

Aversion to Ambiguity and Model Misspecification in Dynamic Stochastic Environments

Lars Peter Hansen* Jianjun Miao†

June 28, 2018

Keywords— uncertainty; robustness; misspecification; ambiguity; risk; entropy

Abstract

Preferences that accommodate aversion to subjective uncertainty and its potential misspecification in dynamic settings are a valuable tool of analysis in many disciplines. By generalizing previous analyses, we propose a tractable approach to incorporating broadly conceived responses to uncertainty. We illustrate our approach on some stylized stochastic environments. By design, these discrete-time environments have revealing continuous-time limits. Drawing on these illustrations, we construct recursive representations of intertemporal preferences that allow for penalized and smooth ambiguity aversion to subjective uncertainty. These recursive representations imply continuous-time limiting Hamilton-Jacobi-Bellman equations for solving control problems in the presence of uncertainty.

In statistics, control theory, decision theory, and economics, the question of how to cope with subjective uncertainty comes into play. While researchers have developed many different approaches, it is standard in decision theory to impose axioms on preferences over choices. Axiomatic decision theory justifies representations of preferences that provide applied researchers with alternative ways to capture uncertainty responses. In applications, however, a decision maker must decide how to calibrate preference parameters, including aversion to ambiguity and to model misspecification. An important issue that arises is how to transport these parameters across alternative environments, a question to which existing axiomatic treatments provide little guidance.

*University of Chicago, lhansen@uchicago

†Boston University, jianjunmiao9@gmail.com

It has become common practice within economics to transport risk aversion parameters and subjective discount rate parameters from one environment to another. Through some examples we suggest that a mechanical implementation of the same approach is not appealing for some forms of aversion to ambiguity and model misspecification that interest us. These examples are not only of interest in their own right, they also provide guidance as to meaningful continuous-time limiting counterparts to Hamilton-Jacobi-Bellman (HJB) equations for which the aversions contribute in the limit.

We build our analysis in multiple steps. We first use a version of Shannon's relative entropy to quantify uncertainty. Specifically, in section 1 we present a relative entropy decomposition of the joint probability distribution over a future observation and an unknown parameter that we find to be revealing when conceptualizing various forms of aversion. In section 2 we pose two alternative static robustness problems using relative entropy penalizations. Both are special cases of the variational preferences axiomatized by Maccheroni et al. (2006) and Strzalecki (2011). The first underlies the known connection between risk sensitive and robust control dating back to Jacobson (1973). As we illustrate, this problem features the potential misspecification of the predictive density familiar from statistics. A second problem targets a concern about prior misspecification. In section 3, we solve these two problems using examples and explore some limits. By changing the exposure to an underlying shock, we distinguish in a sharp way the contributions from the likelihood and prior to the relative entropy of the joint distribution. We provide a way to scale aversion parameters so that their impact remains intact even in the limiting economy.

The limits we explore in section 3 are, by design, valuable inputs into the formulation and calibration of uncertainty preferences in dynamic stochastic environments. For example, in section 4 we study the impact of compounding the assessments of risk, ambiguity, and the potential for misspecification of subjective probabilities in a sequence of simple recursive learning environments. We also explore a convenient parametrization of smooth ambiguity specification of Klibanoff et al. (2005), which we know from Hansen and Sargent (2007) sometimes can be motivated by an aversion to prior uncertainty. We use these calculations to shed light on how to parameterize aversion to ambiguity and model misspecification across alternative dynamic environments, including an environment that emerges as a continuous-time limit. Finally, in section 5 we propose a more general continuous-time specification with distinct forms of aversions to alternative uncertainty components. This provides a counterpart of the discrete-time specifications of Hansen and Sargent (2007), Hayashi and Miao (2011), and Ju and Miao (2012). Moreover, it provides HJB equations

that allow a decision maker to confront uncertainty in its various forms.

1 Entropy penalization

Relative entropy is an attractive and tractable measure of discrepancy between probability measures. We use it to guide our formulation of preferences for a decision maker with an aversion to model ambiguity and concerns about misspecification. As a precursor to formulating such preferences, we explore the relative entropy relations between priors, likelihoods, and predictive densities in a static setting.

We start with a prior distribution π for hypothetical parameter values θ over a set Θ and a likelihood λ that informs us of the density for possible outcomes $y \in \mathcal{Y}$ given θ (with respect to a measure τ). Let \mathbf{P} denote the implied joint probability measure over $\mathcal{Y} \times \Theta$

$$\mathbf{P}(dy, d\theta) = \lambda(y|\theta)\tau(dy)\pi(d\theta).$$

The corresponding predictive density for y that integrates over θ is

$$\phi(y) = \int_{\Theta} \lambda(y|\theta)\pi(d\theta).$$

In our analysis it is the predictive density that the decision maker cares about when taking actions, but he or she has uncertainty about subjective inputs into its construction. We introduce baseline counterparts with $\hat{\cdot}$'s and explore discrepancies relative to these baselines.

The relative entropy discrepancy of the joint distribution for outcomes and parameters is defined as:

$$D(\mathbf{P} \mid \hat{\mathbf{P}}) = \int_{\Theta} \log \left[\frac{d\pi}{d\hat{\pi}}(\theta) \right] \pi(d\theta) + \int_{\Theta} \int_{\mathcal{Y}} \lambda(y|\theta) \log \left[\frac{\lambda(y|\theta)}{\hat{\lambda}(y|\theta)} \right] \tau(dy)\pi(d\theta), \quad (1)$$

where we presume that \mathbf{P} is absolutely continuous with respect to the baseline $\hat{\mathbf{P}}$. This representation gives us two potential contributions to relative entropy, one coming from differences in the prior and the other from differences in the likelihood.

There is an equivalent way to represent the same relative entropy in terms of predictive

densities and posterior distributions where the posterior distribution, π^+ , is

$$\pi^+(d\theta | y) = \frac{\lambda(y | \theta)}{\phi(y)}\pi(d\theta)$$

and $\hat{\pi}^+$ is defined analogously.

2 Relative entropy and robustness problem

In this section we pose and solve two problems that adjust for robustness. These problems are the ingredients to preferences that capture a concern for robustness. While both use relative entropy to constrain the robustness assessment, one does so in a more restricted way in order to target the robust choice of a prior. The solutions to both are special cases of the solution to a more general problem that entails exponential tilting toward lower expected utilities.

Consider the first robust evaluation:

Problem 2.1.

$$\min_{\mathbf{P}} \int U(y)d\mathbf{P} + \kappa D(\mathbf{P} | \hat{\mathbf{P}}).$$

In this problem, $\kappa > 0$ is a parameter that penalizes the search for robustness in the likelihood and prior. The $\kappa = \infty$ limit enforces a commitment to the baseline specifications. Since the function $U(y)$ depends only on y and not θ , the solution distorts only the predictive density

$$\phi^*(y) = \frac{\exp[-\frac{1}{\kappa}U(y)] \hat{\phi}(y)}{\int_{\mathcal{Y}} \exp[-\frac{1}{\kappa}U(y)] \hat{\phi}(y)\tau(dy)}$$

and not the posterior π^+ of θ given y . The resulting minimized objective is

$$-\kappa \log \int_{\mathcal{Y}} \exp\left[-\frac{1}{\kappa}U(y)\right] \hat{\phi}(y)\tau(dy).$$

The exponential tilting towards outcomes y with lower utility is familiar from the extensive literature applying relative entropy penalizations in control and estimation problems.

In particular, Problem 2.1 has the same solution as

$$\min_{\phi} \int_{\mathcal{Y}} U(y)\phi(y)\tau(dy) + \kappa \int_{\mathcal{Y}} \left[\log \phi(y) - \log \hat{\phi}(y) \right] \phi(y)\tau(dy).$$

This latter problem is the static counterpart to the dynamic recursive specification used extensively in robust control theory. See, for instance, Jacobson (1973), Doyle et al. (1989), James (1992), Basar and Bernhard (1995), and Petersen et al. (2000). While the construction of the predictive density $\hat{\phi}$ embeds the reference prior $\hat{\pi}$, prior sensitivity is only confronted in an indirect way by Problem 2.1 through potential modifications in the predictive density.

As an alternative, we target prior robustness by restricting $\lambda = \hat{\lambda}$ and thus omitting specification concerns about the likelihood. One justification for this is that uncertainty about how a prior weights alternative likelihoods already gives us a way to capture many forms of likelihood uncertainty. Let Π denote the set of priors that are absolutely continuous with respect to $\hat{\pi}$ and solve:

Problem 2.2.

$$\min_{\pi \in \Pi} \int \bar{U}(\theta)\pi(d\theta) + \kappa \int_{\Theta} \log \left[\frac{d\pi}{d\hat{\pi}}(\theta) \right] \pi(d\theta)$$

where

$$\bar{U}(\theta) \doteq \int_{\mathcal{Y}} U(y)\hat{\lambda}(y|\theta)\tau(dy).$$

Notice that \bar{U} is constructed without likelihood uncertainty. In using the formulation to investigate prior sensitivity, we penalize only the prior contribution to relative entropy. This formulation follows Hansen and Sargent (2007).

The solution to this problem gives the worst-case prior

$$\pi^*(d\theta) = \frac{\exp \left[-\frac{1}{\kappa} \bar{U}(\theta) \right]}{\int_{\Theta} \exp \left[-\frac{1}{\kappa} \bar{U}(\theta) \right] \hat{\pi}(d\theta)} \hat{\pi}(d\theta),$$

provided that the right-hand side integral is finite. The optimized objective is

$$-\kappa \log \int_{\Theta} \exp \left[-\frac{1}{\kappa} \bar{U}(\theta) \right] \hat{\pi}(d\theta). \tag{2}$$

The solution to this particular problem is recognizable as a smooth ambiguity objective

and a special case of Klibanoff et al. (2005).¹

3 A revealing family of stochastic environments

We study Problems 2.1 and 2.2 for a family of stochastic environments to illustrate the impacts of the penalization. We design these illustrations to serve as a precursor to formulating continuous-time limits to be used in a dynamic setting. They will provide input into formulations of HJB equations for continuous-time decision problems.

Consider a parameterized family of log-normal random variables, indexed by ϵ , and assume a logarithmic utility function U :

$$U(Y_\epsilon) = \log Y_\epsilon, \quad \log Y_\epsilon = y_0 + \mu(\theta)\epsilon - \frac{1}{2}|\varsigma|^2\epsilon + \varsigma \cdot Z_\epsilon, \quad (3)$$

where μ is a measurable function of θ and Z_ϵ is a normally distributed random vector with mean 0 and covariance matrix ϵI .² We find it convenient to think of $\epsilon > 0$ as an interval of time and $\{Z_\epsilon : 0 \leq \epsilon \leq 1\}$ as a multivariate standard Brownian motion.

Under the $\hat{\phi}_\epsilon$ predictive density $\log Y_\epsilon$ is distributed as a mixture of a prior distribution $\hat{\pi}$ for θ and a normal distribution with mean

$$y_0 + \epsilon\mu(\theta) - \frac{\epsilon}{2}|\varsigma|^2$$

and variance $\epsilon|\varsigma|^2$. For pedagogical simplicity, we consider the small ϵ approximation whereby the distribution for $\log Y_\epsilon$ implied by $\hat{\phi}_\epsilon$ is normal with mean:

$$y_0 + \epsilon \int_{\Theta} \mu(\theta)\hat{\pi}(d\theta) - \frac{\epsilon}{2}|\varsigma|^2$$

and variance $\epsilon|\varsigma|^2$. The minimized objective in Problem 2.1 is

$$\begin{aligned} -\kappa \log \int_{\mathcal{Y}} \exp \left[-\frac{1}{\kappa} U(y) \right] \hat{\phi}_\epsilon(y) \tau(dy) \\ \approx y_0 + \epsilon \left[\int_{\Theta} \mu(\theta)\hat{\pi}(d\theta) - \left(\frac{1}{2} + \frac{1}{2\kappa} \right) |\varsigma|^2 \right]. \end{aligned}$$

¹See Segal (1990) and Davis and Paté-Cornell (1994) for earlier motivations for smooth ambiguity decision problems.

²We use $|\cdot|$ to denote the Euclidean norm and I to denote an identity matrix.

Thus, up to the first-order approximation in ϵ , the decision maker targets concerns about misspecified likelihoods in contrast to misspecified priors.

Consider now Problem 2.2, where

$$\bar{U}_\epsilon(\theta) \doteq \int_{\mathcal{Y}} \log(y) \hat{\lambda}_\epsilon(y|\theta) \tau(dy) \quad (4)$$

The minimized objective is given by

$$\begin{aligned} -\kappa \log \int_{\Theta} \exp \left[-\frac{1}{\kappa} \bar{U}_\epsilon(\theta) \right] \hat{\pi}(d\theta) \\ \approx y_o + \epsilon \left[\int_{\Theta} \mu(\theta) \hat{\pi}(d\theta) - \frac{|\zeta|^2}{2} \right], \end{aligned}$$

which is the same as the expected utility level under $\hat{\pi}$ up to the first-order approximation in ϵ . Thus the robustness adjustment vanishes to the first-order in ϵ . In contrast, the risk-aversion adjustment associated with logarithmic utility remains present.

Why is this the case? The probability distribution for Y_ϵ becomes concentrated as ϵ becomes small conditioned on θ . The same is true of the predictive distributions for alternative priors. By contrast, the prior divergence contribution to relative entropy does not depend on ϵ . This difference in how the prior and likelihoods behave as a function of ϵ causes the entropy penalty relative to the utility distortion to converge to infinity as ϵ goes to zero. Thus the solution to Problem 2.2 implies no robustness adjustment up to the first-order approximation.

As an alternative, we alter Problem 2.2 by letting the prior divergence scale in ϵ so that the consequence and cost of altering priors have comparable magnitudes up to the first-order in ϵ . Thus we allow κ to depend on ϵ by setting $\kappa(\epsilon) = \kappa_d \epsilon$ for some constant $\kappa_d > 0$ and consider the problem:

Problem 3.1.

$$\min_{\pi \in \Pi} \int_{\Theta} \bar{U}_\epsilon(\theta) \pi(d\theta) + \epsilon \kappa_d \int_{\Theta} \log \left[\frac{d\pi}{d\hat{\pi}}(\theta) \right] \pi(d\theta).$$

The minimized objective is now given by

$$-\epsilon \kappa_d \log \int_{\Theta} \exp \left[-\frac{1}{\epsilon \kappa_d} \bar{U}_\epsilon(\theta) \right] \hat{\pi}(d\theta).$$

Plugging (4) into this expression yields

$$y_0 + \epsilon \left\{ -\kappa_d \log \int_{\Theta} \exp \left[-\frac{1}{\kappa_d} \mu(\theta) \right] \hat{\pi}(d\theta) \right\} - \frac{\epsilon}{2} |\varsigma|^2.$$

We may interpret the term in $\{\cdot\}$ as certainty equivalent of $\mu(\theta)$ adjusted for robustness. By Jensen's Inequality, it is lower than the Bayesian mean of $\mu(\theta)$ under the baseline prior. Thus the robust prior adjustment incurs a utility cost in terms of the local mean.

The worst-case prior is given by

$$\begin{aligned} \pi_{\epsilon}^*(d\theta) &= \frac{\exp \left[-\frac{1}{\epsilon \kappa_d} \bar{U}_{\epsilon}(\theta) \right] \hat{\pi}(d\theta)}{\int_{\Theta} \exp \left[-\frac{1}{\epsilon \kappa_d} \bar{U}_{\epsilon}(\theta) \right] \hat{\pi}(d\theta)} \\ &\rightarrow \frac{\exp \left[-\frac{1}{\kappa_d} \mu(\theta) \right] \hat{\pi}(d\theta)}{\int_{\Theta} \exp \left[-\frac{1}{\kappa_d} \mu(\theta) \right] \hat{\pi}(d\theta)}, \end{aligned}$$

as $\epsilon \rightarrow 0$. Therefore, by reducing the penalization proportionally to ϵ , the exponential tilting applies to the first derivative of the objective function U with respect to ϵ . In so doing, we are left applying a particular smooth ambiguity adjustment to the local mean of $U(Y)$ conditioned on θ . The worst-case prior puts more weight on the lower values of the mean $\mu(\theta)$.

We see little motivation for holding fixed the robustness penalization as we change ϵ . Quite the contrary, we see a good reason for altering it as suggested in our simple scaling adjustment. This scaling lends itself the study of prior robustness within a continuous-time setting for decision problems.

The same issue emerges with other approaches to ambiguity as in the smooth ambiguity model of Klibanoff et al. (2005). With that perspective, we view $-\frac{1}{\kappa(\epsilon)}$ as an exponential adjustment for ambiguity aversion capturing a differential preference response for exposure to ambiguity induced by a prior over θ , in contrast to risk conditioned on the parameter θ . By holding $\kappa(\epsilon)$ fixed independent of ϵ , we reproduce Skiadas (2013). By instead using $\kappa(\epsilon) = \epsilon \kappa_d$, we obtain a limiting smooth ambiguity adjustment with a hyperbolic parameterization of ambiguity aversion.

4 Recursive risk and smooth ambiguity aversion

We now explore consequences of compounding risk, ambiguity and misspecification aversion over time. We continue to find the environment in section 3 captured by (3) to be revealing from a pedagogical standpoint, allowing us to draw on some of our previous insights. The calculations in this section illustrate how risk compounds in rather different ways than uncertainty about priors over parameters. The resulting differences are pertinent to how we conceive of ambiguity and misspecification aversion in dynamic environments in general, and in continuous-time limit environments in particular.³ Given that there are known ways for ambiguity and misspecification aversion to have an impact on preferences in discrete time, we find it most appealing to adopt parameterizations with more meaningful and revealing continuous-time limits.

In what follows, partition the interval $[0, 1]$ into subintervals of length $1/n$, and let $\epsilon = 1/n$. Let $\mathcal{F}_{j\epsilon}^\epsilon$ be the sigma algebra generated by θ and $\{Y_0, Y_{1\epsilon}, Y_{2\epsilon}, \dots, Y_{j\epsilon}\}$ for $j = 0, 1, 2, \dots, n - 1$.

4.1 Risk aversion

We first consider a recursive construction of risk aversion over subintervals. Introduce a certainty equivalent operator applied to a positive random variable X that is $\mathcal{F}_{t+\epsilon}^\epsilon$ measurable,

$$\mathbb{R}_{\epsilon,t}(X) = (E [X^{1-\gamma} | \mathcal{F}_t^\epsilon])^{\frac{1}{1-\gamma}},$$

where $t = j\epsilon$ and $\gamma > 0$ represents the risk aversion parameter. Note that

$$\log \mathbb{R}_{\epsilon,t}(Y_{t+\epsilon}) = \log Y_t + \epsilon \left[\mu(\theta) - \frac{\gamma}{2} |\varsigma|^2 \right].$$

for $\epsilon > 0$. The risk adjustment $\epsilon\gamma|\varsigma|^2/2$ scales with ϵ .

Introduce a recursive construction for the certainty equivalent of the terminal consumption Y_1 ,

$$U_t = \mathbb{R}_{\epsilon,t}(U_{t+\epsilon}), \quad U_1 = Y_1, \tag{5}$$

³The calculations build in part on prior work by Hansen and Sargent (2011), who first discuss productive ways for robustness concerns to persist in a continuous-time limiting specification. Our analysis here addresses these issues more generally and forges connections to a broader collection of previous contributions.

for $t = 0, \epsilon, \dots, (n - 1)\epsilon$. We show that

$$U_0 = \mathbb{R}_{1,0}(Y_1) = Y_0 \exp \left[\mu(\theta) - \frac{\gamma}{2} |\varsigma|^2 \right],$$

for any value of ϵ . Notice that in the recursive construction on the right-hand side of (5), we use the same risk aversion coefficient γ in the certainty equivalent operator for any ϵ . This simple example shows why the common practice of holding γ fixed across environments (in our case indexed by ϵ) can be sensible, even as we shrink the exposure to risk by looking at small time intervals, due to the law of iterated expectations.

4.2 Parameter learning

Prior to investigating a recursive specification of smooth ambiguity, we remind readers of Bayesian updating within this setting. Analogous to our previous discussion, we let $\mathcal{G}_{j\epsilon}^\epsilon$ be the sigma algebra generated by $\{Y_0, Y_{1\epsilon}, Y_{2\epsilon}, \dots, Y_{j\epsilon}\}$ for $j = 0, 1, 2, \dots, n-1$, but we exclude the random parameter θ from the construction. The resulting family $\{\mathcal{G}_{j\epsilon}^\epsilon : j = 0, 1, 2, \dots, n-1\}$ captures the reduced information structure for a decision maker that does not know the parameter realization. For simplicity let $\mu(\theta) = \theta$.

Suppose that the date-zero prior for θ is normal with mean m_0 and variance q_0 . Only data on Y are used to make inferences about θ . The posterior conditioned on \mathcal{G}_t^ϵ is normal with mean m_t and variance q_t . The Bayesian recursive updating for $m_{t+\epsilon}$ is

$$m_{t+\epsilon} - m_t = \frac{q_t}{|\varsigma|^2 + \epsilon q_t} \left(\log Y_{t+\epsilon} - \log Y_t - \epsilon m_t + \frac{\epsilon}{2} |\varsigma|^2 \right), \quad (6)$$

and the conditional variance update is

$$q_{t+\epsilon} - q_t = \frac{\epsilon(q_t)^2}{|\varsigma|^2 + \epsilon q_t}. \quad (7)$$

Both equations have well-known continuous-time limits. The limiting version of equation (6) for the conditional mean is

$$dm_t = \frac{q_t}{|\varsigma|^2} \left(d \log Y_t - m_t dt + \frac{|\varsigma|^2}{2} dt \right),$$

and the limiting version of equation (7) is

$$\frac{dq_t}{dt} = -\frac{(q_t)^2}{|\varsigma|^2}.$$

4.3 Recursive robust priors

Analogous to (2), define a certainty equivalent robustness adjustment operator

$$\mathbb{B}_{\epsilon,t}(X) = -\kappa(\epsilon) \log E \left[\exp \left(-\frac{1}{\kappa(\epsilon)} X \right) \mid \mathcal{G}_t^\epsilon \right]$$

for a random variable X that is \mathcal{F}_t^ϵ measurable. Here we allow the robustness parameter $\kappa(\epsilon) > 0$ to depend on ϵ . Motivated by the static model in sections 3 and 4, we define the recursive adjustment for robust priors for log utility over the terminal consumption Y_1 ,

$$\log B_t = \mathbb{B}_{\epsilon,t} [E(\log B_{t+\epsilon} \mid \mathcal{F}_t^\epsilon)], \quad B_1 = Y_1, \quad (8)$$

for $t = 0, \epsilon, \dots, (n-1)\epsilon$.

We show in the online Materials and Methods section that

$$\log B_t = \log Y_t + (1-t)m_t + \mathbf{b}_t,$$

where

$$\mathbf{b}_{t+\epsilon} - \mathbf{b}_t = \frac{\epsilon^2}{2\kappa(\epsilon)} q_t \left(\frac{(1-t)q_t + |\varsigma|^2}{|\varsigma|^2 + \epsilon q_t} \right)^2 + \frac{\epsilon|\varsigma|^2}{2}.$$

From this formula we see that the worst-case prior/posterior for θ has a mean given by

$$m_t - \frac{\epsilon}{\kappa(\epsilon)} q_t \left(\frac{(1-t)q_t + |\varsigma|^2}{|\varsigma|^2 + \epsilon q_t} \right).$$

Using the recursive equation (8) and the conjectured form for B_t we obtain

$$\mathbf{b}_{t+\epsilon} - \mathbf{b}_t = \frac{\epsilon^2}{2\kappa(\epsilon)} q_t \left(\frac{(1-t)q_t + |\varsigma|^2}{|\varsigma|^2 + \epsilon q_t} \right)^2 + \frac{\epsilon|\varsigma|^2}{2}.$$

If the robustness parameter $\kappa(\epsilon)$ is held constant independent of ϵ , then robustness vanishes in the continuous-time limit as $\epsilon \rightarrow 0$. By contrast, we let $\kappa(\epsilon) = \kappa_d \epsilon$ as in Problem 3.1 so that penalization used in making robust prior/posterior adjustments scales with ϵ . Take

small ϵ limits and find

$$\frac{db_t}{dt} = \frac{1}{2\kappa_d} q_t \left(\frac{(1-t)q_t + |\varsigma|^2}{|\varsigma|^2} \right)^2 + \frac{|\varsigma|^2}{2}.$$

The first term on the right-hand side is the recursive robustness adjustment that remains present in a continuous-time limit because of how we scale the robustness penalty as a function of ϵ .

4.4 Risk aversion and smooth ambiguity

We next explore “smooth ambiguity adjustments” in Klibanoff et al. (2005) over different time intervals including ones that are arbitrarily small. For X that is \mathcal{F}_t^ϵ measurable, define a certainty equivalent operator

$$\mathbb{A}_{\epsilon,t}(X) = [E(X^{-\alpha(\epsilon)} | \mathcal{G}_t^\epsilon)]^{-\frac{1}{\alpha(\epsilon)}},$$

where $\alpha(\epsilon) > 0$ may depend on ϵ and captures aversion to the uncertainty about the unknown parameter θ .

We now investigate how recursive smooth ambiguity behaves as a function of ϵ . Define smooth ambiguity adjustments recursively for utility over the terminal consumption Y_1 ,

$$A_t = \mathbb{A}_{\epsilon,t}[\mathbb{R}_{\epsilon,t}(A_{t+\epsilon})], \quad A_1 = Y_1. \tag{9}$$

In the online Materials and Methods section, we show that

$$\log A_t = \log Y_t + (1-t)m_t + \mathbf{a}_t$$

where

$$\begin{aligned} \mathbf{a}_{t+\epsilon} - \mathbf{a}_t = & \\ & \frac{1}{2} \left(\frac{|\varsigma|^2 + (1-t)q_t}{|\varsigma|^2 + \epsilon q_t} \right)^2 [\epsilon(\gamma - 1)|\varsigma|^2 + \alpha(\epsilon)\epsilon^2 q_t] \\ & + \frac{\epsilon|\varsigma|^2}{2} \end{aligned}$$

and $\mathbf{a}_1 = 0$.

Consistent with our earlier discussion of smooth ambiguity over small time intervals, by letting $\alpha(\epsilon)$ be the same for all $\epsilon > 0$, smooth ambiguity contributes only a second-order adjustment.⁴ In contrast, by adopting a hyperbolic parameterization: $\alpha(\epsilon) = \frac{\alpha_h}{\epsilon}$ the smooth ambiguity contribution becomes first-order and does not vanish in a continuous-time limit. With this specification, the continuous-time limiting equation becomes

$$\frac{da_t}{dt} = \frac{1}{2} \left(\frac{|\varsigma|^2 + (1-t)q_t}{|\varsigma|^2} \right)^2 [\alpha_h q_t + (\gamma - 1)|\varsigma|^2] + \frac{|\varsigma|^2}{2}.$$

It matches our robust prior recursion $A_t = B_t$ by setting $\gamma = 1$ and $\alpha_h = \frac{1}{\kappa_d}$. Thus the power-power specification of smooth ambiguity adjustments is related to robust prior adjustments through a log transformation.

In obtaining the limiting recursion, we held the risk aversion parameter the same over environments indexed by ϵ while we scaled the ambiguity aversion or robustness parameter with ϵ . While there is no perfect way to transform ambiguity aversion parameters across environments with different discrete-time increments, the embedding we suggest has the virtue of possessing a tractable continuous-time limit. While we may think that setting $\alpha(\epsilon) = \frac{\alpha_h}{\epsilon}$ amounts to imposing “infinite” ambiguity aversion in the continuous-time limit, this is misguided in our view. Notice that in the continuous-time limit α_h scales the variance associated with estimation. As an analogy to our scaling, let $0 < \beta < 1$ denote the discount factor over a unit of time. Then we think of $\beta^{\frac{1}{n}}$ as the discount factor over an interval $\frac{1}{n}$ which in the large n limit implies a unit discount factor in the limit. But in fact the rate $-\log \beta$ continues to play a role of discounting in the continuous-time limit.

5 Implications for stochastic control

Let $(\Omega, \mathcal{F}, \mathbf{P})$ be a probability space and time is continuous over $[0, T]$. Let $\{Y_t : t \geq 0\}$ be an observable stochastic process and θ an unknown parameter. Let $\{\mathcal{F}_t\}$ be the filtration generated by current and past Y_t and θ and $\{\mathcal{G}_t\}$ the filtration generated by current and past Y_t only. The decision maker’s continuation value process (V_t) is adapted to $\{\mathcal{G}_t\}$ and

⁴Researchers Maccheroni et al. (2013) have a revealing but different type of limiting characterization of risk aversion and smooth ambiguity in a static environment with small changes in the uncertainty exposure from a given baseline specification.

satisfies a backward stochastic differential equation:

$$dV_t = V_t \nu_t(\theta) dt + V_t \varsigma_t \cdot dW_t, \quad V_T \text{ given,}$$

where $\{W_t : t \geq 0\}$ is a standard Brownian motion relative to $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbf{P})$. The continuation value process will be used to give recursive construction of the decision-maker's preference over consumption processes. We will derive a HJB equation for the $\{V_t\}$ to incorporate risk aversion, smooth ambiguity aversion and/or prior misspecification aversion.

We find it most convenient to use the logarithm of the continuation value $v_t = \log V_t$, which with an application of Ito's Lemma satisfies

$$dv_t = \nu_t(\theta) dt - \frac{1}{2} |\varsigma_t|^2 dt + \varsigma_t \cdot dW_t.$$

This evolution captures the familiar lognormal adjustment.

Motivated by the analysis in sections 3 through 5, we define risk and an ambiguity adjustment operators:

$$\begin{aligned} \mathbb{R}_{\epsilon,t}^*(v_{t+\epsilon}) &= \frac{1}{1-\gamma} \log E(\exp[(1-\gamma)v_{t+\epsilon}] | \mathcal{F}_t), \\ \mathbb{A}_{\epsilon,t}^*(w_{\epsilon,t}) &= -\frac{\epsilon}{\alpha_h} \log E\left[\exp\left(-\frac{\alpha_h}{\epsilon} w_{\epsilon,t}\right) | \mathcal{G}_t\right], \end{aligned}$$

where $w_{\epsilon,t}$ is \mathcal{F}_t measurable. In what follows, let

$$v_{\epsilon,t}^* = \mathbb{A}_{\epsilon,t}^* \circ \mathbb{R}_{\epsilon,t}^*(v_{t+\epsilon})$$

and $V_{\epsilon,t}^* = \exp(v_{\epsilon,t}^*)$.

We consider a preference specification that includes intermediate consumption and subjective discounting. Suppose that V_t satisfies the commonly used and convenient recursion:

$$V_t = \left([1 - \exp(-\delta\epsilon)] (C_t)^{1-\rho} + \exp(-\delta\epsilon) (V_{\epsilon,t}^*)^{1-\rho} \right)^{\frac{1}{1-\rho}},$$

where $\delta > 0$ is the subjective discount rate, $\{C_t : t \geq 0\}$ is a consumption process adapted to $\{\mathcal{G}_t : t \geq 0\}$, and $\frac{1}{\rho}$ is an intertemporal elasticity of substitution. Using the homogeneity

property of this recursion and taking logarithms, write

$$0 = \frac{1}{1-\rho} \log \left[[1 - \exp(-\delta\epsilon)] \left(\frac{C_t}{V_t} \right)^{1-\rho} + \exp(-\delta\epsilon) \left(\frac{V_{\epsilon,t}^*}{V_t} \right)^{1-\rho} \right]. \quad (10)$$

We derive (heuristically) the limit of the right-hand side of the equal sign of (10). This will characterize a restriction on local evolution of the continuation value process. Notice that

$$\begin{aligned} \mathbb{R}_{\epsilon,t}^* (v_{t+\epsilon} - v_t) &= \mathbb{R}_{\epsilon,t}^* (v_{t+\epsilon}) - v_t, \\ \mathbb{A}_{\epsilon,t}^* (w_{\epsilon,t} - v_t) &= \mathbb{A}_{\epsilon,t}^* (w_{\epsilon,t}) - v_t, \end{aligned}$$

where v_t is \mathcal{G}_t measurable. Extending our previous calculations, we first study the local risk adjustment, $\mathbb{R}_{\epsilon,t}^* (v_{t+\epsilon} - v_t)$:

$$\begin{aligned} &\lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \left(\frac{1}{1-\gamma} \right) \log E(\exp[(1-\gamma)(v_{t+\epsilon} - v_t)] | \mathcal{F}_t) \\ &= \nu_t(\theta) - \frac{\gamma}{2} |\varsigma_t|^2. \end{aligned}$$

Let $w_{\epsilon,t} = \mathbb{R}_{\epsilon,t}^* (v_{t+\epsilon})$ and $w_{0,t} = v_t$, and compute

$$\begin{aligned} \lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \mathbb{A}_{\epsilon,t}^* (w_{\epsilon,t} - v_t) &= -\frac{1}{\alpha_h} \log E(\exp[-\alpha_h \nu_t(\theta)] | \mathcal{G}_t) - \frac{\gamma}{2} |\varsigma_t|^2. \end{aligned}$$

With these intermediate calculations, we construct the local counterpart to (10) by dividing both sides of the equation by ϵ and taking small ϵ limits resulting in

$$\begin{aligned} 0 &= \frac{\delta}{1-\rho} \left[\left(\frac{C_t}{V_t} \right)^{1-\rho} - 1 \right] \\ &\quad - \frac{1}{\alpha_h} \log E(\exp[-\alpha_h \nu_t(\theta)] | \mathcal{G}_t) - \frac{\gamma}{2} |\varsigma_t|^2. \end{aligned} \quad (11)$$

The first term on the right-hand side comes from discounting, intermediate consumption

and intertemporal substitution. The second term captures uncertainty to the drift induced by the unknown parameter vector θ . We motivate this either as a robust prior/posterior adjustment or as an aversion to “smooth ambiguity” in the unknown drift. The third term captures the local risk adjustment coming from exposure to the underlying Brownian motion.

Using this limiting recursion for a HJB equation requires that we use a value function and its derivatives to deduce formulas for $\nu_t(\theta)$ and ς_t as functions of the relevant Markov state vector.

We conclude with four connections to related literature in control theory and economics.

Remark 5.1. *Recursion (11) provides a continuous-time analog to discrete time recursions in Hayashi and Miao (2011) and Ju and Miao (2012). Moreover, it shows how to incorporate ambiguity aversion into the continuous-time specifications of Duffie and Epstein (1992) and Duffie and Lions (1992).*

Remark 5.2. *The ambiguity adjustment:*

$$-\frac{1}{\alpha_h} \log E(\exp[-\alpha_h \nu_t(\theta)] | \mathcal{G}_t) \tag{12}$$

for the drift is a smooth counterpart to a continuous-time ambiguity adjustment with period-by-period constraints when the random set: $H_t = \{\nu_t(\theta) : \theta \in \Theta\}$ is a compact subset of a Euclidean space with probability one. As α_h becomes arbitrarily large, the smooth ambiguity adjustment converges to the minimum over H_t . The result is a continuation value recursion of the form considered by Chen and Epstein (2002).

More generally, $E(\exp[-\alpha_h \nu_t(\theta)] | \mathcal{G}_t)$ is the moment generating function for ν_t when viewed as a function of θ conditioned on \mathcal{G}_t . In general, a moment generating function is not guaranteed to be finite for all α_h . Even when the moment generating function is finite for small α_h , there may only be a compact interval of α_h 's for which the function remains finite. The existence of a finite upper bound for α_h has a recognizable connection to breakdown points in dynamic control theory.

Remark 5.3. *Recursive equation (11) also gives the continuous-time counterpart to a discrete-time recursions in Hansen and Sargent (2007) that captures two forms of robustness. These researchers also note the potentially stark asymmetry between how likelihood and priors contribute to overall entropy and propose separate penalizations. As we discussed previously, smooth ambiguity adjustment (12) has a dual interpretation as the outcome of prior/posterior uncertainty adjustment.*

For the likelihood adjustment, there is a well known link between risk sensitivity and robustness dating back to Jacobson (1973).⁵ Capture the likelihood uncertainty by representing $dW_t = h_t dt + dW_t^h$ under a change of probability measure for which dW_t^h is a Brownian increment and h_t is a local drift distortion. This adjustment alters the implied value function v_t by $\varsigma_t \cdot h_t$. Minimizing the local evolution of the value function by the choice of h_t subject to a penalization, $\frac{\kappa_f}{2}|h_t|^2$, illustrates a link between robustness and risk sensitivity that is familiar in the control literature.

Remark 5.4. To construct the filtration $\{\mathcal{G}_t : t \geq 0\}$ in practice and use it for solving a control problem we must produce a recursive solution for a filtering or estimation problem. While we posed the analysis as one for which θ is an unknown parameter, in fact this parameter could be an actual process designed to capture time variation in some underlying parameters. The sigma algebra \mathcal{G}_t could condition on the entire process or just the process up to time t . Examples of recursive constructions include Zakai equations, Kalman filtering or Wohham filtering depending on the application.

6 Supplemental material

6.1 Problem 2.1

As noted in section 1, there is an equivalent way to represent the relative entropy (1) in terms of posteriors and predictive densities:

$$D(\mathbf{P} \mid \hat{\mathbf{P}}) = \int_{\mathcal{Y}} \int_{\Theta} \log \left[\frac{d\pi^+}{d\hat{\pi}^+}(\theta \mid y) \right] \pi^+(d\theta \mid y) \phi(y) \tau(dy) \\ + \int_{\mathcal{Y}} \log \left[\frac{\phi(y)}{\hat{\phi}(y)} \right] \phi(y) \tau(dy).$$

Provided that priors π and $\hat{\pi}$ are absolutely continuous, the same can be said of posteriors. Since the objective of Problem 2.1 can be expressed in terms of the predictive density, it follows from this entropy decomposition that minimization will distort only this density and not the posterior distribution.

⁵For a continuous-time analysis see, for instance, James (1992).

6.2 Section 4 derivations

In section 4 we asserted that

$$\log B_t = \log Y_t + (1 - t)m_t + \mathbf{b}_t,$$

where

$$\mathbf{b}_{t+\epsilon} - \mathbf{b}_t = \frac{\epsilon^2}{2\kappa(\epsilon)} q_t \left(\frac{(1-t)q_t + |\varsigma|^2}{|\varsigma|^2 + \epsilon q_t} \right)^2 + \frac{\epsilon|\varsigma|^2}{2}$$

and where $\mathbf{b}_1 = 0$. To derive this recursive formula, note that

$$\begin{aligned} E(\log B_{t+\epsilon} \mid \mathcal{F}_t^\epsilon) &= \log Y_t + (1-t)m_t \\ &\quad + \epsilon \left[\frac{(1-t)q_t + |\varsigma|^2}{|\varsigma|^2 + \epsilon q_t} \right] (\theta - m_t) + \mathbf{b}_{t+\epsilon} - \frac{\epsilon|\varsigma|^2}{2}. \end{aligned}$$

Now apply

$$\begin{aligned} \mathbb{B}_{\epsilon,t} [E(\log B_{t+\epsilon} \mid \mathcal{F}_t^\epsilon)] &= \log Y_t + (1-t)m_t \\ &\quad + \mathbf{b}_{t+\epsilon} - \frac{\epsilon|\varsigma|^2}{2} - \frac{\epsilon^2}{2\kappa(\epsilon)} q_t \left(\frac{(1-t)q_t + |\varsigma|^2}{|\varsigma|^2 + \epsilon q_t} \right)^2. \end{aligned}$$

The recursion follows from our formula for $\log B_t$.

In section 4 we also claimed that

$$\log A_t = \log Y_t + (1-t)m_t + \mathbf{a}_t$$

where

$$\begin{aligned} \mathbf{a}_{t+\epsilon} - \mathbf{a}_t &= \\ &\quad \frac{1}{2} \left(\frac{|\varsigma|^2 + (1-t)q_t}{|\varsigma|^2 + \epsilon q_t} \right)^2 [\epsilon(\gamma - 1)|\varsigma|^2 + \alpha(\epsilon)\epsilon^2 q_t] \\ &\quad + \frac{\epsilon|\varsigma|^2}{2} \end{aligned}$$

and $\mathbf{a}_1 = 0$. To deduce this recursion, we first compute

$$\begin{aligned}
\log \mathbb{R}_{\epsilon,t}(A_{t+\epsilon}) &= \frac{1}{1-\gamma} \log E[(A_{t+\epsilon})^{1-\gamma} | \mathcal{F}_t^\epsilon] \\
&= \log Y_t + (1-t)m_t \\
&\quad + \epsilon \left[\frac{|\varsigma|^2 + (1-t)q_t}{|\varsigma|^2 + \epsilon q_t} \right] (\theta - m_t) \\
&\quad + \frac{\epsilon(1-\gamma)}{2} \left(\frac{|\varsigma|^2 + (1-t)q_t}{|\varsigma|^2 + \epsilon q_t} \right)^2 |\varsigma|^2 \\
&\quad + \mathbf{a}_{t+\epsilon} - \frac{\epsilon|\varsigma|^2}{2}.
\end{aligned}$$

Applying the operator $\mathbb{A}_{\epsilon,t}$ and the recursive equation (9), the recursion follows.

The authors thank Michael Barnett, William Brock, Peter Hansen, Massimo Marinacci, and Grace Tsiang for helpful comments and the Alfred P. Sloan Foundation for the support.

References

- Basar, Tamar and Pierre Bernhard. 1995. *H^∞ -Optimal Control and Related Minimax Problems: A Dynamic Game Theory Approach*. Boston, MA: Birkhauser, second ed.
- Chen, Zengjing and Larry Epstein. 2002. Ambiguity, Risk, and Asset Returns in Continuous Time. *Econometrica* 70 (4):1403–1443.
- Davis, Donald B. and M.-Elisabeth Paté-Cornell. 1994. A Challenge to the Compound Lottery Axiom: A two-stage normative structure and comparison to other theories. *Theory and Decision* 37 (3):267–309.
- Doyle, John C, Keith Glover, Pramod P. Khargonekar, and Bruce Francis. 1989. State Space Solutions for H_2 and H_∞ Control Problems. *IEEE Transactions for Automatic Control* 34 (8):831–847.
- Duffie, Darrel and Pierre-Louis Lions. 1992. PDE solutions of stochastic differential utility. *Journal of Mathematical Economics* 21 (6):577 – 606.
- Duffie, Darrell and Larry Epstein. 1992. Stochastic Differential Utility. *Econometrica* 60 (2):353–94.
- Hansen, Lars Peter and Thomas J. Sargent. 2007. Recursive Robust Estimation and Control without Commitment. *Journal of Economic Theory* 136 (1):1–27.
- . 2011. Robustness and Ambiguity in Continuous Time. *Journal of Economic Theory* 146 (3):1195–1223.
- Hayashi, Takashi and Jianjun Miao. 2011. Intertemporal substitution and recursive smooth ambiguity preferences. *Theoretical Economics* 6 (3):423–472.
- Jacobson, D. H. 1973. Optimal Stochastic Linear Systems with Exponential Performance Criteria and Their Relation to Deterministic Differential Games. *IEEE Transactions for Automatic Control* AC-18:1124–131.
- James, M. R. 1992. Asymptotic Analysis of Nonlinear Stochastic Risk-Sensitive Control and Differential Games. *Mathematics of Control, Signals and Systems* 5 (4):401–417.
- Ju, Nengjiu and Jianjun Miao. 2012. Ambiguity, Learning, and Asset Returns. *Econometrica* 80 (2):559–591.

- Klibanoff, P., M. Marinacci, and S. Mukerji. 2005. A Smooth Model of Decision Making Under Uncertainty. *Econometrica* 73 (6):1849–1892.
- Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini. 2006. Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. *Econometrica* 74 (6):1147–1498.
- Maccheroni, Fabio, Massimo Marinacci, and Dorian Ruffino. 2013. Alpha as Ambiguity: Robust Mean-Variance Portfolio Analysis. *Econometrica* 81 (3):1075–1113.
- Petersen, I.R., M.R. James, and P. Dupuis. 2000. Minimax Optimal Control of Stochastic Uncertain Systems with Relative Entropy Constraints. *Automatic Control, IEEE Transactions on* 45 (3):398–412.
- Segal, Uzi. 1990. Two-Stage Lotteries without the Reduction Axiom. *Econometrica* 58 (2):349–377.
- Skiadas, Costis. 2013. Smooth Ambiguity Aversion toward Small Risks and Continuous-Time Recursive Utility. *Journal of Political Economy* 121 (4):775–792.
- Strzalecki, Tomasz. 2011. Axiomatic Foundations of Multiplier Preferences. *Econometrica* 79 (1):47–73.