{bmatrix} \leq 0$$

for all nonnegative integers t and all pairs (x, u).

6.2 A caveat

By using a min-max theorem to verify that it is legitimate to exchange orders of maximization and minimization, it is sometimes possible to provide an *ex post* Bayesian interpretation of a robust decision rule.[17] In our analysis of the "structured uncertainty" contained in this section of our paper, we have not been able to proceed in this way: we have *not* shown that a robust decision rule is a best response to one of the distorted models under consideration. While a min-max theorem does apply to our analysis for a fixed θ, it does so in a way that is difficult to interpret because of the direct impact through the term

$$\frac{1}{2} \begin{bmatrix} 1 & x & u \end{bmatrix} H \begin{bmatrix} 1 \\ x \\ u \end{bmatrix}$$

of u on the allowable relative entropy.

6.3 Capturing doubts expressed through $\xi_u \neq 0$

Consider the special case $\xi_x = \xi_0 = 0$ that in the spirit of Friedman expresses uncertainty about the coefficient β on the control U in the law of motion (1). With these settings of ξ_x, ξ_0 in (25), we focus on

$$(\xi_u)^2 \leq \eta$$

for some η. The robust decision rule that attains the right side of Bellman equation (30) is

$$F(x, \theta) = - \left[\frac{\nu_2 \beta \kappa}{\nu_2 \beta^2 + (\theta - \nu_2 \alpha^2) \eta} \right] x, \tag{34}$$

[17] For example, see Hansen and Sargent (2008, ch. 7).

which has the same form as the Friedman rule (7), except that here

$$\frac{(\theta - \nu_2 \alpha^2)\eta}{\nu_2}$$

plays the role that σ^2 played in rule (7).

By way of comparison, suppose that $\eta = 0$. Then decision rule (34) becomes

$$F(x) = -\left(\frac{\kappa}{\beta}\right) x$$

which is the original rule (3) that emerges from model (1) without model uncertainty and β a known constant.

6.4 A Numerical Example

Consider the following numerical example:

$$X_{t+1} = .8X_t + U_t + .6W_{t+1},$$

which implies that the stationary distribution of X_t is a standard normal. Let $h_{33} = .09$ and the other the entries of the H matrix equal to zee

As suggested in our robustness analysis, we can impose a constraint on adjusted entropy by computing θ to solve

$$\max_\theta V(x; \theta),$$

subject to the appropriate discounted intertemporal version of our adjusted relative entropy constraint. Figure 6 shows an example of $V(x, \theta)$. The maximizing value of θ is approximately 2.36 when the initial value $x = 1$. There is in fact very little dependence on x as is shown in Figure 7. Figure 8 compares three decision rules. Figure 9 plots stationary distributions under three decision rules under the baseline model (1).

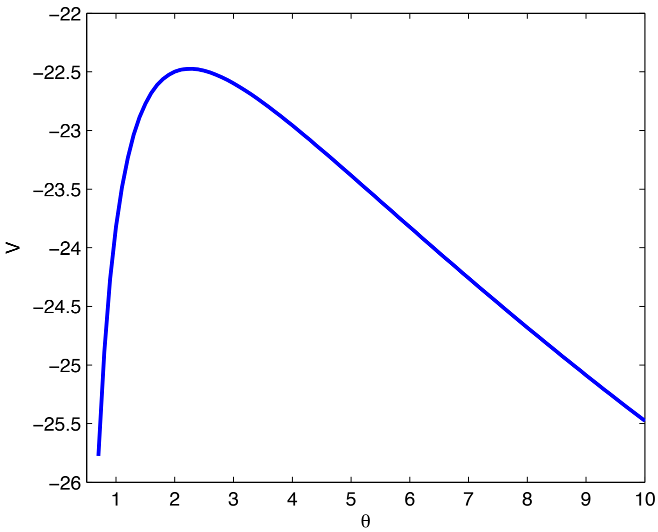

Figure 6: Value $V(x, \theta)$ as a function of θ.

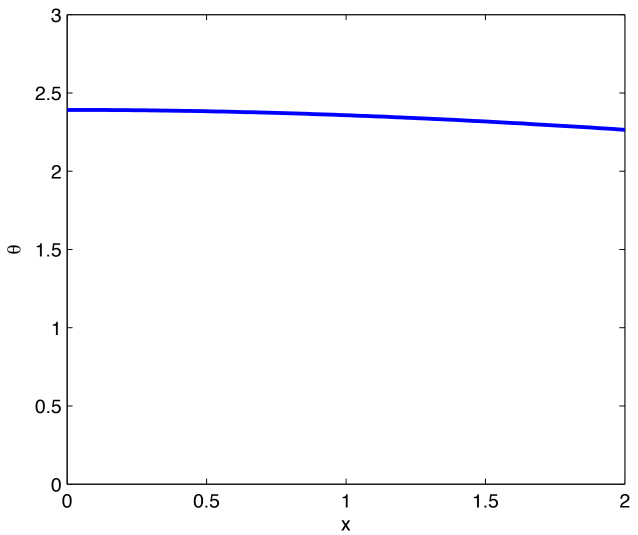

Figure 7: θ as a function of x.

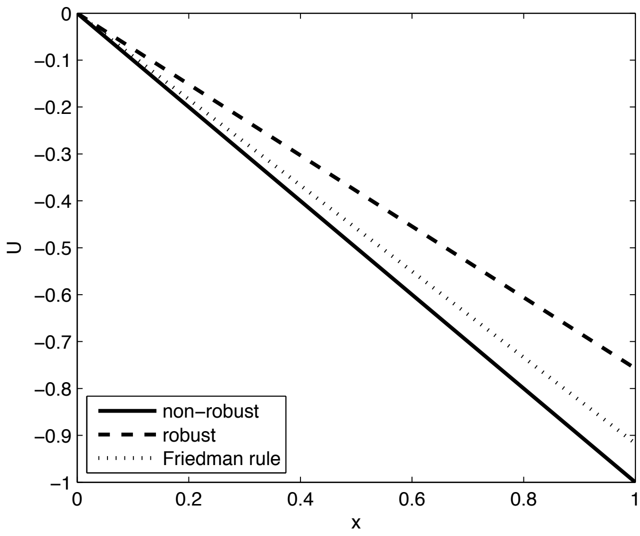

Figure 8: Three decision rules. The robust rule presumes that $\theta = 2.36$.

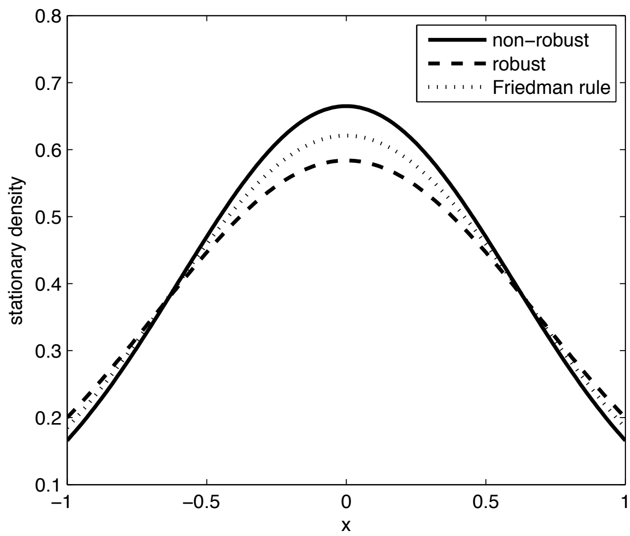

Figure 9: Stationary distribution of x under the baseline model under three decision rules.

7 Concluding remarks

We hope that readers will accept our analysis in the spirit we intend, namely, as hypothetical examples, like Friedman (1953), in which the sensitivities of policy recommendations to details of a model's stochastic structure can be examined.[18] Friedman's finding that caution due to model specification doubts translates into "doing less" as measured by a response coefficient in a decision rule linking an action u to a state x depends on the structure of the baseline model that he assumed. Even within single-agent decision problems like Friedman's but ones with different structures than Friedman's, we know other examples in which "being cautious" translates into "doing more now and picking up the pieces later." (For example, see Sargent (1999), Tetlow and von zur Muehlen (2001), Cogley et al. (2008), and Barlevy (2009)). So the "do-less" flavor of some of our results should be taken with grains of salt.

In more modern macreconomic models than Friedman's, it is essential that there are multiple agents. Refining rational expectations by imputing specification doubts to agents inside or outside a model opens interesting channels of caution beyond those appearing in the present paper. A model builder faces choices about to whom to impute model specification doubts (e.g., the model builder himself or people inside the model) and also what those doubts are about. Early work in multi-agent settings appears in Hansen and Sargent (2012, 2008, ch. 15) and Karantounias (2013).

[18]For example, the "no-increase in caution" finding of section 5 depends on our having set up the objective function to make it a "targetting problem".

A Analysis of a two-period example (by David Evans and Paul Ho)

This appendix analyses a two-period example designed to shed light on the shape of the robust policy functions presented in the infinite horizon section 4 model of a Bayesian decision maker who does not trust the prior distribution (5).

We take as given a continuation value function $V(x_{t+1}) = -\frac{1}{2}x^2$. With $\alpha = 0$, $\tilde{V}$ in section 4 becomes

$$\tilde{V}(x, \beta, u) = -\frac{1}{2}(\kappa x + \beta u)^2.$$

The T^2 operator from section 4 becomes

$$\begin{aligned}
[T^2\tilde{V}](x, u) &= -\theta \log \left[\mathbb{E} \exp \left(-\frac{\exp(-\delta)}{\theta} \tilde{V}(x, \beta, u) | x, u \right) \right] \\
&= -\theta \log \left\{ \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{\infty} \exp \left(\frac{e^{-\delta}}{2\theta} (\kappa x + \beta u)^2 \right) \exp \left(-\frac{(\beta - \mu)^2}{2\sigma^2} \right) d\beta \right\}. \qquad (35)
\end{aligned}$$

The integrand can be combined into

$$\exp \left[-\frac{1}{2} \left(\beta^2 \left(\frac{1}{\sigma^2} - \frac{u^2 e^{-\delta}}{\theta} \right) - 2\beta \left(\frac{u\kappa x e^{-\delta}}{\theta} + \frac{\mu}{\sigma^2} \right) + \frac{\mu^2}{\sigma^2} - \frac{\kappa^2 x^2 e^{-\delta}}{\theta} \right) \right].$$

This integral exists only if

$$\frac{1}{\sigma^2} - \frac{u^2 e^{-\delta}}{\theta} > 0,$$

which immediately implies an upper bound on the control u:

$$|u| \leq \sqrt{\frac{\theta}{\sigma^2 \exp^{-\delta}}}.$$

After several lines of algebra (checked with python code), the integrand in (35) can be written as the following product:

$$\frac{1}{\sqrt{2\pi}\sigma} \exp \left(-\frac{1}{2} \left[\frac{\mu^2}{\sigma^2} - \frac{\kappa x^2 e^{-\delta}}{\theta} - \frac{(\theta\mu + u\kappa x e^{-\delta}\sigma^2)^2}{\sigma^2\theta(\theta - u^2 e^{-\delta}\sigma^2)} \right] \right) \exp \left(-\frac{\left(\beta - \frac{\theta\mu + u\kappa x e^{-\delta}\sigma^2}{\theta - u^2 e^{-\delta}\sigma^2} \right)^2}{2\frac{\sigma^2\theta}{\theta - u^2 e^{-\delta}\sigma^2}} \right)$$

The term on the left is independent of β and so can be pulled outside the integral, while the infinite integral of the term on the right is easily computed to be $\sqrt{2\pi}\sigma\sqrt{\frac{\theta}{\theta-u^2e^{-\delta}\sigma^2}}$. Plugging this into equation (35), we obtain

$$[T^2\tilde{V}](x,u) = \frac{1}{2}\left[\frac{\mu^2\theta}{\sigma^2} - \kappa x^2 e^{-\delta} - \frac{(\theta\mu + u\kappa x e^{-\delta}\sigma^2)^2}{\sigma^2(\theta - u^2e^{-\delta}\sigma^2)} - \log\left(\frac{\theta}{\theta - u^2e^{-\delta}\sigma^2}\right)\right]. \tag{36}$$

We have verified this equation numerically using python code.

Remark A.1. *The worst-case distribution of β_t is*

$$\mathcal{N}\left(\frac{\theta\mu + u\kappa x e^{-\delta}\sigma^2}{\theta - u^2e^{-\delta}\sigma^2}, \frac{\sigma^2\theta}{\theta - u^2e^{-\delta}\sigma^2}\right)$$

The first-order condition for the maximization of the right side of equation (36) with respect to u is

$$\begin{aligned}\frac{\partial[T^2\tilde{V}]}{\partial u} &= \frac{1}{2}\Bigg[-\frac{2(\theta\mu + u\kappa x e^{-\delta}\sigma^2)\kappa x e^{-\delta}\sigma^2}{\sigma^2(\theta - u^2e^{-\delta}\sigma^2)} - \frac{2(\theta\mu + u\kappa x e^{-\delta}\sigma^2)^2 u e^{-\delta}\sigma^2}{\sigma^2(\theta - u^2e^{-\delta}\sigma^2)^2} - \theta\frac{\theta - u^2e^{-\delta}\sigma^2}{\theta}\frac{2\theta u e^{-\delta}\sigma^2}{(\theta - u^2e^{-\delta}\sigma^2)^2} \\ &= \frac{-1}{(\theta - u^2e^{-\delta}\sigma^2)^2}\left[(\theta\mu + u\kappa x e^{-\delta}\sigma^2)\kappa x e^{-\delta}(\theta - u^2e^{-\delta}\sigma^2)\right. \\ &\qquad \left. + (\theta\mu + u\kappa x e^{-\delta}\sigma^2)^2 u e^{-\delta} + \theta u e^{-\delta}\sigma^2(\theta - u^2e^{-\delta}\sigma^2)\right]\end{aligned} \tag{37}$$

Thus, the optimal policy $|u| \leq \sqrt{\frac{\theta}{\sigma^2 e^{-\delta}}}$ solves the following cubic equation

$$(\theta\mu + u\kappa x e^{-\delta}\sigma^2)\kappa x e^{-\delta}(\theta - u^2e^{-\delta}\sigma^2) + (\theta\mu + u\kappa x e^{-\delta}\sigma^2)^2 u e^{-\delta} + \theta u e^{-\delta}\sigma^2(\theta - u^2e^{-\delta}\sigma^2) = 0 \tag{38}$$

Note that as $x \to \pm\infty$, equation (38) will be dominated by the term

$$\begin{aligned}&u\kappa^2x^2e^{-\delta}\sigma^2(\theta - u^2e^{-\delta}\sigma^2) + (\kappa e^{-\delta}\sigma^2)^2u^3e^{-\delta}x^2 \\ &\qquad = ux^2\left[\kappa^2e^{-\delta}\sigma^2(\theta - u^2e^{-\delta}\sigma^2) + (\kappa e^{-\delta}\sigma^2)^2u^2e^{-\delta}\right].\end{aligned}$$

The term in the brackets is positive and bounded away from zero, implying that $u \to 0$ as $x \to \pm\infty$.

Remark A.2. *Because it is always possible to set $u = 0$ and thereby remove the consequences of risk and uncertainty about β, there is no lower bound on $\theta \in (0, \infty)$.*

B Computations (by David Evans)

We approximating the value function $V(x)$ in the Bellman equation (15) over an interval $[-\underline{x}, \overline{x}]$ with cubic splines. For realizations of x outside $[-\underline{x}, \overline{x}]$ a quadratic function that best fit the value function at the interpolation points was assumed. To apply the three mappings underlying (15), we used Gauss-Hermite quadrature to integrate over the Gaussian random variables. The program iterated over the bellman equations until a desired tolerance was achieved. Robustness checks were performed over the number of integration nodes for the quadrature operation.

References

Barlevy, Gadi. 2009. Policymaking under uncertainty: Gradualism and robustness. *Economic Perspectives* (Q II):38–55.

Cogley, Timothy, Riccardo Colacito, Lars Peter Hansen, and Thomas J. Sargent. 2008. Robustness and U.S. Monetary Policy Experimentation. *Journal of Money, Credit and Banking* 40 (8):1599–1623.

Epson, William. 1947. *Seven types of ambiguity.* New York: New Direction Books.

de Finetti, Bruno. 1937. La Prevision: Ses Lois Logiques, Ses Sources Subjectives. *Annales de l'Institute Henri Poincare'* 7:1 – 68. English translation in Kyburg and Smokler (eds.), *Studies in Subjective Probability*, Wiley, New York, 1964.

Friedman, Milton. 1953. The Effects of Full Employment Policy on Economic Stability: A Formal Analysis. In *Essays in Positive Economics*, edited by Milton Friedman, 117–132. Chicago, Illinois: University of Chicago Press.

Hansen, L. P. and T. J. Sargent. 2001. Robust Control and Model Uncertainty. *American Economic Review* 91:60–66.

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

———. 2008. *Robustness.* Princeton, New Jersey: Princeton University Press.

———. 2010. Fragile beliefs and the price of uncertainty. *Quantitative Economics* 1 (1):129–162.

———. 2012. Three types of ambiguity. *Journal of Monetary Economics* 59 (5):422–445.

Hansen, Lars Peter Hansen and Thomas J. Sargent. 2015. Sets of Models and Prices of Uncertainty. University of Chicago manuscript.

Karantounias, Anastasios G. 2013. Managing pessimistic expectations and fiscal policy. *Theoretical Economics* 8 (1).

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. A Smooth Model of Decision Making under Ambiguity. *Econometrica* 73 (6):1849–1892.

———. 2009. Recursive smooth ambiguity preferences. *Journal of Economic Theory* 144 (3):930–976.

Kreps, David M. 1988. *Notes on the Theory of Choice*. Underground Classics in Economics. Westview Press.

Petersen, I. R., M. R. James, and P. Dupuis. 2000. Minimax Optimal Control of Stochastic Uncertain Systems with Relative Entropy Constraints. *IEEE Transactions on Automatic Control* 45:398–412.

Prescott, Edward C. 1972. The Multi-Period Control Problem Under Uncertainty. *Econometrica* 40 (6):1043–58.

Sargent, Thomas J. 1999. Dicsussion of “Policy Rules for Open Economies” by Laurence Ball. In *Monetary Policy Rules*, edited by John Taylor, 144–154. Chicago, Illinois: University of Chicago Press.

Tetlow, Robert J. and Peter von zur Muehlen. 2001. Robust monetary policy with misspecified models: Does model uncertainty always call for attenuated policy? *Journal of Economic Dynamics and Control* 25 (6-7):911–949.

Wald, Abraham. 1939. Contributions to the Theory of Statistical Estimation and Testing Hypotheses. *The Annals of Mathematical Statistics* 10 (4):pp. 299–326.