

Robustness and ambiguity in continuous time*

Lars Peter Hansen[†]

University of Chicago and NBER

Thomas J. Sargent[‡]

New York University and Hoover Institution

January 27, 2011

Abstract

We use statistical detection theory in a continuous-time environment to provide a new perspective on calibrating a concern about robustness or an aversion to ambiguity. A decision maker repeatedly confronts uncertainty about state transition dynamics and a prior distribution over unobserved states or parameters. Two continuous-time formulations are counterparts of two discrete-time recursive specifications of [16]. One formulation shares features of the smooth ambiguity model of [24, 25]. Here our statistical detection calculations guide how to adjust contributions to entropy coming from hidden states as we take a continuous-time limit.

KEYWORDS: ambiguity, robustness, hidden Markov model, likelihood function, entropy, statistical detection error, smooth ambiguity

JEL CODES: C61, D81, D83

*We thank Anmol Bhandari, Ramon Marimon, Cecilia Parlatore Siritto, an editor, and two referees for very helpful comments on earlier drafts. Lars Peter Hansen's contribution was supported in part by the Templeton Foundation.

[†]e-mail: lhansen@uchicago.edu

[‡]Corresponding author. E-mail: ts43@nyu.edu, phone: 212 998 3548, fax: 212 995 4186.

1 Introduction

Following Frank Knight, I use the word “uncertainty” to mean random variation according to an unknown probability law. . . . An increase in uncertainty aversion increases the multiplicity of subjective distributions. An uncertainty neutral decision maker has only one subjective distribution and so is Bayesian. [5], p. 1

We contribute to this volume honoring Truman Bewley, whom we both greatly admire, by refining some of our recent efforts to construct workable models of what [3, 4, 5] called Knightian uncertainty.¹ Bewley’s papers revoked Savage’s completeness axiom and modeled choices among incompletely ordered alternatives by assuming that the decision maker has multiple probability models and a ‘status quo’ choice. The status quo is preferred to another choice unless it yields higher expected utility under all of those probability models. Our work shares motivations with Bewley’s, but we take a different approach to modeling an agent who uses multiple probability models. Our decision maker starts with a single approximating model. But because he does not trust that model, the decision maker investigates the utility consequences of other models in hopes of constructing a decision rule that is robust to misspecification.

We propose two continuous-time recursive specifications of robust control problems with hidden state variables, some of which can be interpreted as unknown parameters. Each specification is a continuous-time limit of a discrete-time problem proposed by [16]. The first is a recursive counterpart to formulations in the robust control literature with a continuation value function that depends on hidden states, while the second reinterprets the recursive utility model of [26] and [24, 25] in terms of concerns about robustness when a continuation value function depends only on observed information.²

We quantify the decision maker’s ambiguity by constructing a worst-case model and assessing its statistical discrepancy from the benchmark model. We do this by modifying statistical discrimination methods developed by [29] and [1]³ so that they apply to our continuous-time setting. We use these statistical model detection

¹For more testimonies to our admiration of Truman Bewley, see the chapters on what we call ‘Bewley models’ in [31], ch. 6 and [28], ch. 17.

²The first is a limiting version of recursions (20)-(21) of [16] and the second is a version of recursion (23).

³[29] and [1] built on [8].

calculations to formulate entropy penalties that enable us to represent preferences that express model ambiguity and concerns about statistical robustness. These statistical measures of model discrimination tell us how to calibrate a decision maker’s ambiguity. The ability to quantify model ambiguity is important because promising empirical results in [23], [9], and [18] make us want to understand how much ambiguity aversion is *a priori* reasonable.

The remainder of this paper is organized as follows. Sections 2 and 3 describe the stochastic model, information structure, and a continuous time filtering problem. Section 4 describes alternative representations of entropy and establishes that entropy over alternative posterior distributions is the pertinent concept for decision making. A peculiar consequence of this finding is that corrections representing smooth ambiguity over hidden states vanish in a continuous time limit. We trace this outcome to how various contributions to entropy are scaled with respect to the accumulation of information over time as reflected in likelihood ratios. Section 5 describes two robust control problems that sustain concerns about model misspecification even in their continuous-time limits. For the subsection 5.1 formulation in which continuation values depend only on observed information, a key aspect is how our statistical model detection calculations tell us to rescale with the passage of time those contributions to entropy that come from unknown distributions of hidden states. Section 6 extends the analysis to hidden Markov settings with time-varying states. We study models with both continuous hidden states (Kalman filtering) and discrete hidden states (Wonham filtering). Section 7 extends outcomes from the statistical detection literature to justify our way of rescaling contributions to entropy in the section 5.1 formulation in which continuation values depend only on observed information. Section 8 then describes how our proposal for rescaling relates to characterizations of smooth ambiguity. Section 9 offers some concluding remarks, while appendix A solves a discrete-state entropy problem.

2 The benchmark model specification

To set out essential components of our formulation, we start with a relatively simple stochastic specification that allows either for multiple models or for unknown parameters that generate a continuous-time trajectory of data. We have designed things so that recursive filtering can be readily characterized. Our continuous time specification

reveals important aspects of modeling and calibrating a preference for robustness. In section 6, we will consider extensions that incorporate more general learning about a possibly evolving hidden state.

Consider a continuous-time specification with model uncertainty. The state dynamics are described by a collection of diffusions indexed by ι :

$$dY_t = \mu(Y_t, \iota)dt + \sigma(Y_t)dW_t. \tag{1}$$

The parameter ι , which indexes a model, is hidden from the decision-maker, who does not observe Y . The matrix $\sigma(Y)$ is nonsingular, implying that the Brownian increment dW_t would be revealed if ι were known. Initially, we assume that the unknown model ι is in a finite set \mathcal{I} , a restriction we impose mainly for notational simplicity. Later we will also consider examples with a continuum of models, in which case we shall think of ι as an unknown parameter.

To characterize preferences under model uncertainty, we suppose that Y is beyond the control of the decision maker. We display Bellman equations that can be solved for value functions that express model specification doubts. We stress how these Bellman equations allow us to distinguish aspects of the model specification that concern the decision maker, in particular, the stochastic transition laws for the hidden states versus the decision maker's subjective prior over those hidden states. Two different risk-sensitivity parameters express these distinct specification suspicions. In this paper we focus on specifying continuous-time Bellman equations that express a decision maker's specification doubts; but we do not provide rigorous statements of the conditions for existence of solutions of the Bellman equations. We warn readers that existence is a loose end that will have to be addressed when applying our formulations.⁴ Nevertheless, for particular examples that especially interest us, we provide explicit solutions of the Bellman equations.

⁴We know from the structure of verification theorems that in general we would have to add some technical restrictions about transition dynamics, utility functions, and entropy constraints. Typically there are tradeoffs among these restrictions, and sometimes the special structure of the decision problem can be exploited to good advantage.

3 Filtering without ambiguity

A filtering problem in which the model ι is unknown to the decision-maker is central to our analysis. A key ingredient is the log-likelihood conditioned on the unknown model ι . For the moment, we ignore concerns about model misspecification.

3.1 Log-likelihoods

Log-likelihoods are depicted relative to some measure over realized values of the data. Consider a counterpart to (1) in which the drift is zero, but the same σ matrix governs conditional variability. This process has a local evolution governed by the stochastic differential equation

$$dY_t = \sigma(Y_t)dW_t.$$

We use a solution that takes Y_0 as given to induce a probability measure τ over vector valued continuous functions on $[0, t]$. Likelihoods are constructed from densities with respect to τ . In particular, the local evolution of the log-likelihood is⁵

$$d \log L_t(\iota) = \mu(Y_t, \iota)' [\sigma(Y_t)\sigma(Y_t)']^{-1} dY_t - \frac{1}{2} \mu(Y_t, \iota)' [\sigma(Y_t)\sigma(Y_t)']^{-1} \mu(Y_t, \iota) dt.$$

We initialize the likelihood by specifying $L_0(\iota)$. We stipulate that $L_0(\iota)$ includes contributions both from the density of Y_0 conditioned on ι and from a prior probability $\pi_0(\iota)$ assigned to model ι , and so $L_t(\iota)$ is a likelihood for $[Y_u : 0 < u \leq t]$ conditional on Y_0 and ι times a prior over ι conditioned on Y_0 .

3.2 Posterior

As a consequence, the date t posterior for ι is

$$\pi_t(\iota) = \frac{L_t(\iota)}{\sum_{\iota} L_t(\iota)}.$$

The date t posterior for the local mean of dY_t is

$$\bar{\mu}_t = \sum_{\iota} \pi_t(\iota) \mu(Y_t, \iota).$$

⁵For example, see [27] chapter 9.

Then

$$dY_t = \bar{\mu}_t dt + \sigma(Y_t) d\bar{W}_t, \quad (2)$$

and

$$d\bar{W}_t = \sigma(Y_t)^{-1} [\mu(Y_t, \iota) - \bar{\mu}_t] dt + dW_t \quad (3)$$

is a multivariate standard Brownian motion relative to a filtration \mathcal{Y} generated by Y that does not include knowledge of the model ι . Since he does not know the model ι , the increment $d\bar{W}_t$ is the date t innovation pertinent to the decision-maker. Equation (2) gives the stochastic evolution of Y in terms of this innovation. The decomposition of the increment to the innovation process \bar{W} with reduced information \mathcal{Y} into the two components on the right side of (3) will play a significant role later.

The evolution for likelihood \bar{L}_t of the mixture (of means) model (2)-(3) is

$$\begin{aligned} d \log \bar{L}_t &= (\bar{\mu}_t)' [\sigma(Y_t) \sigma(Y_t)']^{-1} dY_t - \frac{1}{2} (\bar{\mu}_t)' [\sigma(Y_t) \sigma(Y_t)']^{-1} \bar{\mu}_t dt \\ &= (\bar{\mu}_t)' [\sigma(Y_t) \sigma(Y_t)']^{-1} \sigma(Y_t) d\bar{W}_t + \frac{1}{2} (\bar{\mu}_t)' [\sigma(Y_t) \sigma(Y_t)']^{-1} \bar{\mu}_t dt. \end{aligned}$$

We shall use the following example.

Example 3.1. Consider an additively separable drift $\mu(Y, \iota) = \nu(Y) + \iota$ so that

$$dY_t = \nu(Y_t) dt + \iota dt + \sigma dW_t.$$

Suppose that $\mathcal{I} = \mathbb{R}^n$ for some n and impose a normal prior π_0 over ι , so that

$$\log \pi_0(\iota) = -\frac{1}{2} (\iota - \bar{\iota}_0)' \Lambda_0 (\iota - \bar{\iota}_0) + \text{constant}$$

where Λ_0 is the prior precision matrix and $\bar{\iota}_0$ is the prior mean.

For this example,

$$d \log L_t(\iota) = [\nu(Y_t) + \iota]' (\sigma \sigma')^{-1} dY_t - \frac{1}{2} [\nu(Y_t) + \iota]' (\sigma \sigma')^{-1} [\nu(Y_t) + \iota] dt.$$

We form a likelihood conditioned on Y_0 and use the prior over ι to initialize $\log L_0(\iota)$. Since the log-likelihood increment and the logarithm of the prior are both quadratic in ι , it follows that the posterior density for ι is normal.

To obtain the evolution of posterior probabilities, note that

$$dL_t(\iota) = L_t(\iota)[\nu(Y_t) + \iota]'(\sigma\sigma')^{-1}dY_t.$$

Integrating with respect to ι , we obtain

$$d\bar{L}_t = \left[\int L_t(\iota)[\nu(Y_t) + \iota]d\iota \right]' (\sigma\sigma')^{-1}dY_t.$$

Then the posterior probability density

$$\pi_t(\iota) = \frac{L_t(\iota)}{\bar{L}_t}$$

evolves as

$$\begin{aligned} d\pi_t(\iota) &= \frac{L_t(\iota)}{\bar{L}_t} \left([\nu(Y_t) + \iota]' - \int \frac{L_t(\iota)}{\bar{L}_t} [\nu(Y_t) + \iota]'d\iota \right) (\sigma\sigma')^{-1}dY_t \\ &\quad - \frac{L_t(\iota)}{\bar{L}_t} [\nu(Y_t) + \iota]'(\sigma\sigma')^{-1} \left(\int \frac{L_t(\iota)}{\bar{L}_t} [\nu(Y_t) + \iota]'d\iota \right)' dt \\ &\quad + \frac{L_t(\iota)}{\bar{L}_t} \left[\int \frac{L_t(\iota)}{\bar{L}_t} [\nu(Y_t) + \iota]'d\iota \right] (\sigma\sigma')^{-1} \left(\int \frac{L_t(\iota)}{\bar{L}_t} [\nu(Y_t) + \iota]'d\iota \right) dt \\ &= \pi_t(\iota) (\iota - \bar{\iota}_t)' (\sigma\sigma')^{-1} [dY_t - \nu(Y_t)dt - \bar{\iota}_t dt], \end{aligned} \tag{4}$$

where $\bar{\iota}_t$ is the posterior mean of ι . Multiplying both sides of (4) by ι and integrating with respect to ι , the posterior mean evolves as

$$d\bar{\iota}_t = \Sigma_t(\sigma\sigma')^{-1} [dY_t - \nu(Y_t)dt - \bar{\iota}_t dt],$$

where Σ_t is the posterior covariance matrix for ι .

To get a formula for Σ_t , first note that the evolution of the logarithm of the posterior density implied by (4) is

$$\begin{aligned} d \log \pi_t(\iota) &= (\iota - \bar{\iota}_t)' (\sigma\sigma')^{-1} [dY_t - \nu(Y_t)dt - \bar{\iota}_t dt] \\ &\quad - \frac{1}{2} (\iota - \bar{\iota}_t)' (\sigma\sigma')^{-1} (\iota - \bar{\iota}_t) dt. \end{aligned}$$

This follows because $dY_t - \nu(Y_t)dt - \bar{\iota}_t dt$ is a Brownian increment under the \mathcal{Y} filtration with instantaneous covariance $\sigma\sigma'$. The evolution for $d \log \pi_t(\iota)$ follows from Ito's

formula. Integrating between zero and t , the log-density is quadratic in ι , and hence ι is normally distributed. From the resulting quadratic form in ι , the date t precision matrix $\Lambda_t = (\Sigma_t)^{-1}$ is

$$\Lambda_t = \Lambda_0 + t(\sigma\sigma')^{-1},$$

where Λ_0 is the prior precision matrix.

4 Relative entropy

Statistical discrimination and large-deviation theories⁶ underlie relative entropy's prominent role in dynamic stochastic robust control theory. In this section, we construct relative entropies and discuss some of their implications in the continuous-time stochastic setting of sections 2 and 3. This will set the stage for explorations of alternative ways of scaling different contributions to entropy that will concern us in subsequent sections.

4.1 Factorization

Recall that $L_t(\iota)$ is the likelihood as a function of ι and scaled to incorporate the prior probability of ι . This likelihood depends implicitly on the record $y \doteq [Y_u, 0 \leq u \leq t]$ of observed states between zero and t . Let $f_t(y, \iota)$ be the density constructed so that $f_t(Y_u, 0 \leq u \leq t, \iota) = L_t(\iota)$, where y is our hypothetical realization of the Y process between dates zero and t . Similarly, the marginal density is

$$g_t(y) = \sum_{\iota} f_t(y, \iota),$$

which implies that $g_t(y) = \bar{L}_t$. Let $\tilde{f}_t(y, \iota)$ be an alternative joint density for ι and Y observed between dates zero and t . Then relative entropy is

$$\text{ent}(\tilde{f}_t) = \sum_{\iota} \int \left[\log \tilde{f}_t(y, \iota) - \log f_t(y, \iota) \right] \tilde{f}_t(y, \iota) d\tau(y).$$

⁶See [12] for an exposition and development of these tools.

We can factor the joint density $f_t(y, \iota)$ as

$$f_t(y, \iota) = g_t(y) \left[\frac{f_t(y, \iota)}{g_t(y)} \right] = g_t(y) \psi_t(\iota|y),$$

and similarly for the alternative $\tilde{\tau}$ densities. Notice that $\psi_t(\iota|y) \equiv \pi_t(\iota)$, the posterior probability of ι . These density factorizations give rise to an alternative representation of entropy:

$$\begin{aligned} \text{ent}(\tilde{g}_t, \tilde{\psi}_t) &= \int \left(\sum_{\iota} \tilde{\psi}_t(\iota|y) \left[\log \tilde{\psi}_t(\iota|y) - \log \psi_t(\iota|y) \right] \right) \tilde{g}_t(y) d\tau(y) \\ &+ \int [\log \tilde{g}_t(y) - \log g_t(y)] \tilde{g}_t(y) d\tau(y). \end{aligned}$$

For a fixed t , consider the following *ex ante* decision problem:

Problem 4.1.

$$\max_a \min_{\tilde{g}_t, \tilde{\psi}_t} \int \sum_{\iota} U[a(y), y, \iota] \tilde{\psi}_t(\iota|y) \tilde{g}_t(y) d\tau(y) + \theta \text{ent}(\tilde{g}_t, \tilde{\psi}_t)$$

where U is a concave function of action a .

This is a static *robust control* problem. We refer to it as an *ex ante* problem because the objective averages across *all* data that could *possibly* be realized. But the decision a depends only on the data that are *actually* realized. Just as in Bayesian analysis, prior to computing the worst case \tilde{g}_t we can find optimal a and $\tilde{\psi}_t$ by solving the following *conditional problem*:

Problem 4.2.

$$\max_a \min_{\tilde{\psi}_t} \sum_{\iota} U[a(y), y, \iota] \tilde{\psi}_t(\iota|y) + \theta \sum_{\iota} \tilde{\psi}_t(\iota|y) \left[\log \tilde{\psi}_t(\iota|y) - \log \psi_t(\iota|y) \right]$$

separately for each value of y without simultaneously computing \tilde{g}_t . Solving this conditional problem for each y gives a robust decision for a .

In the conditional problem, we perturb only the posterior $\psi_t(\iota|y)$.⁷ This simplification led [16] to perturb the *outcome* of filtering, namely, the posterior distribution,

⁷In practice it suffices to compute $a(y)$ for y given by the realization of Y between dates zero and t .

rather than the likelihood and prior.

Thus, we have established that for decision making, the pertinent concept of relative entropy over an interval of time reduces to a measure of entropy over the posterior:

$$\text{ent}^*(\tilde{\pi}_t) = \sum_{\iota} \tilde{\pi}_t(\iota) [\log \tilde{\pi}_t(\iota) - \log \pi_t(\iota)] \quad (5)$$

where $\pi_t(\iota) = \psi_t(\iota|Y_u, 0 \leq u \leq t)$.

4.2 Forward-looking and backward-looking concerns for misspecification

We follow [16] in considering forward-looking model misspecifications (i.e., misspecifications of *future* shocks to states conditional on the entire state) as well as backward-looking misspecifications (i.e., misspecifications of the distribution of hidden states arising from filtering). We say forward-looking because of how worst-case distortions of this type depend on value functions. We say backward-looking because filtering processes historical data.

To distinguish these two forms of misspecification, consider first a joint density for $Y_{t+\epsilon} - Y_t$, $[Y_u : 0 < u \leq t]$ and ι and conditioned on Y_0 . We decompose the process of conditioning in terms of the product of three densities: i) the density for $Y_{t+\epsilon} - Y_t$ conditioned on Y_t and ι , ii) the density for $[Y_u, 0 < u \leq t]$ conditioned on ι and Y_0 , and iii) a prior over ι conditioned on Y_0 . To characterize misspecifications using relative entropy, we consider contributions from all three sources. Since the decision-maker conditions on $[Y_u : 0 < u \leq t]$, in section 4.1 we established that it suffices to focus on the relative entropy of the date t “posterior” (conditioned on information available at time t) defined in (5). Specifically the relative entropy measure is:

$$\begin{aligned} & \sum_{\iota} \tilde{\pi}_t(\iota) \int \left[\log \tilde{\ell}_{\epsilon,t}(z|y, \iota) - \log \ell_{\epsilon,t}(z|y, \iota) \right] \tilde{\ell}_{\epsilon,t}(z|y, \iota) dz \\ & + \sum_{\iota} \tilde{\pi}_t(\iota) [\log \tilde{\pi}_t(\iota) - \log \pi_t(\iota)] \end{aligned} \quad (6)$$

where $\ell_{\epsilon,t}(z|y, \iota)$ is the density for $Y_{t+\epsilon} - Y_t$ conditioned on $[Y_u : 0 < u \leq t]$ and ι and z is a realized value of $Y_{t+\epsilon} - Y_t$. Thus, in our study of misspecification we focus on two component densities: a) the density for $Y_{t+\epsilon} - Y_t$ conditioned on $[Y_u : 0 < u \leq t]$

and ι , and b) the density for ι conditioned on $[Y_u : 0 < u \leq t]$, which is the posterior for ι given date t information. Since we are using date t as the point of reference, we refer to distortions in a) as forward-looking and distortions in b) as backward looking.

Under the benchmark or approximating model, the density $\ell_{\epsilon,t}(z|Y_0^t, \iota)$ is (approximately) normal with mean $\epsilon\mu(Y_t, \iota)$ and variance $\epsilon\sigma(Y_t)^2$ and Y_0^t is short-hand notation for $[Y_u : 0 < u \leq t]$. In the small ϵ or continuous-time limit, the distortions have a very simple structure. The original Brownian motion increment dW_t is altered by appending a drift, as occurs in the continuous time case with observed Markov states discussed by [22], [7], [1], and [19]. That is, in the distorted model, $dW_t = h_t(\iota)dt + d\tilde{W}_t$, where $d\tilde{W}_t$ is a standard Brownian increment. The corresponding contribution to relative entropy scaled by $\frac{1}{\epsilon}$ is

$$\lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \int \left[\log \tilde{\ell}_{\epsilon,t}(z|Y_0^t, \iota) - \log \ell_{\epsilon,t}(z|Y_0^t, \iota) \right] \tilde{\ell}_{\epsilon,t}(z|Y_0^t, \iota) dz = \frac{1}{2} |h_t(\iota)|^2.$$

Thus, the forward-looking contribution to entropy is $\frac{1}{2}|h_t(\iota)|^2 dt$, which scales linearly with the time increment dt , and the limiting version of the combined entropy in (6) is (heuristically)

$$\sum_{\iota} \tilde{\pi}_t(\iota) \left[\frac{1}{2} |h_t(\iota)|^2 dt + \log \tilde{\pi}_t(\iota) - \log \pi_t(\iota) \right].$$

A key observation for us is that because it reflects the history of data up to t , the contribution to entropy that comes from distorting the distribution $\pi_t(\iota)$ of ι , namely,

$$\sum_{\iota} \tilde{\pi}_t(\iota) [\log \tilde{\pi}_t(\iota) - \log \pi_t(\iota)],$$

does *not* scale linearly with the time increment dt . As we explain in detail in the next subsection, this has the consequence of making it very costly in terms of entropy to distort the distribution of ι , causing adjustments for an unknown distribution of ι to vanish in a continuous-time limit.

4.3 The scaling problem

Associated with the two distortions h_t and $\tilde{\pi}_t$ is a distorted conditional mean:

$$\tilde{\mu}_t = \sum_{\iota} \tilde{\pi}_t(\iota) [\sigma(Y_t)h_t(\iota) + \mu(Y_t, \iota)].$$

Notice that $\tilde{\mu}_t$ is influenced by both $h_t(\iota)$ and $\tilde{\pi}_t$.

The difference in how the entropies are scaled implies that to achieve a given distorted mean $\tilde{\mu}_t$ at a minimum cost in terms of entropy

$$\sum_{\iota} \tilde{\pi}_t(\iota) \left(\frac{1}{2} |h_t(\iota)|^2 dt + [\log \tilde{\pi}_t(\iota) - \log \pi_t(\iota)] \right)$$

we should set $h_t(\iota) = \sigma(Y_t)^{-1}(\tilde{\mu}_t - \bar{\mu}_t)$ and leave the model probabilities unchanged, thus setting $\tilde{\pi}_t(\iota) = \pi_t(\iota)$. These choices make the resulting minimized entropy be scaled linearly by dt .

This outcome has important consequences for recursive specifications of decision problems that are designed to express a decision maker's concerns about misspecification. When a single entropy restriction is imposed over the joint distribution of the data and model, and when the objective of a probability distorting minimizing player is to distort the drift $\bar{\mu}_t$, it is evidently too costly in terms of entropy to induce *any* distortion to the current-period posterior as an intermediate step that aims ultimately to distort $\bar{\mu}_t$.

We will explore two responses that are designed to allow a decision maker to express ambiguity about the prior $\pi_t(\iota)$ in a continuous time limit. One reduces the cost, while another enhances the benefits that accrue to the fictitious evil agent imagined to be perturbing the prior over the hidden state ι . Our formulation in section 5.1 weights the entropy contribution from $\tilde{\pi}_t$ by dt ; our formulation in section 5.2 makes changes in $\tilde{\pi}_t$ more consequential to the decision-maker by altering the forward-looking objective function. These two alternative amendments lead to continuous-time versions of two distinct discrete-time formulations in [16], with section 5.1 corresponding to recursion (23) of [16] and section 5.2 corresponding to their recursion (20)-(21).⁸

⁸A third approach is to apply robust control theory to the decision problem by directly perturbing the transition probabilities for $Y_{t+\epsilon} - Y_t$ conditioned on $[Y_u : 0 \leq u \leq t]$ without explicit reference to an estimation problem. In this approach, learning and filtering are regarded just as means to compute benchmark probabilities but are not used independently to make contributions to perturbed

5 Robustness-inspired alternatives to expected utility

In this section, we describe recursive representations of preferences that express concerns about robustness in a continuous time setting. We do this by constructing continuous-time versions of two discrete-time recursive specifications of preferences for robustness proposed by [16]. The two specifications differ in the information sets used to compute benchmark value functions.

5.1 Continuation values that don't depend on ι

First, we consider a counterpart to recursion (23) in [16], where continuation values do not depend on the hidden state. The evolution of a continuation-value process V adapted to \mathcal{Y} is

$$dV_t = \bar{\eta}_t dt + \varsigma_t \cdot d\bar{W}_t.$$

For (time additive) discounted utility, the drift satisfies

$$\bar{\eta}_t = \delta V_t - \delta U(C_t),$$

where δ is the subjective rate of discount and U is the instantaneous utility function.

Recall that

$$d\bar{W}_t = \sigma(Y_t)^{-1}[\mu(Y_t, \iota) - \bar{\mu}_t]dt + dW_t. \tag{7}$$

It is convenient to view the right-hand side of (7) as a two-stage lottery. One stage is modeled as a Brownian increment and the other stage is a lottery over ι with probabilities given by the date t posterior distribution. This distribution over ι induces a distribution for the increment $\sigma(Y_t)^{-1}[\mu(Y_t, \iota) - \bar{\mu}_t]dt$.⁹ In turn two distortions to $d\bar{W}_t$.

Remark 5.1. *The two contributions to $d\bar{W}_t$ displayed on the right side of (7) behave very differently over small time intervals because the variance of a Brownian increment dW_t scales with dt , while the variance of $\sigma(Y_t)^{-1}[\mu(Y_t, \iota) - \bar{\mu}_t]dt$ scales with probabilities.*

⁹ As we shall see in section 8, the behavior of the component $\sigma(Y_t)^{-1}[\mu(Y_t, \iota) - \bar{\mu}_t]dt$ on the right side of (7) will also motivate our way of parameterizing ambiguity about the distribution of the hidden state or parameter.

$(dt)^2$. For small time intervals, variation in the Brownian increment dW_t dominates the variation in $d\bar{W}_t$.

We next consider distortions to the two components on the right side of (7).

5.1.1 Distortion to dW_t

Introduce a drift distortion $h_t(\iota)dt$ so that $dW_t = h_t(\iota)dt + d\hat{W}_t$, where $d\hat{W}_t$ is a standard Brownian increment. Recall that local entropy is $\frac{|h_t(\iota)|^2}{2}dt$. Find a worst-case model by solving

Problem 5.2.

$$\min_{h_t(\iota)} \left\{ \varsigma_t \cdot h_t(\iota) + \frac{\theta_1}{2} |h_t(\iota)|^2 \right\},$$

where θ_1 is an entropy penalty parameter.

The minimizer

$$\tilde{h}_t = -\frac{1}{\theta_1} \varsigma_t,$$

is independent of ι . Our first adjustment for robustness of the continuation value drift is

$$\bar{\eta}_t = \delta V_t - \delta U(C_t) + \frac{|\varsigma_t|^2}{2\theta_1}.$$

As discussed in [19], this distorted drift outcome is a special case of the variance multiplier specification of [11] in which our $\frac{1}{\theta_1}$ becomes their multiplier on the conditional variance of the continuation value.¹⁰

5.1.2 Distortion to $\sigma(Y_t)^{-1}[\mu(Y_t, \iota) - \bar{\mu}_t]$

We distort the other component of $d\bar{W}_t$ on the right side of (7) by altering the posterior $\pi_t(\iota)$. Once again, we penalize distortions in terms of relative entropy. Since changing π_t alters the drift, we adopt a common scaling for both the instantaneous return and the contribution (5) to relative entropy. We accomplish this by multiplying the object $\text{ent}^*(\tilde{\pi}_t)$ defined in equation (5) by dt in order to constructing a penalty for the minimizing agent. This adjustment sharply distinguishes the treatments of the two entropy penalties that [16] include in their discrete-time specification, one that measures unknown dynamics, the other that measures unknown states. By altering

¹⁰[11] construct continuous-time limiting versions of the recursive utility specification of [26].

the filtering probabilities we induce a change in the drift of the continuation value given by

$$(\varsigma_t)' \sigma(Y_t)^{-1} \sum_{\iota} [\tilde{\pi}_t(\iota) - \pi_t(\iota)] \mu(Y_t, \iota).$$

Since this term and the relative entropy measure $\text{ent}^*(\tilde{\pi}_t)$ are scaled by dt , we are led to solve:

Problem 5.3.

$$\min_{\tilde{\pi}_t} (\varsigma_t)' \sigma(Y_t)^{-1} \sum_{\iota} [\tilde{\pi}_t(\iota) - \pi_t(\iota)] \mu(Y_t, \iota) + \theta_2 \text{ent}^*(\tilde{\pi}_t)$$

where θ_2 is a penalty parameter.

The minimized objective is

$$(\varsigma_t)' \sigma(Y_t)^{-1} \bar{\mu}_t + \theta_2 \log \left(\sum_{\iota} \pi_t(\iota) \exp \left[-\frac{1}{\theta_2} (\varsigma_t)' \sigma(Y_t)^{-1} \mu(Y_t, \iota) \right] \right).$$

Result 5.4. *The drift $\bar{\eta}_t$ and local Brownian exposure vector ς_t for the continuation value process $\{V_t\}$ satisfy*

$$\begin{aligned} \bar{\eta}_t = & \delta V_t - \delta U(C_t) + \frac{|\varsigma_t|^2}{2\theta_1} + (\varsigma_t)' \sigma(Y_t)^{-1} \bar{\mu}_t \\ & + \theta_2 \log \left(\sum_{\iota} \pi_t(\iota) \exp \left[-\frac{1}{\theta_2} (\varsigma_t)' \sigma(Y_t)^{-1} \mu(Y_t, \iota) \right] \right). \end{aligned} \quad (8)$$

The total contribution from the two distortions, namely,

$$\frac{|\varsigma_t|^2}{2\theta_1} + \theta_2 \log \left(\sum_{\iota} \pi_t(\iota) \exp \left[-\frac{1}{\theta_2} (\varsigma_t)' \sigma(Y_t)^{-1} \mu(Y_t, \iota) \right] \right)$$

is necessarily nonnegative.

Example 5.5. *Reconsider the filtering problem in example 3.1. Given that the bench-*

mark posterior is normal, formulas for exponentials of normals imply

$$\begin{aligned} & \theta_2 \log \left(\int \pi_t(\iota) \exp \left[-\frac{1}{\theta_2} (\varsigma_t)' \sigma(Y_t)^{-1} \mu(Y_t, \iota) \right] d\iota \right) \\ &= - (\varsigma_t)' \sigma(Y_t)^{-1} \bar{\mu}_t + \frac{1}{2\theta_2} (\varsigma_t)' \sigma(Y_t)^{-1} \Sigma_t [\sigma(Y_t)^{-1}]' \varsigma_t. \end{aligned} \quad (9)$$

For this example, the composite penalty can be decomposed as

$$\begin{array}{ccc} \frac{|\varsigma_t|^2}{2\theta_1} & + & \frac{1}{2\theta_2} (\varsigma_t)' \sigma(Y_t)^{-1} \Sigma_t [\sigma(Y_t)^{-1}]' \varsigma_t \\ \uparrow & & \uparrow \\ \text{misspecified} & & \text{misspecified} \\ \text{dynamics} & & \text{state estimation} \end{array}$$

5.2 Continuation values that do depend on ι

We now turn to a continuous-time version of recursions (20)-(21) in [16] where continuation values *do* depend on the hidden state.¹¹ This structure will impel us to compute two value functions, one to be called V that conditions on the hidden state, another to be called \hat{V} that conditions only on observable information. Suppose that the continuation value is constructed knowing ι . Thus, we write

$$dV_t(\iota) = \eta_t(\iota) dt + \varsigma_t(\iota) \cdot dW_t.$$

5.2.1 Distortion to dW_t

We again append a drift to the Brownian increment subject to an entropy penalty and solve:

Problem 5.6.

$$\min_{h_t(\iota)} \left\{ \varsigma_t(\iota) \cdot h_t(\iota) + \frac{\theta_1}{2} |h_t(\iota)|^2 \right\},$$

where θ_1 is an entropy penalty parameter.

The minimizer is

$$\tilde{h}_t(\iota) = -\frac{1}{\theta_1} \varsigma_t(\iota).$$

¹¹[15] describe a distinct formulation that is explicitly linked to the robust control literature. It imposes commitment to prior distortions.

The optimized $h_t(\iota)$ will typically depend on ι . As an outcome, we are led to the equation:

$$\eta_t(\iota) = \delta V_t(\iota) - \delta U(C_t) + \frac{|\varsigma_t(\iota)|^2}{2\theta_1}.$$

We compute the continuation value subject to this restriction, imposing an appropriate terminal condition. In this way, we adjust for robustness with respect to the motion of the state variables.

5.2.2 Distortion to $\sigma(Y_t)^{-1}[\mu(Y_t, \iota) - \bar{\mu}_t]$

To adjust for robustness in *estimation*, i.e., for robustness with respect to the probability density of ι , we solve

Problem 5.7.

$$\hat{V}_t = \min_{\tilde{\pi}_t} \sum_{\iota} \tilde{\pi}_t(V_t(\iota) + \theta_2 [\log \tilde{\pi}_t(\iota) - \log \pi_t(\iota)])$$

Result 5.8. *The date t continuation value \hat{V}_t that solves problem 5.7 is*

$$\hat{V}_t = -\theta_2 \log \left(\sum_{\iota} \exp \left[-\frac{1}{\theta_2} V_t(\iota) \right] \pi_t(\iota) \right).$$

where the drift $\eta_\tau(\iota)$ and shock exposure $\varsigma_\tau(\iota)$ for the complete information continuation value process $\{V_\tau(\iota) : \tau \geq t\}$ are restricted by

$$\eta_\tau(\iota) = \delta V_\tau(\iota) - \delta U(C_\tau) + \frac{|\varsigma_\tau(\iota)|^2}{2\theta_1}.$$

We use the continuation value \hat{V}_t to rank alternative consumption processes from the perspective of time t . Since we apply the robustness adjustment for estimation to the continuation value and not its stochastic increment, we do not scale relative entropy linearly by dt in the minimization problem on the right side of the equation defining problem 5.7.

In this formulation, the value function for the fictitious evil agent can be constructed recursively, but the induced preferences for the decision-maker change over time. As in [30], we confront this intertemporal change in preferences by having our decision-maker play a dynamic game with future versions of himself.

5.3 Examples

We illustrate the two recursions when the state dynamics in example 3.1 are specialized so that $U(C_t) = \log C_t = H \cdot Y_t$, where

$$dY_t = \Delta Y_t dt + \iota dt + \sigma dW_t.$$

Example 5.9. Consider first the subsection 5.1 formulation in which the continuation value does not depend on ι . Guess that

$$V_t = \lambda \cdot Y_t + \kappa \cdot \bar{v}_t + \phi_t.$$

Then

$$\begin{aligned} \bar{\eta}_t &= (\Delta Y_t + \bar{v}_t) \cdot \lambda + \frac{d\phi_t}{dt} \\ \varsigma_t &= \sigma' \lambda + \sigma^{-1} \Sigma_t \kappa. \end{aligned}$$

Moreover, since the posterior for ι is normal, it follows from (8) and (9) that

$$\bar{\eta}_t = \delta V_t - \delta H Y_t + \frac{|\varsigma_t|^2}{2\theta_1} + \frac{1}{2\theta_2} (\varsigma_t)' \sigma^{-1} \Sigma_t (\sigma^{-1})' \varsigma_t.$$

Combining these findings, we know

$$\begin{array}{rclcl} \Delta' \lambda & = & \delta \lambda & - & \delta H \\ \lambda & = & \delta \kappa & & \\ \frac{d\phi_t}{dt} & = & \delta \phi_t & + & \frac{|\varsigma_t|^2}{2\theta_1} + \frac{1}{2\theta_2} (\varsigma_t)' \sigma^{-1} \Sigma_t (\sigma^{-1})' \varsigma_t \\ & & & & \uparrow \qquad \qquad \qquad \uparrow \\ & & & & \text{misspecified} \qquad \qquad \text{misspecified} \\ & & & & \text{dynamics} \qquad \qquad \text{state estimation.} \end{array}$$

The first equation can be solved for λ and then the second for κ . To compute ϕ_t , we solve a first-order differential equation forward:

$$\phi_t = - \int_0^\infty \exp(-\delta u) \left[\frac{|\varsigma_{t+u}|^2}{2\theta_1} + \frac{1}{2\theta_2} (\varsigma_{t+u})' \sigma^{-1} \Sigma_{t+u} (\sigma^{-1})' \varsigma_{t+u} \right] du$$

The fact that ϕ_t is negative reflects an aversion to uncertainty (a.k.a. a concern

about misspecified estimation and models). Welfare implications of this aversion can be quantified using the methods of [2] and [6].

Example 5.10. Now consider the subsection 5.2 formulation in which the continuation value depends on the model ι . Guess

$$V_t(\iota) = \lambda \cdot Y_t + \kappa \cdot \iota + \phi$$

where

$$\begin{aligned} \Delta' \lambda &= \delta \lambda - \delta H \\ \lambda &= \delta \kappa \\ \delta \phi &= -\frac{1}{2\theta_1} \lambda' \sigma \sigma' \lambda. \end{aligned}$$

Then

$$\hat{V}_t = \lambda \cdot Y_t + \kappa \cdot \bar{\iota}_t \quad -\frac{1}{2\delta\theta_1} \lambda' \sigma \sigma' \lambda \quad -\frac{1}{2\theta_2} \kappa' \Sigma_t \kappa.$$

\uparrow
misspecified
dynamics

\uparrow
misspecified
state estimation

Here, our robust-adjusted continuation value includes two negative terms, one that adjusts for model misspecification and another that adjusts for estimation based on a possibly misspecified model.

6 Hidden-state Markov models

This section extends the previous Markov setup by letting ι itself be governed by Markov transitions. This motivates the decision maker to learn about a moving target ι_t . Bayesian learning carries an asymptotic rate of learning linked in interesting ways to the tail behavior of detection error probabilities. For expositional convenience, we use special Markov settings that imply quasi-analytical formulas for filtering.

6.1 Kalman filtering

Consider a linear model in which some time-varying Markov states are hidden and in which the Markov state follows the linear law of motion

$$\begin{aligned}dX_t &= AX_t dt + BdW_t \\dY_t &= DX_t dt + Fdt + GdW_t\end{aligned}$$

where dY_t is observed. The random vector $DX_t + F$ plays the role of $\mu(Y_t) + \iota_t$ in the previous section and is partially hidden from the decision-maker. The Kalman filter provides a recursive solution to the filtering problem. We abstract from one aspect of time variation by letting the prior covariance matrix for the state vector Σ_0 equal its limiting value. Recursive filtering implies the innovations representation for Y_t with the conditional mean \bar{X}_t of X_t obeying

$$\begin{aligned}d\bar{X}_t &= A\bar{X}_t dt + \bar{B}d\bar{W}_t \\dY_t &= \mu_t dt + \sigma d\bar{W}_t,\end{aligned}\tag{10}$$

where σ is nonsingular, $\mu_t = D\bar{X}_t + F$ is the drift for the signal increment dY_t ,

$$\begin{aligned}\sigma\sigma' &= GG' \\ \bar{B} &= (\Sigma D' + BG')(GG')^{-1}\sigma = (\Sigma D' + BG')(\sigma')^{-1} \\ d\bar{W}_t &= \sigma^{-1} (dY_t - D\bar{X}_t - Fdt) = \sigma^{-1} [GdW_t + D(X_t - \bar{X}_t)dt].\end{aligned}$$

The matrix Σ is the limiting covariance matrix, which we assume exists and is nonsingular. Log consumption is $\log C_t = H \cdot Y_t$ and the benchmark preferences are discounted expected logarithmic utility, as in section 5.3.

6.1.1 Continuation values that do not depend on hidden states

Consider the case in which continuation values do not depend on the hidden states, as in section 5.1.

Guess a solution

$$V_t = \lambda \cdot \bar{X}_t + H \cdot Y_t + \phi.$$

The innovation for the continuation value is

$$\begin{aligned} (\lambda' \bar{B} + H' \sigma) d\bar{W}_t &= (\lambda' \bar{B} + H' \sigma) \sigma^{-1} [GdW_t + D(X_t - \bar{X}_t)dt], \\ &= [\lambda'(\Sigma D' + BG')(GG')^{-1} + H'] [GdW_t + D(X_t - \bar{X}_t)dt] \end{aligned}$$

and the drift for the continuation value under our guess is

$$\bar{\eta}_t = \lambda \cdot (A\bar{X}_t) + H \cdot (D\bar{X}_t) + H \cdot F. \quad (11)$$

From the specification of preferences, the drift η satisfies

$$\begin{aligned} \bar{\eta}_t &= \delta V_t - \delta \log C_t + \frac{1}{2\theta_1} [\lambda'(\Sigma D' + BG')(GG')^{-1} + H'] GG' [(GG')^{-1}(D\Sigma + GB')\lambda + H] \\ &\quad + \frac{1}{2\theta_2} [\lambda'(\Sigma D' + BG')(GG')^{-1} + H'] D\Sigma D' [(GG')^{-1}(D\Sigma + GB')\lambda + H]. \end{aligned} \quad (12)$$

Equating coefficients on \bar{X}_t as given from (11) and (12) gives

$$A'\lambda + D'H = \delta\lambda$$

Therefore, $\lambda = (\delta I - A')^{-1} D'H$.

As inputs into constructing detection-error probabilities, we require the worst-case distortions. The worst-case drift for dW_t is

$$\tilde{h}_t = -\frac{1}{\theta_1} G' [(GG')^{-1}(D\Sigma + GB')\lambda + H]$$

and the distorted mean \tilde{X}_t for X_t is

$$\tilde{X}_t = \bar{X}_t - \frac{1}{\theta_2} \Sigma D' [(GG')^{-1}(D\Sigma + GB')\lambda + H].$$

Combining these distortions gives

$$d\bar{W}_t = -(\bar{G})^{-1} \left(\frac{1}{\theta_1} GG' + \frac{1}{\theta_2} D\Sigma D' \right) [(GG')^{-1}(D\Sigma + GB')\lambda + H] + d\tilde{W}_t, \quad (13)$$

where $d\tilde{W}_t$ is a multivariate standard Brownian motion under the distorted probability

law. Both drift distortions are constant. Substituting (13) into (10) gives the implied distorted law of motion for the reduced information structure generated by the signal history. In particular, the distorted drift $\tilde{\mu}_t$ for dY_t is:

$$\tilde{\mu}_t = D\bar{X}_t - G \frac{1}{\theta_1} [(D\Sigma + GB')\lambda + (GG')H] - \frac{1}{\theta_2} D\Sigma D' [(GG')^{-1}(D\Sigma + GB')\lambda + H].$$

6.1.2 Continuation values that depend on the hidden states

As in section 5.2, we now turn to the case in which continuation values depend on the hidden states. Guess a solution

$$V_t = \lambda \cdot X_t + H \cdot Y_t + \phi.$$

The innovation to the continuation value is

$$(B'\lambda + G'H) \cdot dW_t,$$

and the drift is

$$\eta_t = \lambda \cdot (AX_t) + H \cdot (DX_t + F).$$

This drift satisfies

$$\eta_t = \delta V_t - \delta H \cdot X_t + \frac{1}{2\theta_1} |B'\lambda + G'H|^2.$$

The vector λ is the same as in the limited information case, and the worst-case model prescribes the following drift to the Brownian increment dW_t :

$$\tilde{h}_t = -\frac{1}{\theta_1} (B'\lambda + G'H).$$

The robust state estimate \tilde{X}_t is

$$\tilde{X}_t = \bar{X}_t - \frac{1}{\theta_2} \Sigma \lambda.$$

Thus, the combined distorted drift for dY_t is

$$\tilde{\mu}_t = D\bar{X}_t - \frac{1}{\theta_1} G(B'\lambda + G'H) - \frac{1}{\theta_2} D\Sigma \lambda.$$

The drift distortions in both robustness specifications are constant. While the pa-

parameterization of a hidden-state Markov model leading to this outcome is convenient, the constant distortions make it empirically limiting when it comes to accounting for history-dependent market prices of uncertainty.¹² Partly for that reason, we consider next environments that imply state dependent distortions in the probabilities that incidentally can lead to time-varying and state-dependent market prices of uncertainty.

6.2 Wonham filtering

Suppose that

$$dY_t = \Delta Y_t dt + \iota_t dt + \sigma dW_t$$

where $\iota_t = \Gamma Z_t$ and Z_t follows a discrete-state Markov chain with intensity matrix A . The realized values of Z_t are coordinate vectors, i.e., a basis of vectors with 1 in one location, zeroes elsewhere. This is a Wonham filtering problem where the signal increment is now $dY_t - \Delta Y_t dt$. The local population regression of Z_t onto $dY_t - \Delta Y_t dt - \Gamma \bar{Z}_t dt$ has the conditional regression coefficient vector

$$K_t = \left[\text{diag}(\bar{Z}_t) - \bar{Z}_t \bar{Z}_t' \right] G' (\sigma \sigma')^{-1}.$$

The notation “diag” denotes a diagonal matrix with the entries of the argument in the respective diagonal positions. Then the recursive solution to the filtering problem is

$$\begin{aligned} dY_t &= \Delta Y dt + \Gamma \bar{Z}_t dt + \sigma d\bar{W}_t \\ d\bar{Z}_t &= A' \bar{Z}_t dt + K_t \sigma d\bar{W}_t, \end{aligned}$$

where the innovation

$$d\bar{W}_t = \sigma^{-1} \Gamma (Z_t - \bar{Z}_t) dt + dW_t$$

is an increment to a multivariate standard Brownian motion.

¹²This observation motivated [18] and [14] to explore alternative specifications in their empirical applications that give rise to history dependent contributions of model uncertainty to risk prices. We could induce time dependence by initializing the state covariance matrix away from its invariant limit, but this induces only a smooth alteration in uncertainty prices that does not depend on the data realizations.

6.2.1 Continuation values

When continuation values do not depend on hidden states, the continuation value function must be computed numerically. The full information value function, however, is of the form

$$V_t = \lambda \cdot Y_t + \kappa \cdot Z_t + \xi.$$

Appendix A derives formulas for κ and λ when the logarithm of consumption is given by $H \cdot X_t$ and the instantaneous utility function is logarithmic (a unitary intertemporal elasticity of substitution). Given solutions for λ and κ , appendix A also provides formulas for the worst-case drift distortion \tilde{h} and the worst-case intensity matrix \tilde{A} for the continuous-time Markov chain.

6.2.2 Worst-case state estimate

It remains to compute the worst-case state estimate. For this, we solve

$$\begin{aligned} \min_{\{\tilde{Z}_{i,t}\}_i} \sum_i \tilde{Z}_{i,t} \left[\kappa_i + \theta_2 \left(\log \tilde{Z}_{i,t} - \log \bar{Z}_{i,t} \right) \right] \\ \text{subject to } \sum_i \tilde{Z}_{i,t} = 1. \end{aligned}$$

The minimizer

$$\tilde{Z}_{i,t} \propto \bar{Z}_{i,t} \exp \left(-\frac{\kappa_i}{\theta_2} \right)$$

“tilts” \tilde{Z}_t towards states with smaller continuation values.

6.2.3 Combined distortion

Given the initial robust state estimate and the worst-case dynamics, we again apply the Wonham filter to obtain:

$$\begin{aligned} dY_t &= \Delta Y_t dt + \Gamma \tilde{Z}_t dt + \sigma \tilde{h} dt + \sigma d\tilde{W}_t \\ d\tilde{Z}_t &= \tilde{A}' \tilde{Z}_t dt + \tilde{K}_t (dY_t - \Delta Y_t dt - \Gamma \tilde{Z}_t dt - \sigma \tilde{h} dt) \end{aligned}$$

where

$$\tilde{K}_t = \left[\text{diag} \left(\tilde{Z}_t \right) - \tilde{Z}_t \tilde{Z}_t' \right] \Gamma' (\sigma \sigma')^{-1}.$$

7 Statistical discrimination

This section uses insights from the statistical detection literature to defend the rescaling of entropy contributions recommended in sections 5, 6, and 8. [8] used likelihood ratios to discriminate among competing statistical models. [29] extended [8]’s analysis to apply to continuous-time Markov processes with observable states. We follow [1], [14], and [17], ch. 9 in using such methods to quantify how difficult it is to distinguish worst-case models from the decision maker’s benchmark models. We extend the analysis of [1] to allow for hidden Markov states.

For the hidden-state Markov models of section 6, the log-likelihood ratio between the worst-case and benchmark models evolves as

$$\begin{aligned} d \log \tilde{L}_t - d \log \bar{L}_t &= (\tilde{\mu}_t - \bar{\mu}_t)' (\sigma \sigma')^{-1} dY_t \\ &\quad - \frac{1}{2} (\tilde{\mu}_t)' (\sigma \sigma')^{-1} \tilde{\mu}_t dt \\ &\quad + \frac{1}{2} (\bar{\mu}_t)' (\sigma \sigma')^{-1} \bar{\mu}_t dt, \end{aligned}$$

where the specification of $\bar{\mu}_t$ and $\tilde{\mu}_t$ depends on the specific hidden state. Equivalently, the evolution of the likelihood ratio can be written as

$$\begin{aligned} d \log \tilde{L}_t - d \log \bar{L}_t &= (\tilde{\mu}_t - \bar{\mu}_t)' (\sigma')^{-1} d\bar{W}_t \\ &\quad - \frac{1}{2} (\tilde{\mu}_t - \bar{\mu}_t) (\sigma \sigma')^{-1} (\tilde{\mu}_t - \bar{\mu}_t) dt, \end{aligned}$$

which makes the likelihood ratio a martingale with respect to the reduced information filtration \mathcal{Y} generated the process Y under the benchmark model probabilities. The alternative specifications in section 6 imply different formulas for the conditional means of the hidden states and worst-case adjustments, but the continuous-time likelihood ratio for each of them shares this common structure.

[8] used the expected value of the likelihood ratio to a power $0 < \alpha < 1$ to bound the limiting behavior of the probability that the likelihood ratio exceeds alternative thresholds.¹³ In our setting, this approach leads us to study the behavior of the

¹³Chernoff’s calculation was an early application of Large Deviation Theory.

conditional expectation of $M_t(\alpha) = \left(\tilde{L}_t/L_t\right)^\alpha$. The logarithm of $M_t(\alpha)$ evolves as:

$$\begin{aligned} d \log M_t(\alpha) &= \alpha(\tilde{\mu}_t - \bar{\mu}_t)'(\sigma')^{-1}d\bar{W}_t \\ &\quad - \frac{\alpha}{2}(\tilde{\mu}_t - \bar{\mu}_t)(\sigma\sigma')^{-1}(\tilde{\mu}_t - \bar{\mu}_t) \\ &= \alpha(\tilde{\mu}_t - \bar{\mu}_t)'(\sigma')^{-1}d\bar{W}_t - \frac{\alpha^2}{2}(\tilde{\mu}_t - \bar{\mu}_t)(\sigma\sigma')^{-1}(\tilde{\mu}_t - \bar{\mu}_t) \\ &\quad + \left(\frac{\alpha^2 - \alpha}{2}\right)(\tilde{\mu}_t - \bar{\mu}_t)(\sigma\sigma')^{-1}(\tilde{\mu}_t - \bar{\mu}_t). \end{aligned}$$

This shows that $M_t(\alpha)$ can be factored into two components identified by the last two lines. The first component $M_t^1(\alpha)$ evolves as

$$d \log M_t^1(\alpha) = \alpha(\tilde{\mu}_t - \bar{\mu}_t)'(\sigma')^{-1}d\bar{W}_t - \frac{\alpha^2}{2}(\tilde{\mu}_t - \bar{\mu}_t)(\sigma\sigma')^{-1}(\tilde{\mu}_t - \bar{\mu}_t)$$

and is a local martingale in levels. The second component $M_t^2(\alpha)$ evolves as

$$d \log M_t^2(\alpha) = \left(\frac{\alpha^2 - \alpha}{2}\right)(\tilde{\mu}_t - \bar{\mu}_t)(\sigma\sigma')^{-1}(\tilde{\mu}_t - \bar{\mu}_t)dt.$$

Since it has no instantaneous exposure to $d\bar{W}_t$, this second component is locally predictable. An additive decomposition in logarithms implies a multiplicative decomposition in levels. Since the conditional expectation of the martingale component does not grow, the local growth rate is fully encoded in the locally predictable component.

7.1 Local discrimination

[1] used the $d \log M_t^2(\alpha)$ component for a fully observed Markov process to define the local rate of statistical discrimination. This local rate gives a statistical measure of how easy it is to discriminate among competing models using historical data, in the sense that it bounds the rate at which the probability of making a mistake in choosing between two models decreases as the sample size grows. The counterpart of this local rate for a hidden state model is

$$\frac{1}{8}(\tilde{\mu}_t - \bar{\mu}_t)(\sigma\sigma')^{-1}(\tilde{\mu}_t - \bar{\mu}_t) \tag{14}$$

This rate attains

$$\max_{\alpha} \left(\frac{\alpha - \alpha^2}{2} \right) (\tilde{\mu}_t - \bar{\mu}_t) (\sigma \sigma')^{-1} (\tilde{\mu}_t - \bar{\mu}_t)$$

where the objective is maximized by setting $\alpha = 1/2$. Since [1] consider complete information models, they focus only on the drift distortion to an underlying Brownian motion. But for us, states are hidden from the decision-maker, so robust estimation necessarily gives rise to an additional contribution to the local rate of statistical discrimination between models.¹⁴

For the stochastic specification with an invariant ι and continuation values that do not depend on the unknown model,

$$\begin{aligned} \tilde{\mu}_t - \bar{\mu}_t = & \quad -\frac{1}{\theta_1} \sigma(Y_t) \varsigma_t \quad + \quad [\sum_{\iota} \tilde{\pi}_t(\iota) \mu(Y_t, \iota)] - \bar{\mu}(Y_t) \\ & \quad \uparrow \qquad \qquad \qquad \uparrow \\ & \text{misspecified} \qquad \qquad \text{misspecified} \\ & \text{dynamics} \qquad \qquad \text{model estimation} \end{aligned}$$

where $\tilde{\pi}_t$ solves problem 5.3. This formula displays the two contributions to statistical discrimination and via the second term on the right shows how a concern about misspecified model estimation alters the local rate of statistical discrimination. Recall that in our recursive formulation, we scaled the contribution of entropy from the posterior over the unknown model by dt . Evidently this scaling balances the contributions to the detection error rate in a way designed to make both components be of comparable magnitudes. The impact of the misspecified model estimation will vanish over time as the decision maker learns ι , however.

Consider next the Kalman filtering model with continuation values that do not depend on the hidden state. In this case, we showed that

$$\begin{aligned} \tilde{\mu}_t - \bar{\mu}_t = & \quad -\frac{1}{\theta_1} G(B'\lambda + G'H) \qquad \qquad -\frac{1}{\theta_2} D\Sigma\lambda. \\ & \quad \uparrow \qquad \qquad \qquad \uparrow \\ & \text{misspecified} \qquad \qquad \text{misspecified} \\ & \text{dynamics} \qquad \qquad \text{state estimation} \end{aligned}$$

¹⁴The term $\sigma^{-1}(\tilde{\mu}_t - \bar{\mu}_t)$ entering (14) is also the uncertainty component to the price of local exposure to the vector $d\bar{W}_t$. See [1] and [18].

Now the second term on the right shows how a concern about misspecified state estimation alters the local rate of statistical discrimination. Since the hidden state evolves over time, the impact of the second term will not vanish. Both contributions, however, are time invariant.

Finally, consider the Wonham filtering model with continuation values that do not depend on the discrete hidden state. Now both distortions must be computed numerically. They depend on the vector of vector Y_t of observables and on probabilities over the hidden states, persist through time, and have comparable magnitudes.

An analogous set of results can be obtained when continuation values depend on the unknown model or hidden states.

7.2 Long-run discrimination

For the Kalman filtering model, the local discrimination rate is constant and necessarily coincides with its long-term counterpart. For the Wonham filtering model with discrete hidden states, the local discrimination rate is state dependent. However, it has a limiting discrimination rate that is state independent. [29] construct this rate for a fully observed Markov state. The filtering problem associated with a hidden state Markov process implies an alternative Markov process in which the hidden state is replaced by the density of the hidden state conditional on the history of signals. The recursive representation of the motion of that posterior distribution shows how the posterior responds to new information. We apply the approach suggested by [29] to the Markov process constructed from the recursive solution to the filtering problem.

We know the likelihood ratio and its evolution under the benchmark model. In particular, notice that we can construct \tilde{Z}_t as a function of \bar{Z}_t given the vector κ . The long-run rate maximizes the limit

$$\rho(\alpha) = \lim_{t \rightarrow \infty} \frac{1}{t} \log E [M_t(\alpha) | Y_0 = y, \bar{Z}_0 = z],$$

by choice of α . The rate $\rho(\alpha)$ can also be characterized by finding the dominant eigenvalue for the generator of a semigroup of operators. Operator t in this semigroup maps a function of (y, z) into another function of (y, z) defined by

$$E \left[\frac{M_t(\alpha)}{M_0(\alpha)} \varphi(Y_t, \bar{Z}_t) | Y_0 = y, \bar{Z}_0 = z \right].$$

The pertinent eigenvalue problem is

$$E \left[\frac{M_t(\alpha)}{M_0(\alpha)} \varphi(Y_t, \bar{Z}_t) | Y_0 = y, \bar{Z}_0 = z \right] = \exp[-\rho(\alpha)t] \varphi(y, z)$$

where φ is restricted to be a positive function.¹⁵ Since this equation must hold for all t , there is a local counterpart that requires solving a second-order partial differential equation. Again, concerns about misspecified dynamics and misspecified state estimation both contribute to the asymptotic rate. Following [29], the long-term counterpart to Chernoff entropy is:

$$\max_{0 \leq \alpha \leq 1} \rho(\alpha).$$

To relate this to our discussion of local discrimination rates, for a given α , suppose that M_t^α is a martingale (not just a local martingale). This multiplicative martingale is associated with a change in probability measure, one that preserves the Markov structure. With the change of measure, we may use the approach pioneered by Donsker and Varadhan to characterize $\rho(\alpha)$ as the solution to a problem that maximizes the average local discrimination rate subject to a relative entropy constraint for the probability measures used in computing the average.¹⁶ The benchmark probability distribution used in forming the relative entropy criterion is that of the stationary distribution for the Markov process under the change of measure, presuming that such a distribution exists. Given this averaging property, the comparability of magnitudes that we described for the local discrimination rates carries over to the global rates.

8 Smooth adjustments for model uncertainty

Thus far, we have described links between entropy-based robust control problems and recursive utility in continuous time. In this section, we explore an analogous link of robust control theory to the smooth ambiguity decision theoretic model of [24, 25]. [16] mentioned this link for the discrete-time model counterpart to the decision problem described above in 5.1. In particular, instead of directly introducing perturbations in probabilities as we do in robust control theory, we can use an exponential adjustment

¹⁵There may be multiple solutions to this eigenvalue problem but there are well known ways to select the appropriate solution. See [20].

¹⁶See, for instance, [10].

for model ambiguity and view it as a special case of [25] parameterized by θ_2 . That is, we impose the exponential adjustment *directly* without viewing it as the outcome of a multiplier problem that involves θ_2 -penalized minimization over perturbations to probabilities. To isolate the link to decision-theoretic models with smooth ambiguity, we set θ_1 to infinity. Here we explore the link between the decision problem in section 5.1 and a version of the problem studied by [25] appropriately modified to match our continuous-time specification. For simplicity, we feature the case in which learning is about an invariant parameter or model indicator, ι .

Consider a sequence of stochastic environments indexed by a parameter ϵ that sets the time gap between observations. It is perhaps simplest to think of $\epsilon = 2^{-j}$ for nonnegative integers j and for a fixed j construct a stochastic process of observations at dates $0, \epsilon, 2\epsilon, \dots$. Incrementing from j to $j + 1$ divides the sampling interval in half. To simplify notation and approximation, we use the continuous-time formulation structures set out in subsection 5.1 with continuation values not depending on the unknown parameter, but we consider consumption choices made on the ϵ -spaced grid just described.¹⁷

Given knowledge of ι , the decision maker uses expected utility preferences:

$$\begin{aligned} V_t(\iota) &= [1 - \exp(-\epsilon\delta)] E \left[\sum_{k=0}^{\infty} \exp(-k\epsilon\delta) U[C_{t+k\epsilon}] | \mathcal{Y}_t, \iota \right] \\ &= [1 - \exp(-\epsilon\delta)] U(C_t) + \exp(-\epsilon\delta) E [V_{t+\epsilon}(\iota) | \mathcal{Y}_t, \iota]. \end{aligned}$$

Thus,

$$\begin{aligned} E [V_{t+\epsilon}(\iota) | \mathcal{Y}_t, \iota] - V_t(\iota) &= [\exp(\epsilon\delta) - 1] V_t(\iota) - [\exp(\epsilon\delta) - 1] U(C_t) \\ &\approx \epsilon\delta V_t - \epsilon\delta U(C_t). \end{aligned} \tag{15}$$

To extend these preferences to accommodate smooth ambiguity as in [25], we first consider preferences over one step-ahead continuation plans. The one step-ahead construction begins with a consumption process at date t : $(C_t, C_{t+\epsilon}, C_{t+2\epsilon}, \dots)$ and forms a new stochastic sequence: $(C_t, C_{t+\epsilon}, C_{t+\epsilon}, \dots)$ that is constant over time from

¹⁷An alternative approach would be to sample $\{Y_t\}$ for each choice of ϵ and solve the corresponding filtering problem in discrete-time for each ϵ .

time period $t + \epsilon$ forward. Therefore, the date $t + \epsilon$ continuation value $V_{t+\epsilon}$ is

$$V_{t+\epsilon} = U(C_{t+\epsilon}).$$

Because no C 's beyond $t + \epsilon$ affect $V_{t+\epsilon}$, knowledge of ι plays no role in the valuation from date $t + \epsilon$ forward. For a given ι , [25] define the one-step ahead certainty equivalent $\check{C}_t(\iota)$ to be the solution to

$$[1 - \exp(-\epsilon\delta)]U(C_t) + \exp(-\epsilon\delta)U[\check{C}_t(\iota)] = [1 - \exp(-\epsilon\delta)]U(C_t) + \exp(-\epsilon\delta)E[V_{t+\epsilon}|\mathcal{Y}_t, \iota].$$

Therefore,

$$\check{C}_t(\iota) = U^{-1}(E[V_{t+\epsilon}|\mathcal{Y}_t, \iota]).$$

For future reference, define

$$\check{V}_t(\iota) = U[\check{C}_t(\iota)] = E[V_{t+\epsilon}|\mathcal{Y}_t, \iota]. \quad (16)$$

[25] refer to $\check{C}_t(\iota)$ as the second-order act associated with $(C_t, C_{t+\epsilon}, C_{t+2\epsilon} \dots)$, and they impose an assumption of subjective expected utility over second-order acts, where the probability measure is the date t posterior for ι . See their Assumption 7.

To form a bridge between the formulation of [25] and our work, we would let the function that they use to represent expected utility over second-order acts have the exponential form¹⁸

$$U^*(C) = -\exp[-\gamma U(C)]. \quad (17)$$

Then the date t objective applied to second-order acts is the expected utility

$$-E(\exp[-\gamma \check{V}_t(\iota)]|\mathcal{Y}_t).$$

The continuation value certainty equivalent of \check{V}_t is

$$-\frac{1}{\gamma} \log E(\exp[-\gamma \check{V}_t(\iota)]|\mathcal{Y}_t).$$

¹⁸Note that [24] allow for more general utility functions U^* .

Assembling these components gives the recursive representation

$$\begin{aligned} V_t &= [1 - \exp(-\epsilon\delta)]U(C_t) - \exp(-\epsilon\delta)\frac{1}{\gamma} \log E [\exp [-\gamma\check{V}_t(\iota)] | \mathcal{Y}_t] \\ &= [1 - \exp(-\epsilon\delta)]U(C_t) - \exp(-\epsilon\delta)\frac{1}{\gamma} \log E [\exp (-\gamma E [V_{t+\epsilon} | \mathcal{Y}_t, \iota]) | \mathcal{Y}_t], \end{aligned}$$

where we have substituted for $\check{V}_t(\iota)$ from (16). [25] use dynamic consistency to extend this preference representation beyond second-order acts.

8.1 A smooth ambiguity adjustment that vanishes in continuous time

Consider now a continuous-time approximation. Take a continuation value process $\{V_t\}$ with drift $\eta_t(\iota)$ conditioned on ι . For example, take $V_\tau = U(C_{t+\epsilon})$ $\tau \geq t + \epsilon$ as in the construction of second-order acts. Using a continuous-time approximation,

$$\check{V}_t(\iota) = E [V_{t+\epsilon} | \mathcal{Y}_t, \iota] \approx V_t + \epsilon\eta_t(\iota)$$

Then

$$-\frac{1}{\gamma} \log E (\exp [-\gamma\check{V}_t(\iota)] | \mathcal{Y}_t) \approx V_t + \epsilon\bar{\eta}_t - \frac{1}{\gamma} \log E \left(\exp [-\gamma\epsilon[\eta_t(\iota) - \bar{\eta}_t]] | \mathcal{Y}_t \right) \quad (18)$$

where

$$\bar{\eta}_t = E [\eta_t(\iota) | \mathcal{Y}_t]$$

is the date t drift for the process $\{U(C_\tau) : \tau \geq t\}$ under the filtration $\{\mathcal{Y}_\tau : \tau \geq t\}$ that omits knowledge of ι . Since $[\eta_t(\iota) - \bar{\eta}_t]$ has conditional mean zero,

$$\frac{1}{\gamma} \log E (\exp [-\gamma\epsilon[\eta_t(\iota) - \bar{\eta}_t]] | \mathcal{Y}_t)$$

contributes only an ϵ^2 term. This can be seen by using the power series expansion for the exponential and taking conditional expectations of the terms. The first-order term in ϵ is zero because $[\eta_t(\iota) - \bar{\eta}_t]$ has conditional mean zero, so the second-order term dominates for small ϵ . The limiting counterpart to (15) scaled by $\frac{1}{\epsilon}$ is

$$\bar{\eta}_t = \delta V_t - \delta U(C_t).$$

The parameter γ drops out of this equation in the limit and there is no adjustment for ambiguity. This calculation reaffirms a finding of [32] that a smooth ambiguity adjustment completely dissipates in a continuous-time limit.

8.2 A smooth (in a derivative) ambiguity adjustment that survives in continuous time

As an alternative, suppose that we adjust the utility function U^* over second-order acts simultaneously with ϵ . In particular, we replace γ with $\frac{\gamma}{\epsilon}$ on the right-hand side of (18):

$$U(C_t) + \epsilon \bar{\eta}_t - \frac{\epsilon}{\gamma} \log E [\exp (-\gamma [\eta_t(\iota) - \bar{\eta}_t]) | \mathcal{Y}_t].$$

This leads to

$$\bar{\eta}_t - \frac{1}{\gamma} \log E [\exp (-\gamma [\eta_t(\iota) - \bar{\eta}_t]) | \mathcal{Y}_t] = \delta V_t - \delta C_t,$$

or

$$\bar{\eta}_t = \delta V_t - \delta C_t + \frac{1}{\gamma} \log E [\exp (-\gamma [\eta_t(\iota) - \bar{\eta}_t]) | \mathcal{Y}_t]. \quad (19)$$

Our alternative adjustment thus makes the concern about ambiguity remain in a continuous-time limit.

To connect to our subsection 5.1 analysis of continuous-time versions of preferences for robustness, write the local evolution for the continuation value under the filtration \mathcal{Y} as:

$$dV_t = \bar{\eta}_t dt + \varsigma_t \cdot d\bar{W}_t$$

and recall from equation (7) that¹⁹

$$d\bar{W}_t = dW_t + \sigma(Y_t)^{-1} [\mu(Y_t, \iota) - \bar{\mu}(Y_t)] dt.$$

Thus,

$$\eta_t(\iota) = \bar{\eta}_t + (\varsigma_t)' \sigma(Y_t)^{-1} [\mu(Y_t, \iota) - \bar{\mu}(Y_t)] dt,$$

¹⁹Here we see the important role of decomposition (7) alluded to in footnote 9.

and the drift for our version of smooth ambiguity can be expressed as

$$\begin{aligned} \bar{\eta}_t = & \delta V_t - \delta U(C_t) + (\varsigma_t)' \sigma(Y_t)^{-1} \bar{\mu}(Y_t) \\ & + \frac{1}{\gamma} \log \left(\sum_{\iota} \pi_t(\iota) \exp \left[-\gamma (\varsigma_t)' \sigma(Y_t)^{-1} \mu(Y_t, \iota) \right] \right). \end{aligned}$$

This formulation coincides with (8) when $\theta_1 = \infty$ and $\theta_2 = \frac{1}{\gamma}$.

8.2.1 Senses of smoothness

As we showed in (19), the continuous-time limit of our scaling makes a smooth exponential ambiguity adjustment to the *derivative* of the continuation value. The way that this limit U^* in (18) depends on the sampling interval ϵ leads us to raise the important question of the extent to which we can hope to parameterize ambiguity preferences through U^* in a way that plausibly remains fixed across alternative environments. When we follow [25] and use U^* in (18) and expected utility over second-order acts, the ambiguity adjustment vanishes in the continuous time limit computed in subsection 8.1. The reason that the adjustment disappears in the limit is that the impact of uncertainty about ι on transition distributions for the Markov state vanishes too quickly as we shrink the sampling interval ϵ to zero. In order to sustain an ambiguity adjustment in continuous time, we have increased the curvature of U^* as we have diminished the sampling interval. In terms of the [25] analysis, we performed this adjustment because of how uncertainty about ι is manifested in the constructed second-order acts.

The robustness model counterpart to the smooth ambiguity model is our subsection 5.1 formulation of preferences in which continuation values do not depend on hidden states or parameters. Our study of the properties of statistical detection has led us directly to suggest an alternative scaling in which ambiguity aversion becomes reflected in the local evolution of the continuation value. The statistical detection link under our proposed rescaling is an important outcome for us because of how it conveniently links to the strategy that [1] proposed for quantifying model ambiguity by calibrating θ .

8.2.2 Example

In order to illustrate the impact of our proposed adjustment, return to the first example in section 5.3. We showed that the value function has the form

$$V_t = \lambda \cdot Y_t + \kappa \cdot \bar{l}_t + \phi_t$$

and reported formulas for λ , κ and ϕ_t . Under the ambiguity interpretation

$$\phi_t = -\gamma \int_0^\infty \exp(-\delta u) [\zeta_{t+u}'(\sigma^{-1})\Sigma_{t+u}(\sigma^{-1})'\zeta_{t+u}] du$$

scales linearly in the ambiguity parameter γ .

What lessons do we learn from this? The exposure of a continuation value to model uncertainty diminishes proportionally to ϵ as ϵ shrinks to zero. But the risks conditioned on a model ι , namely, $W_{t+\epsilon} - W_t$, have standard deviations that scale as $\sqrt{\epsilon}$, and these risks come to dominate the uncertainty component. By replacing the ambiguity parameter γ with $\frac{\gamma}{\epsilon}$, we offset this diminishing importance of ambiguity when we move toward approximating a continuous-time specification.

The consequences of our rescaling proposal for preferences over consumption processes are apparently reasonable. This is because even though we drive the ambiguity aversion parameter $\frac{\gamma}{\epsilon}$ to infinity in the continuous-time limit, the local uncertainty exposure of continuation values and consumption diminish simultaneously at comparable rates. This calculation provokes further thoughts about how separately to calibrate both a decision maker's ambiguity aversion and his risk aversion. The empirical analyses of [23] and [9] study implications of smooth ambiguity models. Our continuous-time limiting investigations motivate the discrete-time ambiguity aversion parameter γ in terms of the local uncertainty that confronts a decision maker.

Our example sets an intertemporal substitution elasticity equal to unity with the consequence that a proportionate change in the consumption process leads to a change in the constant term of the continuation value equal to the logarithm of the proportionality factor. If we modify the intertemporal substitution elasticity parameter using the [25] formulation in conjunction with our exponential risk adjustment, this homogeneity property no longer applies. In discrete-time, [23] and [21] propose an alternative recursion with ambiguity that preserves the homogeneity property just described.

9 Concluding remarks

All versions of max-min expected utility models, including the recursive specifications of robustness developed in [16], assign a special role to a worst-case model. We advocate a strategy for quantifying model ambiguity in particular applications by imagining that a decision maker thinks about the statistical plausibility of that model worst-case model. How he thinks about statistically discriminating it from other models should affect his robust decision making procedure.

In our formulations, robust decision makers use an approximating model as a benchmark around which they put a cloud of perturbed models with respect to which they want robust valuations and decisions. To characterize how “robust” the decision maker wants to be, [1] and [17], ch. 19 employed measures of statistical discrepancy between worst-case models and benchmark models. This paper has pushed that use of detection error calculations further by employing them to justify a proposal for a new way of scaling contributions to entropy in continuous time hidden Markov models with robust decision making. The link to statistical detection that is preserved under our proposed way of scaling is helpful because it permits us to use the detection error probabilities advocated by [1] as a way of quantifying model ambiguity as captured by our penalty parameter θ .

We have explored robustness to two alternative types of misspecifications. The first type is misspecified dynamics as reflected in distributions of current and future states and signals conditioned on current states. The second type is misspecified dynamics of the histories of the signals and hidden state variables occurring in filtering problems. We investigate the impact of both forms of misspecification on decision making and to understand better the impact of relative entropy restraints on the model misspecifications that a decision maker chooses to explore. We considered parameterizations of preferences for robustness for which both types of misspecification are parts of the problem of designing robust decision rules in the continuous-time limit. We have advocated specific continuous-time formulations for robust decision problems in light of how both types of misspecification contribute to measures of statistical discrepancy.

By construction, model ambiguity survives in our limiting formulations. As a consequence, our continuous-time formulations respond to a challenge that [32] posed for recursive formulations of smooth ambiguity like one of [24, 25]. The smoothness

in [24, 25] refers to averaging using probabilities as weights. Expressing ambiguity in this way avoids the kinks in indifference curves that are present when ambiguity is expressed using the max-min expected utility theory of [13]. [32] showed how a smooth concern for ambiguity vanishes in a particular continuous-time limit, but the [32]’s limit is distinct from ours. Indeed, our statistical detection approach tells us to explore a different continuous-time limit than the one criticized by [32]. We recognize that the sense in which our limiting preferences remain *smooth* is delicate. They *are* smooth, in the sense of taking a weighted average of a *derivative* (and not the *level*) of a continuation value using probabilities as weights.

A Discrete-state entropy problem

In this appendix, we follow [1] in computing the worst-case distortion of a discrete-state Markov chain. The continuation value is assumed to be given by

$$V_t = \lambda \cdot Y_t + \kappa \cdot Z_t + \xi.$$

Consider a matrix R with nonnegative entries. Each row has at least one strictly positive entry and the row gives transition probabilities conditioned that a jump occurs up to a proportionality factor. The proportionality factor gives the state dependent jump intensity. The implied intensity matrix A is

$$A = R - \text{diag} \{R\mathbf{1}_n\}$$

where $\mathbf{1}_n$ is an n -dimensional vector with all entries equal to one, and the matrix of transition probabilities is $\exp(tA)$ over an interval of time t . Consider an alternative specification of R given by $S \otimes R$, where S has all positive entries and \otimes denotes entry-by-entry multiplication.

The combined conditional relative entropy for the drift distortion for the Brownian motion and the distortion to the intensity matrix is

$$\text{ent}(h, S) = \frac{|h|^2}{2} + z \cdot \text{vec} \left[\sum_j r_{ij} (1 - s_{ij} + s_{ij} \log s_{ij}) \right]$$

where “vec” denotes a vector formed by stacking the numbers in its argument. Notice that

$$1 - s_{ij} + s_{ij} \log s_{ij} \geq 0$$

since $s_{ij} \log s_{ij}$ is convex and lies above its gradient approximation at $s_{ij} = 1$. Thus, as expected $\text{ent}(h, S)$ is nonnegative. The associated distorted drift for the continuation value inclusive of the entropy penalty is

$$\lambda \cdot (\Delta y + \Gamma z + \sigma h) + z \cdot [(R \otimes S)\kappa] - (z \cdot \kappa) (z \cdot [(R \otimes S)\mathbf{1}_n]) + \theta_1 \text{ent}(h, s).$$

To compute the worst-case model for the state dynamics, we minimize this expression by choice of the vector h and matrix S .

The worst-case model appends the drift

$$\tilde{h} = -\frac{1}{\theta_1} \sigma' \lambda$$

to the Brownian increment dW_t and includes a multiplicative distortion \tilde{S} to the matrix R :

$$\tilde{s}_{ij} = \exp\left(-\frac{1}{\theta_1} \kappa_j + \frac{1}{\theta_1} \kappa_i\right).$$

The minimized drift inclusive of the robustness penalty is

$$\lambda \cdot (\Delta y) - \frac{1}{2\theta_1} \lambda' \sigma \sigma' \lambda + \frac{z' \Delta \exp\left(-\frac{1}{\theta_1} \kappa\right)}{z' \exp\left(-\frac{1}{\theta_1} \kappa\right)},$$

where $\exp\left(-\frac{1}{\theta_1} \kappa\right)$ is a vector with entries given by exponentials of the entries in the vector argument. The drift of the value function must satisfy

$$\lambda \cdot (\Delta y + \Gamma z) - \frac{1}{2\theta_1} \lambda' \sigma \sigma' \lambda + \frac{z' A \exp\left(-\frac{1}{\theta_1} \kappa\right)}{z' \exp\left(-\frac{1}{\theta_1} \kappa\right)} = \delta(\lambda \cdot y + \kappa \cdot z + \xi) - \delta H \cdot y,$$

which gives equations to be solved for λ and κ .

References

- [1] Evan W. Anderson, Lars Peter Hansen, Thomas J. Sargent, A quartet of semi-groups for model specification, robustness, prices of risk, and model detection, *Journal of the European Economic Association* 1 (1) (2003) 68–123.
- [2] Francisco Barillas, Lars Peter Hansen, Thomas J. Sargent, Doubts or variability? *Journal of Economic Theory* 144 (6) (2009) 2388 – 2418.
- [3] Truman F. Bewley, Knightian decision theory: Part i, (1986) Cowles Foundation discussion paper 807.
- [4] Truman F. Bewley, Knightian decision theory, part ii: Intertemporal problems, (1987) Cowles Foundation discussion paper 835.
- [5] Truman F. Bewley, Knightian decision theory and econometric inference, (1988) Cowles Foundation discussion paper 868.
- [6] Simone Cerreia, Fabio Maccheroni, Massimo Marinacci, Luigi Montrucchio, Uncertainty averse preferences, (2008) working paper, Collegio Carlo Alberto.
- [7] Zengjing Chen, Larry Epstein, Ambiguity, risk, and asset returns in continuous time, *Econometrica* 70 (2002) 1403 – 1443.
- [8] H. Chernoff, A measure of the asymptotic efficiency for tests of a hypothesis based on the sum of observations, *Annals of Mathematical Statistics* 23 (1952) 493–507.
- [9] Fabrice Collard, Sujoy Mukerji, Kevin Sheppard, Jean-Marc Tallon, Ambiguity and the historical equity premium, (2008) universite des Sciences Sociales, University of Oxford and Universite Paris 1-Pantheon-Sorbonne.
- [10] Monroe Donsker, S. R. S. Varadhan, On a variational formula for the principal eigenvalue for operators with maximum principle, *Proceedings of the National Academy of Sciences* 72 (1975) 780–783.
- [11] Darrell Duffie, Larry G. Epstein, Stochastic differential utility, *Econometrica* 60 (2) (1992) 353–394.

- [12] Paul Dupuis, Richard S. Ellis, A weak convergence approach to the theory of large deviations, Wiley Series in Probability and Statistics Wiley-Interscience (1997).
- [13] I. Gilboa, D. Schmeidler, Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics* 18 (1989) 141–153.
- [14] Lars Peter Hansen, Beliefs, doubts and learning: Valuing macroeconomic risk, *American Economic Review* 97 (2) (2007) 1–30.
- [15] Lars Peter Hansen, Thomas J. Sargent, Robust estimation and control under commitment, *Journal of Economic Theory* 124 (2) (2005) 258 – 301.
- [16] Lars Peter Hansen, Thomas J. Sargent, Recursive robust estimation and control without commitment, *Journal of Economic Theory* 136 (1) (2007) 1–27.
- [17] Lars Peter Hansen, Thomas J. Sargent, Robustness, Princeton University Press, Princeton, New Jersey (2008).
- [18] Lars Peter Hansen, Thomas J. Sargent, Fragile beliefs and the price of uncertainty, *Quantitative Economics* 1 (1) (2010) 129–162.
- [19] Lars Peter Hansen, Thomas J. Sargent, Gauhar Turmuhambetova, Noah Williams, Robust control and model misspecification, *Journal of Economic Theory* 128 (1) (2006) 45–90.
- [20] Lars Peter Hansen, José A. Scheinkman, Long-term risk: an operator approach, *Econometrica* 77 (1) (2009) 177–234.
- [21] Takashi Hayashi, Jianjun Miao, Intertemporal substitution and recursive smooth ambiguity preferences, *Theoretical Economics*, forthcoming.
- [22] M. R. James, Asymptotic analysis of nonlinear stochastic risk sensitive control and differential games, *Mathematics of Control, Signals and Systems*, 5 (1992) 401–417.
- [23] Nengjui Ju, Jianjun Miao, Ambiguity, Learning and asset returns, (2010) Boston University and Hong Kong University of Science and Technology.

- [24] Peter Klibanoff, Massimo Marinacci, Sujoy Mukerji, A smooth model of decision making under ambiguity, *Econometrica* 73 (6) (2005) 1849–1892.
- [25] Peter Klibanoff, Massimo Marinacci, Sujoy Mukerji, Recursive smooth ambiguity preferences, *Journal of Monetary Economics* 144 (2009) 930–976.
- [26] D. M. Kreps, E. L. Porteus, Temporal resolution of uncertainty and dynamic choice, *Econometrica* 46 (1978) 185–200.
- [27] Robert S. Liptser, Albert N. Shiryaev, *Statistics of Random Processes*, Springer, Berlin second edition (2001).
- [28] Lars Ljungqvist, Thomas J. Sargent, *Recursive Macroeconomic Theory*, Second Edition, MIT Press, Cambridge, Ma. (2004).
- [29] C. M. Newman, B. W. Stuck, Chernoff bounds for discriminating between two Markov processes, *Stochastics* 2 (1979) 139 – 153.
- [30] Bezalel Peleg, Menahem E. Yaari, On the existence of a consistent course of action when tastes are changing, *The Review of Economic Studies* 40 (3) (1973) 391–401.
- [31] Thomas J. Sargent, *Dynamic Macroeconomic Theory*, Harvard University Press, Cambridge, Ma. (1987).
- [32] Costis Skiadas, Smooth ambiguity aversion toward small risks and continuous-time recursive utility, (2009) Northwestern University.