

QUANTIFICATION OF MODEL UNCERTAINTY ON PATH-SPACE VIA GOAL-ORIENTED RELATIVE ENTROPY

JEREMIAH BIRRELL*, MARKOS A. KATSOUAKIS AND LUC REY-BELLET

Abstract. Quantifying the impact of parametric and model-form uncertainty on the predictions of stochastic models is a key challenge in many applications. Previous work has shown that the relative entropy rate is an effective tool for deriving path-space uncertainty quantification (UQ) bounds on ergodic averages. In this work we identify appropriate information-theoretic objects for a wider range of quantities of interest on path-space, such as hitting times and exponentially discounted observables, and develop the corresponding UQ bounds. In addition, our method yields tighter UQ bounds, even in cases where previous relative-entropy-based methods also apply, *e.g.*, for ergodic averages. We illustrate these results with examples from option pricing, non-reversible diffusion processes, stochastic control, semi-Markov queueing models, and expectations and distributions of hitting times.

Mathematics Subject Classification. 62F35, 62B10, 60G40, 60J60, 93E20, 91G20.

Received March 27, 2020. Accepted October 5, 2020.

1. INTRODUCTION

Probabilistic models are widely used in engineering, finance, physics, chemistry, and many other fields. Models of real-world systems inherently carry uncertainty, either in the value of model parameters, or more general non-parametric (*i.e.*, model-form) uncertainty. Such uncertainty can stem from fitting a model to data or from approximating a more complicated, intractable model with one that is simpler and computationally tractable. It is often important to estimate or bound the effect of such uncertainties on quantities-of-interest (QoIs) computed from the model, *i.e.*, to obtain uncertainty quantification (UQ) bounds. This can be especially difficult for dynamical/time-series problems over a long time-horizon, such as the computation of hitting times, reaction rates, cost functions, and option values. When the uncertainty is parametric and small, linearization and perturbative methods are effective tools for quantifying the impact of uncertainty on model predictions; see, *e.g.*, [2, 4, 21, 28, 31, 36, 48, 49, 51]. However, new general methods are needed when the uncertainties are large or non-parametric. In this paper, we present a framework for addressing this problem for a wide range of path-space QoIs.

To illustrate the challenges inherent in model-form uncertainty on path-space, consider the following examples (see Sect. 5 below for further detail):

Keywords and phrases. Uncertainty quantification, relative entropy, non-reversible diffusion processes, semi-Markov queueing models, stochastic control.

Department of Mathematics and Statistics, University of Massachusetts Amherst, Amherst, MA 01003, USA.

*Corresponding author: birrell@math.umass.edu

- (1) In mathematical finance, the value of an asset, X_t , in a variable-interest-rate environment may be modeled by a stochastic differential equation of the form

$$dX_t = (r + \Delta r(t, X_t, Y_t))X_t dt + \sigma X_t dW_t, \quad r, \sigma > 0. \quad (1.1)$$

Many QoIs, such as option values, involve hitting times for X_t , and hence are naturally phrased on path-space with a potentially large time horizon. When the variable rate perturbation, Δr , is influenced by some other process, Y_t , that is incompletely known (such as in the Vasicek model), one may wish to approximate equation (1.1) with a simpler, tractable model, *e.g.*, geometric Brownian motion (obtained by setting $\Delta r \equiv 0$). If the true Δr is potentially large or has an incompletely known structure then perturbative methods are unable to provide reliable error estimates. One needs new, non-perturbative methods to obtain bounds on option values, and other QoIs.

- (2) In queuing systems, Markovian models are widely used due to their mathematical simplicity, even in cases where real-world data shows that waiting times are non-exponential, and hence the systems are non-Markovian [16]. To quantify the model uncertainty, one must estimate/bound the error incurred from using a Markovian approximation to a non-Markovian system. More generally, an approximate model may have a fundamentally different mathematical structure than the “true” model; such scenarios can be difficult to study through perturbative methods.

Variational-principle-based methods, relating expectations of a quantity of interest to information-theoretic divergences, have proven to be effective tools for deriving UQ bounds. While they are applicable to parametric uncertainty, such methods are particularly powerful when applied to systems with model-form uncertainty. Methods have been developed using relative entropy [15, 17, 27, 30, 39] and Rényi divergences [6, 21]. For stochastic processes, the relative entropy rate has been used to derive UQ bounds on ergodic averages [10, 20, 29, 35, 46]. Ergodic averages constitute an important class of QoIs on path space, but many other key QoIs are not covered by these previous methods, such as the expectation of hitting times and exponentially discounted observables. Here, we develop a new general framework for path-space UQ by identifying a larger class of information-theoretic objects (the so-called goal-oriented relative entropies) and proving an accompanying new bound, Theorem 3.6. This framework can be applied to qualitatively new regimes, such as the aforementioned hitting times or to ergodic averages in the presence of unbounded perturbations. In addition, even when the previous methods apply (*e.g.*, for ergodic averages with bounded perturbations), the framework developed here provides quantitatively tighter bounds.

In this work, UQ will refer to the following mathematical problem: One has a baseline model, described by a probability measure P , and an alternative model, \tilde{P} . For a given real-valued function (*i.e.*, QoI), F , one’s goal is to bound the error incurred from using P in place of \tilde{P} , *i.e.*,

$$\text{Bound the bias:} \quad E_{\tilde{P}}[F] - E_P[F]. \quad (1.2)$$

(Here, and in the following, E_Q will denote the expectation under the probability measure Q .) F is the quantity one is interested in, such as a hitting time or an ergodic average. The baseline model, P , is typically an approximate but “tractable” model, meaning one can calculate QoIs exactly, or it is relatively inexpensive to simulate. However, it generally contains many sources of error and uncertainty; it may depend on parameters with uncertain values (obtained from experiment, Monte-Carlo simulation, variational inference, etc.) or is obtained *via* some approximation procedure (coarse graining, neglecting memory terms, linearization, asymptotic approximation, etc.). Any QoI computed from P therefore has significant uncertainty associated with it. In contrast, \tilde{P} , is thought of as the “true” (or at least, a more precise) model, but due to its complexity or a lack of knowledge, it is intractable. For instance, in the finance example discussed above, \tilde{P} is the distribution on path space of the solution to equation (1.1), and one possible choice of P is the distribution of geometric Brownian motion (*i.e.*, the solution to equation (1.1) with $\Delta r \equiv 0$). In the queuing example, \tilde{P} is the distribution of the “true” non-Markovian model on path-space, while P is the distribution of the Markovian approximation.

TABLE 1. Examples of QoIs to which our methods apply.

	QoI
1	$\tilde{P}(\tau \leq T)$
2	$E_{\tilde{P}}[\tau]$
3	$E_{\tilde{P}}[\int_0^\infty f_s \lambda e^{-\lambda s} ds]$
4	$E_{\tilde{P}}[T^{-1} \int_0^T f(X_t) dt]$

The results in [15, 17, 27, 30, 39] produce bounds on the general UQ problem (1.2) in terms of the relative entropy, *i.e.*, Kullback–Leibler (KL) divergence, $R(\tilde{P}||P)$, which quantifies the discrepancy between the models \tilde{P} and P . When comparing stochastic processes, $\{X_t\}_{t \in [0, \infty)}$ and $\{\tilde{X}_t\}_{t \in [0, \infty)}$, one often finds that their distributions on path-space, $P_{[0, \infty)}$ and $\tilde{P}_{[0, \infty)}$, have infinite relative entropy. Hence, when dealing with an unbounded time horizon, these previous methods do not produce (nontrivial) UQ bounds. As we will see, this infinite relative entropy is often an indication that one is using the wrong information-theoretic quantity. In [20] it was shown how to overcome this issue when the QoI is an ergodic average; there, the relative entropy rate was found to be the proper information-theoretic quantity for bounding ergodic averages. See Section 2 for further discussion of this, and other background material. In this paper, we derive a general procedure for identifying information-theoretic quantities, matched to a particular QoI, that can then be used to obtain UQ bounds.

The intuition underlying our approach is conveyed by the following example: Suppose one wants to bound the \tilde{P} -expectation of a stopping time, τ . Even if τ is unbounded, the QoI $F = \tau$ depends only on the path of the process up to time τ , and not on the path for all $t \in [0, \infty)$. Hence, one would expect UQ bounds on τ to only require knowledge of the dynamics up to the stopping time; a bound on $R(\tilde{P}_{[0, \infty)}||P_{[0, \infty)})$ should not be necessary. Rather, one should be able to construct UQ bounds using the relative entropy between the baseline and alternative processes *up to the stopping time*:

$$R(\tilde{P}|_{\mathcal{F}_\tau}||P|_{\mathcal{F}_\tau}), \tag{1.3}$$

where \mathcal{F}_τ is the filtration for the process up to the stopping-time τ ; intuitively, \mathcal{F}_τ captures the information content of the process up to the stopping time. Equation (1.3) is an example of a goal-oriented relative entropy, for the QoI τ .

Our results will make the above intuition rigorous and general, and we will provide appropriate information-theoretic objects for computing UQ bounds on a variety of path-space QoIs, generalizing equation (1.3). Moreover, we will show how partial information regarding the structure of the alternative model, \tilde{P} , can be used to make such bounds *computable*. The method is developed in an abstract framework in Section 3, before being specializing to path-space QoIs in Section 4. The following list shows several classes of QoIs to which our methods apply (see also Tab. 1):

- (1) Cumulative distribution function (CDF) of a stopping time.
- (2) Expected value of a stopping time.
- (3) Expectation of exponentially discounted QoIs.
- (4) Expectation of time averages.

We will also show how one can bound goal-oriented relative entropies, such as equation (1.3), often a key step in obtaining computable UQ bounds. Finally, we illustrate these results through several examples in Section 5:

- (1) The distribution and expectation of hitting times of Brownian motion with constant drift, as compared to Brownian motion perturbed by a non-constant drift; see Section 5.1 and Appendix G.
- (2) UQ bounds for invariant measures of non-reversible stochastic differential equations (SDEs); see Section 5.2.
- (3) Robustness of linear-quadratic stochastic control under non-linear perturbations; see Section 5.3.

- (4) Robustness of continuous-time queueing models under non-exponential waiting-time perturbations (*i.e.*, semi-Markov perturbations); see Section 5.4.
- (5) The value of American put options in a variable-interest-rate environment (including the Vasicek model); see Section 5.5 and Appendix I.

2. BACKGROUND ON UNCERTAINTY QUANTIFICATION *via* INFORMATION-THEORETIC VARIATIONAL PRINCIPLES

In this section we record important background on the information-theoretic-variational-principle approach to UQ; the new methods presented in this paper (starting in Sect. 3) will be built upon these foundations.

First, we fix some notation. Let P be a probability measure on a measurable space (Ω, \mathcal{F}) . $\overline{\mathbb{R}}$ will denote the extended reals and we will refer to a random variable $F : \Omega \rightarrow \overline{\mathbb{R}}$ as a quantity-of-interest (QoI). If $F \in L^1(P)$ we write

$$\widehat{F} = F - E_P[F] \tag{2.1}$$

for the centered quantity of mean 0. The cumulant generating function of F is defined by

$$\Lambda_P^F(c) = \log E_P[e^{cF}], \tag{2.2}$$

where \log denotes the natural logarithm and we use the continuous extensions of $\exp(x)$ to $\overline{\mathbb{R}}$ and of $\log(x)$ to $[0, \infty]$.

Recall also that the relative entropy (*i.e.*, KL divergence) of another probability measure, \tilde{P} , with respect to P is defined by

$$R(\tilde{P}||P) = \begin{cases} E_{\tilde{P}} \left[\log \left(\frac{d\tilde{P}}{dP} \right) \right] & \text{if } \tilde{P} \ll P \\ +\infty & \text{otherwise} \end{cases}. \tag{2.3}$$

It has the property of a divergence, that is $R(\tilde{P}||P) \geq 0$ and $R(\tilde{P}||P) = 0$ if and only if $\tilde{P} = P$; see, for example, [19] for this and further properties.

The starting point of our approach to UQ on path-space is a KL-based UQ bound, which we here call the *UQ information inequality*. The bound was derived in [17, 20]; a similar inequality is used in the context of concentration inequalities, see *e.g.* [12], and was also used independently in [15, 27]. We summarize the proof in Appendix A for completeness: Let $F : \Omega \rightarrow \mathbb{R}$, $F \in L^1(\tilde{P})$. Then

$$\pm E_{\tilde{P}}[F] \leq \inf_{c>0} \left\{ \frac{1}{c} \Lambda_{\tilde{P}}^F(\pm c) + \frac{1}{c} R(\tilde{P}||P) \right\}. \tag{2.4}$$

If F is also in $L^1(P)$ then one can apply equation (2.4) to the centered QoI, \widehat{F} , to obtain centered UQ bounds, *i.e.*, bounds on $E_{\tilde{P}}[\widehat{F}] - E_P[\widehat{F}]$. Equation (2.4) is to be understood as two formulas, one with all the upper signs and another with all the lower signs. This remark applies to all other uses of \pm in this paper. The bounds (2.4) have many appealing properties, including tightness over relative-entropy neighborhoods [17, 20], a linearization formula in the case of small relative entropy [20, 39], and a divergence property [20]; see also [14, 15, 21].

In [10, 20, 35], equation (2.4) was used to derive UQ bounds for ergodic averages on path-space, both in discrete and continuous time. The goal of the present work is to extend these methods to apply to more general path-space QoIs. As motivation, here we provide a summary of one of the results from [20]: Let \mathcal{X} be a Polish space and consider two distributions, $P_{[0,T]}$ and $\tilde{P}_{[0,T]}$, on path-space up to time T , $C([0,T], \mathcal{X})$; these are the baseline and alternative models respectively. Assuming that the alternative process starts in an invariant distribution, $\tilde{\mu}^*$, and the baseline process starts in an arbitrary distribution, μ , one obtains the following UQ bound on ergodic averages (see Sect. 3.2 in [20] for this result and further generalizations):

Let $f : \mathcal{X} \rightarrow \mathbb{R}$ be a bounded observable and for $T > 0$ define the time average $A_T = T^{-1} \int_0^T f(X_t) dt$, where X_t is evaluation at time t for a path in $C([0, T], \mathcal{X})$. Then

$$\pm \left(E_{\tilde{P}_{[0, T]}}[A_T] - E_{P_{[0, T]}}[A_T] \right) \leq \inf_{c > 0} \left\{ \frac{1}{cT} \Lambda_{\tilde{P}_{[0, T]}}^{\hat{A}_T}(\pm cT) + \frac{1}{c} H(\tilde{P} \| P) + \frac{1}{cT} R(\tilde{\mu}^* \| \mu) \right\}, \quad (2.5)$$

where $H(\tilde{P} \| P)$ is the relative entropy rate:

$$H(\tilde{P} \| P) \equiv \lim_{T \rightarrow \infty} \frac{1}{T} R(\tilde{P}_{[0, T]} \| P_{[0, T]}). \quad (2.6)$$

To arrive at equation (2.5), one applies equation (2.4) to $T\hat{A}_T$, divides by T , and employs the chain-rule of relative entropy to write the result in terms of the relative entropy rate. Under appropriate assumptions, one can show that the upper bound remains finite as $T \rightarrow \infty$. In particular, the relative entropy rate between the alternative and baseline processes is positive and finite in many cases (see the online supplement to [20]), and hence provides an appropriate information-theoretic quantity for controlling ergodic averages in the long-time regime.

3. UQ BOUNDS VIA GOAL ORIENTED INFORMATION THEORY

Equation 2.5 provides UQ bounds on ergodic averages, but there are many other classes of path-space QoIs that are important in applications; see Table 1. Naively applying the UQ information inequality, equation (2.4), to one of these other path-space QoIs that has an infinite time horizon (*e.g.*, $F = \tau$ where τ is an unbounded stopping-time) will generally lead to trivial (*i.e.*, infinite) bounds. This is due to the fact that $R(\tilde{P}_{[0, \infty)} \| P_{[0, \infty)})$ is infinite in most cases (see Eq. (2.6) and the subsequent discussion). For ergodic averages, the argument leading to equation (2.5) circumvents this difficulty by expressing the bound in terms of the relative entropy rate. Although the derivation of (2.5) cannot be generalized to other QoIs, the idea of obtaining finite bounds by utilizing an alternative information theoretic quantity can be generalized, and we will do so in this section.

A naive application of equation (2.4) ignores a hidden degree of freedom: the choice of sigma algebra on which one defines the models. More specifically, if the QoI, F , is measurable with respect to \mathcal{G} , a sub sigma-algebra of the full sigma algebra \mathcal{F} on Ω , then one can write

$$E_{\tilde{P}}[F] - E_P[F] = E_{\tilde{P}|_{\mathcal{G}}}[F] - E_{P|_{\mathcal{G}}}[F], \quad (3.1)$$

where $Q|_{\mathcal{G}}$ denotes the restriction of a measure Q to the sub sigma-algebra \mathcal{G} ; \mathcal{G} captures the information content of F and working on \mathcal{G} will lead to UQ bounds that are targeted at the QoI. More specifically, one can apply equation (2.4) to the baseline model $P|_{\mathcal{G}}$ and alternative model $\tilde{P}|_{\mathcal{G}}$, thereby obtaining a UQ bound in terms of a *goal-oriented relative entropy* $R(\tilde{P}|_{\mathcal{G}} \| P|_{\mathcal{G}})$. We use the term goal-oriented in order to emphasize the fact that the information-theoretic quantity is tailored to the QoI under consideration, through the use of \mathcal{G} . As we will show, the data-processing inequality (see [40]) implies

$$R(\tilde{P}|_{\mathcal{G}} \| P|_{\mathcal{G}}) \leq R(\tilde{P} \| P), \quad (3.2)$$

and there are many situations where the inequality is strict. Thus one often obtains tighter UQ bounds by using a goal-oriented relative entropy. The above idea is quite general, and so we develop the theory in this section without any reference to path-space. We also prove new (quasi)convexity properties regarding optimization problems of the form (2.4).

Theorem 3.1 (Goal-Oriented Information Inequality). *Let P and \tilde{P} be probability measures on (Ω, \mathcal{F}) and $F : \Omega \rightarrow \mathbb{R}$.*

(1) Let \mathcal{G} be a sub sigma algebra of \mathcal{F} and $F \in L^1(\tilde{P})$ be \mathcal{G} -measurable. Then

$$\pm E_{\tilde{P}}[F] \leq \inf_{c>0} \left\{ \frac{1}{c} \Lambda_{\tilde{P}}^F(\pm c) + \frac{1}{c} R(\tilde{P}|_{\mathcal{G}}\|P|_{\mathcal{G}}) \right\} \quad (3.3)$$

$$= \inf_{\eta>0} \left\{ \eta \Lambda_{\tilde{P}}^F(\pm \eta^{-1}) + \eta R(\tilde{P}|_{\mathcal{G}}\|P|_{\mathcal{G}}) \right\}, \quad (3.4)$$

where we interpret $-\infty + \infty \equiv \infty$. The objective function $c \mapsto \frac{1}{c} \Lambda_{\tilde{P}}^F(\pm c) + \frac{1}{c} R(\tilde{P}|_{\mathcal{G}}\|P|_{\mathcal{G}})$ is quasiconvex and the objective function $\eta \mapsto \eta \Lambda_{\tilde{P}}^F(\pm \eta^{-1}) + \eta R(\tilde{P}|_{\mathcal{G}}\|P|_{\mathcal{G}})$ is convex.

(2) If $\mathcal{G} \subset \mathcal{H} \subset \mathcal{F}$ are sub sigma-algebras then

$$R(\tilde{P}|_{\mathcal{G}}\|P|_{\mathcal{G}}) \leq R(\tilde{P}|_{\mathcal{H}}\|P|_{\mathcal{H}}) \leq R(\tilde{P}\|P). \quad (3.5)$$

Proof. (1) For a real-valued $F \in L^1(\tilde{P})$, write

$$E_{\tilde{P}}[F] = E_{\tilde{P}|_{\mathcal{G}}}[F] \quad (3.6)$$

and then apply equation (2.4) to arrive at equation (3.3). The bound can then be extended to $\overline{\mathbb{R}}$ -valued F by taking limits. Next, by changing variables $c = 1/\eta$ in the infimum we arrive at equation (3.4). To see that the objective function in (3.4) is convex, first recall that the cumulant generating function $\Lambda_{\tilde{P}}^F(\pm c)$ is convex (this can be proven *via* Hölder's inequality). The perspective of a convex function is convex (see, *e.g.*, page 89 in [13] and note that the result can easily be extended to $\overline{\mathbb{R}}$ -valued convex functions), and therefore $\eta \mapsto \eta \Lambda_{\tilde{P}}^F(\pm \eta^{-1})$ is convex. Adding a linear term preserves convexity, hence $\eta \Lambda_{\tilde{P}}^F(\pm \eta^{-1}) + \eta R(\tilde{P}|_{\mathcal{G}}\|P|_{\mathcal{G}})$ is convex in η . Finally, quasiconvexity of the objective function in (3.3) follows from the general fact that if $h : (0, \infty) \rightarrow \overline{\mathbb{R}}$ is convex then $c \mapsto h(1/c)$ is quasiconvex on $(0, \infty)$. To see this, let $c_1, c_0 > 0$ and $\lambda \in (0, 1)$. Then $1/(\lambda c_1 + (1-\lambda)c_0) = s c_1^{-1} + (1-s)c_0^{-1}$ for some $s \in (0, 1)$. Therefore, convexity of h implies

$$h(1/(\lambda c_1 + (1-\lambda)c_0)) \leq s h(c_1^{-1}) + (1-s) h(c_0^{-1}) \leq \max\{h(c_1^{-1}), h(c_0^{-1})\}. \quad (3.7)$$

(2) If $\mathcal{G} \subset \mathcal{H} \subset \mathcal{F}$ are sub sigma-algebras then we can write $P|_{\mathcal{G}} = (id)_*(P|_{\mathcal{H}})$ (we let id denote the identity on Ω , thought of as a measurable map from (Ω, \mathcal{H}) to (Ω, \mathcal{G}) , and $\psi_* P$ denotes the distribution of the random quantity ψ), and similarly for \tilde{P} . The data processing inequality (see Thm. 14 in [40]) implies $R(\psi_* \tilde{P} \| \psi_* P) \leq R(\tilde{P} \| P)$ for any random quantity ψ , and so we arrive at equation (3.5). \square

Remark 3.2. The objective functions in both (3.3) and (3.4) are unimodal, due to (quasi)convexity. In the majority of our theoretical discussions we will use the form (3.3) of the UQ bound, but the convex form (3.4) is especially appealing for computational purposes, as the optimization can be done numerically *via* standard techniques and with guarantees on the convergence. In the examples in Section 5 we will generally perform the optimization numerically, using either a line search or Nelder-Mead simplex method [38, 43].

Remark 3.3. Equation (3.5) is a version of the data processing inequality, which holds in much greater generality than the form stated above. In particular, it holds for any f -divergence [40]. For relative entropy, the data-processing inequality can be obtained from the chain rule, together with the fact that marginalization can only reduce the relative entropy.

In particular, one can bound the probability of an event $A \in \mathcal{G}$ by applying Theorem 3.1 to $F = 1_A$ (see the example in Sect. 5.5.1):

Corollary 3.4. Let \mathcal{G} be a sub sigma-algebra of \mathcal{F} and $A \in \mathcal{G}$. Then

$$\pm \left(\tilde{P}(A) - P(A) \right) \leq \inf_{c>0} \left\{ \frac{1}{c} \log \left(P(A) e^{\pm c(1-P(A))} + (1-P(A)) e^{\mp c P(A)} \right) + \frac{1}{c} R(\tilde{P}|_{\mathcal{G}}\|P|_{\mathcal{G}}) \right\}. \quad (3.8)$$

Remark 3.5. UQ bounds based on Corollary 3.4 are far from optimal for very rare events (*i.e.*, when $P(A) \ll 1$). See [6, 21] for a related approach to UQ, using Rényi divergences instead of relative entropy, which produces tighter bounds for rare events.

Theorem 3.1 is particularly useful because of the manner in which it splits the UQ problem into two sub-problems:

- (1) Compute or bound the cumulant generating function of the QoI with respect to the baseline model, $\Lambda_P^F(c)$. (Note that there is no dependence on the alternative model, \tilde{P} .)
- (2) Bound the relative entropy of the alternative model with respect to the baseline model, $R(\tilde{P}|_{\mathcal{G}}\|P|_{\mathcal{G}})$, on an appropriate sub sigma-algebra, \mathcal{G} . (The primary difficulty here is the presence of the intractable/unknown model \tilde{P} .)

We generally consider the baseline model, P , and hence sub-problem (1), to be “tractable”. In practice, computing or bounding the CGF under P may still be a difficult task. Bounds on CGFs are often used in the derivation of concentration inequalities; see [12] for an overview of this much-studied problem. For prior uses of CGF bounds in UQ, see [10, 30, 44]. With these points in mind, our focus will primarily be on the second problem; our examples in Section 5 will largely involve baseline models for which the CGF is relatively straightforward to compute.

Theorem 3.6, below, shows how one can use partial information/bounds on the intractable model, \tilde{P} , to produce a computable UQ bound, thus addressing sub-problem 2. The key step is to select a sub sigma algebra \mathcal{G} and a \mathcal{G} -measurable $G \in L^1(\tilde{P})$ for which the following relative entropy bound holds:

$$R(\tilde{P}|_{\mathcal{G}}\|P|_{\mathcal{G}}) \leq E_{\tilde{P}}[G]. \quad (3.9)$$

Theorem 3.6. *Let $F : \Omega \rightarrow \overline{\mathbb{R}}$, $F \in L^1(\tilde{P})$, be \mathcal{G} -measurable, where $\mathcal{G} \subset \mathcal{F}$ is a sub sigma-algebra, and suppose we have a \mathcal{G} -measurable real-valued $G \in L^1(\tilde{P})$ that satisfies $R(\tilde{P}|_{\mathcal{G}}\|P|_{\mathcal{G}}) \leq E_{\tilde{P}}[G]$. Then*

$$\pm E_{\tilde{P}}[F] \leq \inf_{c>0} \left\{ \frac{1}{c} \log E_P [\exp(\pm cF + G)] \right\} \quad (3.10)$$

$$= \inf_{\eta>0} \left\{ \eta \log E_P [\exp(\pm \eta^{-1}F + G)] \right\}. \quad (3.11)$$

The objective function $(0, \infty) \rightarrow (-\infty, \infty]$, $c \mapsto \frac{1}{c} \log E_P [\exp(\pm cF + G)]$ is quasiconvex and the objective function $(0, \infty) \rightarrow (-\infty, \infty]$, $\eta \mapsto \eta \log E_P [\exp(\pm \eta^{-1}F + G)]$ is convex.

Remark 3.7. The bounds (3.10) (or, equivalently, (3.11)) are tight over the set of allowed G 's; see Appendix B for a precise statement and proof of this fact.

Proof. Given $c_0 > 0$, apply equation (3.3) from Theorem 3.1 and the assumption (3.9) to the \mathcal{G} -measurable QoI, $F_{\pm c_0} \equiv F \pm c_0^{-1}G$. Doing so yields

$$\begin{aligned} \pm E_{\tilde{P}}[F \pm c_0^{-1}G] &\leq \inf_{c>0} \left\{ \frac{1}{c} \Lambda_P^{F_{\pm c_0}}(\pm c) + \frac{1}{c} R(\tilde{P}|_{\mathcal{G}}\|P|_{\mathcal{G}}) \right\} \\ &\leq \inf_{c>0} \left\{ \frac{1}{c} \Lambda_P^{F_{\pm c_0}}(\pm c) + \frac{1}{c} E_{\tilde{P}}[G] \right\}. \end{aligned} \quad (3.12)$$

In particular, bounding the right-hand-side by the value at $c = c_0$, we can cancel the $E_{\tilde{P}}[G]$ term (which is assumed to be finite) to find

$$\pm E_{\tilde{P}}[F] \leq \frac{1}{c_0} \Lambda_P^{F_{\pm c_0}}(\pm c_0). \quad (3.13)$$

Taking the infimum over all $c_0 > 0$ yields (3.10). Equation (3.11) then follows by changing variables $c = 1/\eta$. To prove convexity of the objective function in equation (3.11), first note that the map $c \mapsto \log \int e^{\pm cF} d\mu$ is convex on $(0, \infty)$ for all positive measures, μ ; this can be proven using Hölder's inequality in the same way that one proves convexity of the CGF. Applying this to $d\mu = e^G dP$ we see that $\log E_P [\exp(\pm cF + G)]$ is convex in c . Therefore, by again using convexity of the perspective of a convex function, we find that $\eta \mapsto \eta \log E_P [\exp(\pm \eta^{-1}F + G)]$ is convex. Finally, using the quasiconvexity property proven above in equation (3.7), we see that $c \mapsto c^{-1} \log E_P [\exp(\pm cF + G)]$ is quasiconvex, thus completing the proof. \square

Remark 3.8. Previously, a primary strategy for using the non-goal-oriented bound (2.4) in UQ was to find a function G such that $R(\tilde{P}||P) \leq E_{\tilde{P}}[G]$, compute an explicit upper bound $G \leq D \in \mathbb{R}$, bound $R(\tilde{P}||P) \leq D$, and use this to obtain the (in principle) computable bounds

$$\pm E_{\tilde{P}}[F] \leq \inf_{c>0} \left\{ \frac{1}{c} \Lambda_P^F(\pm c) + \frac{D}{c} \right\}. \quad (3.14)$$

Of course, if G is not bounded then the resulting UQ bound is uninformative. The new result (3.10) is tighter than the above described strategy, as can be seen by bounding $G \leq D$ in the exponent of equation (3.10). In addition, (3.10) can give finite results even when G is not bounded, as we demonstrate in the examples in Sections 5.2 and 5.5.1.

Theorem 3.6 takes a relative entropy bound of the form (3.9) (which involves an expectation under the intractable model \tilde{P}) and produces a UQ bound which only involves expectations under the tractable model, P (see Eq. (3.10)). However, in practice it is not enough to use just any G that satisfies equation (3.9). If the relative entropy on \mathcal{G} is finite then the (trivial) choice of $G = \log(d\tilde{P}|_{\mathcal{G}}/dP|_{\mathcal{G}})$ always satisfies equation (3.9) (with equality). Rather, the practical question is whether one can find a *tractable* G that satisfies equation (3.9); the trivial choice $G = \log(d\tilde{P}|_{\mathcal{G}}/dP|_{\mathcal{G}})$ typically *does not* satisfy this tractability requirement. Given the fact that \tilde{P} may only be partially known, a tractable G will generally need to incorporate bounds/partial information regarding \tilde{P} .

The choice of an appropriate \mathcal{G} goes hand-in-hand with the choice of G . The bounds in Theorem 3.1 become tighter as one makes \mathcal{G} smaller, with the tightest bound obtained when $\mathcal{G} = \sigma(F)$, where $\sigma(F)$ is the sigma algebra generated by F (i.e., generated by $\{F^{-1}((a, \infty]) : a \in \mathbb{R}\}$), in which case the relative entropy is $R(\tilde{P}|_{\sigma(F)}||P|_{\sigma(F)}) = R(F_*\tilde{P}||F_*P)$. However, one must again balance the desire for tightness against the need for *computable* bounds. It is generally very difficult to find a tractable G for the relative entropy on $\sigma(F)$, and so a larger sub sigma-algebra must be used. While the task of finding a suitable \mathcal{G} and G is problem specific, for QoIs on path space there are general strategies one can follow. We discuss such strategies in Section 4 and Appendix D below; see also the concrete examples in Section 5.

In some cases, the right-hand-side of the UQ bound (3.10) is still difficult to compute, and so it can be useful to first bound $R(\tilde{P}|_{\mathcal{G}}||P|_{\mathcal{G}})$ and then use the UQ bound from Theorem 3.1 directly; for example, see the analysis of the Vasicek model in Section 5.5.1. The following corollary, obtained by applying Theorem 3.6 to $F = G$ and then reparametrizing the infimum, is useful for this purpose:

Corollary 3.9. *Suppose $G : \Omega \rightarrow \mathbb{R}$, $G \in L^1(\tilde{P})$ is \mathcal{G} -measurable for some sub sigma-algebra $\mathcal{G} \subset \mathcal{F}$, and $R(\tilde{P}|_{\mathcal{G}}||P|_{\mathcal{G}}) \leq E_{\tilde{P}}[G]$. Then*

$$R(\tilde{P}|_{\mathcal{G}}||P|_{\mathcal{G}}) \leq \inf_{\lambda>1} \{(\lambda - 1)^{-1} \Lambda_P^G(\lambda)\}. \quad (3.15)$$

Note that the optimization in (3.15) is over $\lambda > 1$, as opposed to over $c > 0$ like in Theorems 3.1 and 3.6.

Remark 3.10. Sensitivity analysis of a QoI to perturbations within a parametric family, P_θ , $\theta \in U \subset \mathbb{R}^d$, was also studied in [20]. The extension to goal-oriented relative entropy proceeds similarly, assuming one can

find a density for the restricted measure with respect to some common dominating measure: $dP_\theta|_{\mathcal{G}} = p_\theta^{\mathcal{G}} d\mu|_{\mathcal{G}}$. Examples of such densities for various classes of Markov processes can be found in Appendix D. Once one has $p_\theta^{\mathcal{G}}$, the computation of sensitivity bounds proceeds as in [20], and so we make no further comments on sensitivity analysis here.

Remark 3.11. When P is part of a tractable parametric family, $\{P_\theta\}_{\theta \in U}$, the following technique can be used to improve the UQ bounds: If one can carry out the procedure of Theorem 3.6 for each $\theta \in U$ then, by minimizing over θ , one obtains

$$\pm E_{\tilde{P}}[F] \leq \inf_{c>0} \left\{ \frac{1}{c} \inf_{\theta \in U} \log E_{P_\theta} [\exp(\pm cF + G_\theta)] \right\}. \quad (3.16)$$

See Figure I.1 in Appendix I for an example that employs this idea.

4. GOAL-ORIENTED UQ FOR QOIS UP TO A STOPPING TIME

We now specialize the methods of Section 3 to the study of path-space QoIs. Concrete examples can be found in Section 5, but we first discuss the guiding principles and common themes that underlie these applications. Section 4.1 sets up the framework for UQ on path-space. Specifically, we consider QoIs up to a stopping time. Many important problem types fit under this umbrella, such as discounted observables, the CDF of a stopping time, expectation of a stopping time, as well as the previously studied ergodic averages (see Tab. 1); Sections 4.2–4.4 discuss several of these problem types in further detail.

4.1. General setting for UQ on path space up to a stopping time

The general setting in which we derive path-space UQ bounds is as follows:

Assumption 4.1. *Suppose:*

- (1) $(\Omega, \mathcal{F}_\infty, \{\mathcal{F}_t\}_{t \in \mathcal{T}})$ is a filtered probability space, where $\mathcal{T} = [0, \infty)$, or $\mathcal{T} = \mathbb{Z}_0$. We define $\bar{\mathcal{T}} \equiv \mathcal{T} \cup \{\infty\}$.
- (2) P and \tilde{P} are probability measures on $(\Omega, \mathcal{F}_\infty)$.
- (3) $F : \bar{\mathcal{T}} \times \Omega \rightarrow \bar{\mathbb{R}}$, written $F_t(\omega)$, is progressively measurable (progressive), i.e., F is measurable and $F|_{[0,t] \times \Omega}$ is $\mathcal{B}([0,t]) \otimes \mathcal{F}_t$ -measurable for all $t \in \mathcal{T}$ (intervals refer to subsets of \mathcal{T}).
- (4) $\tau : \Omega \rightarrow \bar{\mathcal{T}}$ is a \mathcal{F}_t -stopping time.

We will derive UQ bounds for a process F_t , stopped at τ , which we denote by F_τ . First recall:

Lemma 4.2. F_τ is \mathcal{F}_τ -measurable, where

$$\mathcal{F}_\tau = \{A \in \mathcal{F}_\infty : \forall t \in \mathcal{T}, A \cap \{\tau \leq t\} \in \mathcal{F}_t\} \quad (4.1)$$

is the filtration up to the stopping-time τ , a sub sigma-algebra of \mathcal{F}_∞ .

For a general stopped QoI, F_τ , the following (uncentered) UQ bounds immediately follow from Theorems 3.1 and 3.6 with the choices $\mathcal{G} = \sigma(F_\tau)$ or $\mathcal{G} = \mathcal{F}_\tau$:

Corollary 4.3. *Suppose $F_\tau \in L^1(\tilde{P})$. Then*

$$\begin{aligned} \pm E_{\tilde{P}}[F_\tau] &\leq \inf_{c>0} \left\{ \frac{1}{c} \Lambda_P^{F_\tau}(\pm c) + \frac{1}{c} R((F_\tau)_* \tilde{P} \| (F_\tau)_* P) \right\} \\ &\leq \inf_{c>0} \left\{ \frac{1}{c} \Lambda_P^{F_\tau}(\pm c) + \frac{1}{c} R(\tilde{P}|_{\mathcal{F}_\tau} \| P|_{\mathcal{F}_\tau}) \right\}. \end{aligned} \quad (4.2)$$

If $R(\tilde{P}|_{\mathcal{F}_\tau} \| P|_{\mathcal{F}_\tau}) \leq E_{\tilde{P}}[G]$ for some \mathcal{F}_τ -measurable real-valued $G \in L^1(\tilde{P})$ then

$$\pm E_{\tilde{P}}[F_\tau] \leq \inf_{c>0} \left\{ \frac{1}{c} \log E_P [\exp(\pm c F_\tau + G)] \right\}. \quad (4.3)$$

The objective functions in (4.2) and (4.3) are quasiconvex

Remark 4.4. This result shows that for stopped QoIs, $R(\tilde{P}|_{\mathcal{F}_\tau} \| P|_{\mathcal{F}_\tau})$ is an appropriate goal-oriented relative entropy. One can similarly obtain centered variants of (4.2) and (4.3) by working with $F_\tau - E_P[F_\tau]$ (we omit the details). As in Theorems 3.1 and 3.6, one can change variables $\eta = 1/c$ to obtain convex objective functions.

Corollary 4.3 provides general-purpose UQ bounds for path-space QoIs up to a stopping time. In the following subsections, we discuss the nuances involved in deriving UQ bounds for several of the classes of QoIs from Table 1. In some cases we directly apply Corollary 4.3, while others benefit from a slightly altered approach. Not all of our examples in Section 5 below will fit neatly into one of these categories, but rather, they illustrate several important use-cases and strategies.

4.2. Expectation of a stopping time

Here we consider the expectation of a stopping time, τ . For $\tau \in L^1(P) \cap L^1(\tilde{P})$, (the centered variant of) equation (4.2) in Corollary 4.3 applied to $F_t = t$ yields

$$\pm(E_{\tilde{P}}[\tau] - E_P[\tau]) \leq \inf_{c>0} \left\{ \frac{1}{c} \Lambda_{\tilde{P}}(\pm c) + \frac{1}{c} R(\tilde{P}|_{\mathcal{F}_\tau} \| P|_{\mathcal{F}_\tau}) \right\}. \quad (4.4)$$

For unbounded $\tau \in L^1(P) \cap L^1(\tilde{P})$, it is often useful to use Lemma C.4 from Appendix C to obtain:

$$\pm(E_{\tilde{P}}[\tau] - E_P[\tau]) \leq \inf_{c>0} \left\{ \frac{1}{c} \Lambda_{\tilde{P}}(\pm c) + \frac{1}{c} \liminf_{n \rightarrow \infty} R(\tilde{P}|_{\mathcal{F}_{\tau \wedge n}} \| P|_{\mathcal{F}_{\tau \wedge n}}) \right\}. \quad (4.5)$$

To obtain a computable bound from equation (4.5), first note that a relative entropy bound of the following form often follows from Girsanov's theorem (see Appendix D):

$$R(\tilde{P}|_{\mathcal{F}_{\tau \wedge n}} \| P|_{\mathcal{F}_{\tau \wedge n}}) \leq \eta_0 + K E_{\tilde{P}}[\tau \wedge n] \quad (4.6)$$

for some $\eta_0, K \in [0, \infty)$ and all $n \in \mathbb{Z}^+$. One can then use Theorem 3.6 with $F = \tau \wedge n$ and $G = \eta_0 + K(\tau \wedge n)$. Taking $n \rightarrow \infty$ *via* the monotone and dominated convergence theorems and then reparameterizing the infimum results in the following:

Corollary 4.5. *Assume the relative entropy satisfies a bound of the form equation (4.6) for all $n \in \mathbb{Z}^+$. Then*

$$- \inf_{c>0} \left\{ (c + K)^{-1} (\Lambda_{\tilde{P}}(-c) + \eta_0) \right\} \leq E_{\tilde{P}}[\tau] \leq \inf_{\lambda>K} \left\{ (\lambda - K)^{-1} (\Lambda_{\tilde{P}}(\lambda) + \eta_0) \right\}. \quad (4.7)$$

Again, note that the upper and lower bounds only involve the tractable model, P . We illustrate this technique in a concrete example in Appendix G.

4.3. Time-integral/discounted QoIs

Here we consider exponentially discounted QoIs. In the language of Assumption 4.1, given a progressive process f_s and some $\lambda > 0$ we define

$$F_t = \int_0^t f_s \lambda e^{-\lambda s} ds \quad (4.8)$$

and will consider the stopped process F_τ , with stopping-time $\tau = \infty$. Such QoIs are often used in control theory (see page 38 in [37]) and economics (see page 64 in [45] and page 147 in [47]). Here we derive an alternative UQ bound to Corollary 4.3 that is applicable to such QoIs:

Theorem 4.6. Let $\lambda > 0$ and define $\pi(ds) = \lambda e^{-\lambda s} ds$. Suppose $f : [0, \infty) \times \Omega \rightarrow \bar{\mathbb{R}}$ is progressively measurable

$$\text{and either } f_s \geq 0 \text{ for } \pi\text{-a.e. } s \text{ or } E_{\tilde{P}} \left[\int_0^\infty |f_s| d\pi \right] < \infty. \quad (4.9)$$

Then

$$\begin{aligned} \pm E_{\tilde{P}} \left[\int_0^\infty f_s \lambda e^{-\lambda s} ds \right] &\leq \int_0^\infty \inf_{c>0} \left\{ \frac{1}{c} \Lambda_{\tilde{P}}^{f_s}(\pm c) + \frac{1}{c} R(\tilde{P}|_{\mathcal{F}_s} \| P|_{\mathcal{F}_s}) \right\} \lambda e^{-\lambda s} ds \\ &\leq \inf_{c>0} \left\{ \int_0^\infty \frac{1}{c} \Lambda_{\tilde{P}}^{f_s}(\pm c) \lambda e^{-\lambda s} ds + \frac{1}{c} \int_0^\infty R(\tilde{P}|_{\mathcal{F}_s} \| P|_{\mathcal{F}_s}) \lambda e^{-\lambda s} ds \right\}. \end{aligned} \quad (4.10)$$

Remark 4.7. Theorem 4.6 can be generalized to other discounting measures, π , and other stopping times; see Appendix E.

Proof. First suppose $E_{\tilde{P}} \left[\int_0^\infty |f_s| d\pi \right] < \infty$. This allows us to use Fubini's theorem to compute

$$E_{\tilde{P}} \left[\int_0^\infty f_s \lambda e^{-\lambda s} ds \right] = \int_0^\infty E_{\tilde{P}} [f_s] \pi(ds) \quad (4.11)$$

and implies $E_{\tilde{P}} [f_s]$ is finite for π -a.e. s . For each s , f_s is \mathcal{F}_s -measurable, hence Theorem 3.1 implies

$$\pm E_{\tilde{P}} [f_s] \leq \inf_{c>0} \left\{ \frac{1}{c} \Lambda_{\tilde{P}}^{f_s}(\pm c) + \frac{1}{c} R(\tilde{P}|_{\mathcal{F}_s} \| P|_{\mathcal{F}_s}) \right\} \quad (4.12)$$

for π -a.e. s . Integrating over s gives the first inequality in equation (4.10). The second follows by pulling the infimum outside of the integral.

If $f_s \geq 0$ then repeat the above calculations for $f_s^N \equiv f_s 1_{f_s \leq N} + N 1_{f_s > N}$ and take $N \rightarrow \infty$, using the dominated and monotone convergence theorems. \square

The last line in equation (4.10) consists of two terms, a discounted moment generating function and a discounted relative entropy,

$$D \equiv \int_0^\infty R(\tilde{P}|_{\mathcal{F}_s} \| P|_{\mathcal{F}_s}) \lambda e^{-\lambda s} ds. \quad (4.13)$$

Equation (4.13) is the same information-theoretic quantity that was defined in [47] (see page 147), where it was proposed as a measure of model uncertainty for control problems. Theorem 4.6 provides a rigorous justification for its use in UQ for exponentially discounted QoIs.

Remark 4.8. $\pi(dx) = \lambda e^{-\lambda s} ds$ is a probability measure, and so one could alternatively apply Theorem 3.1 to the product measures $\pi \times P$ and $\pi \times \tilde{P}$. However, Jensen's inequality implies that the bound equation (4.10) is tighter. Alternatively, one might attempt to use Corollary 4.3 in place of Theorem 4.6, in which case one generally finds that the relative-entropy term, $R(\tilde{P}|_{\mathcal{F}_\infty} \| P|_{\mathcal{F}_\infty})$, is infinite and so the corresponding UQ bound is trivial and uninformative. On the other hand, the bound equation (4.10) is often nontrivial; see the example in Section 5.3 below.

4.4. Time-averages

Finally, we remark that ergodic averages are a special case of the current framework: Let $T > 0$, $F_t = t^{-1} \int_0^t f_s ds$, and $\tau = T$ (one can similarly treat the discrete-time case). F_τ is \mathcal{F}_T -measurable and so, if $F_\tau \in L^1(P) \cap L^1(\tilde{P})$, one can apply Corollary 4.3 to \tilde{F}_τ . Reparameterizing the infimum $c \rightarrow cT$, one finds that the bound is the same as equation (2.5), which was previously derived in [20] and used in [10]. As we will demonstrate in the example in Section 5.2 below, one can often use Theorem 3.6 to obtain tighter and more general UQ bounds on ergodic averages than those obtained from equation (2.5).

5. EXAMPLES

We now apply the goal-oriented UQ methods developed above to several examples:

- (1) Hitting times for perturbations of Brownian motion in Section 5.1.
- (2) UQ bounds for invariant measures of non-reversible SDEs in Section 5.2.
- (3) Expected cost in stochastic control systems in Section 5.3.
- (4) Long-time behavior of semi-Markov perturbations of an $M/M/\infty$ queue in Section 5.4.
- (5) Pricing of American put options with variable interest rate in Section 5.5.

5.1. Perturbed Brownian motion

We first illustrate the use of our method, and several of its features, with a simple example wherein many computations can be done explicitly and exactly; the example consists of perturbations to Brownian-motion-with-constant-drift. Specifically, take the baseline model to be the distribution on path-space, $C([0, \infty), \mathbb{R})$, of a \mathbb{R} -valued Wiener process (Brownian motion), W_t , with constant drift, $\mu > 0$, starting from 0, *i.e.*, the distribution of the following process:

Baseline Model:
$$X_t = \mu t + W_t. \tag{5.1}$$

The alternative model will be the distribution of the solution to a SDE of the form:

Alternative Models:
$$d\tilde{X}_t = (\mu + \beta(t, \tilde{X}_t, \tilde{Y}_t, \tilde{Z}))dt + d\tilde{W}_t, \quad \tilde{X}_0 = 0, \tag{5.2}$$

where \tilde{W}_t is also a Wiener process. We allow the drift perturbation, β , to depend on an additional \mathbb{R}^k -valued process, \tilde{Y}_t , and also on independent external data, \tilde{Z} (see Appendix D.3 for further discussion of the type of perturbations we have in mind). More specific assumptions on β will be stated later. We note that the methods developed here do not allow one to perturb the diffusion, as the corresponding relative entropy is infinite due to lack of absolute continuity (see, *e.g.*, page 80 in [24]).

In Section 5.1.1 we will study the cumulative distribution function of τ_a , the level- a hitting time:

QoI 1:
$$\tilde{P}^0(\tau_a \leq T), \quad T > 0, \quad \text{where } \tau_a[y] \equiv \inf\{t : y_t = a\}. \tag{5.3}$$

Note that τ_a is a stopping time on path-space, $C([0, \infty), \mathbb{R})$. \tilde{P}^0 denotes the distribution of \tilde{X}_t on path-space and similarly for P^0 ; the superscript 0 indicates that the initial condition is 0.

We will also derive UQ bounds for the expected level- a hitting time,

QoI 2:
$$E_{\tilde{P}^0}[\tau_a], \quad a > 0. \tag{5.4}$$

This second QoI is simpler, as we can use the general result of Section 4.2. Therefore we relegate the details to Appendix G.

5.1.1. Hitting time distribution: benefits of being goal-oriented

We use the QoI (5.3) to illustrate the benefits of using a goal-oriented relative entropy, versus a non-goal-oriented counterpart. Here we specialize to perturbed SDEs, equation (5.2), whose distributions, \tilde{P}^0 , satisfy the following.

Alternative Models: Perturbations by drifts, β , such that, for some $\alpha > 0$:
$$\tag{5.5}$$

$$R(\tilde{P}^0|_{\mathcal{F}_{\tau_a \wedge T}} \| P^0|_{\mathcal{F}_{\tau_a \wedge T}}) \leq \frac{1}{2} \alpha^2 E_{\tilde{P}^0}[\tau_a \wedge T].$$

For example, equation (5.5) holds if $\|\beta\|_\infty \leq \alpha$ (assuming that the models satisfy the assumptions required to obtain the relative entropy *via* Girsanov's theorem; see Appendix D.3 for details).

The QoI, $F \equiv 1_{\tau_a \leq T}$, is $\mathcal{F}_{\tau_a \wedge T}$ -measurable, and so $R(\tilde{P}^0|_{\mathcal{F}_{\tau_a \wedge T}} \| \tilde{P}^0|_{\mathcal{F}_{\tau_a \wedge T}})$ is an appropriate goal-oriented relative entropy; the bound equation (5.5) implies that we can take $\mathcal{G} = \mathcal{F}_{\tau_a \wedge T}$ and $G = \frac{1}{2}\alpha^2(\tau_a \wedge T)$ in Theorem 3.6. This gives the following goal-oriented UQ bound:

$$\pm \tilde{P}^0(\tau_a \leq T) \leq \inf_{c>0} \left\{ \frac{1}{c} \log E_{P^0} \left[\exp \left(\pm c 1_{\tau_a \leq T} + \frac{1}{2} \alpha^2(\tau_a \wedge T) \right) \right] \right\}. \quad (5.6)$$

For comparison, we also compute the non-goal-oriented bound that follows from the UQ information inequality (2.4). The QoI is a function of the path up to time T , so we apply this proposition to the relative entropy $R(\tilde{P}_T^0 \| P_T^0)$ (the subscript T denotes taking the distribution on path-space $C([0, T], \mathbb{R})$).

Remark 5.1. One can attempt to use the full distributions on $C([0, \infty), \mathbb{R})$ in equation (2.4), but that only results in trivial bounds, as the maximum of $R(\tilde{P}^0 \| P^0)$ over the collection of alternative models (5.5) is infinite; this can be seen by taking a constant perturbation $\beta = \alpha$.

Repeating these computations under the assumption $\|\beta\|_\infty \leq \alpha$, but using the non-goal-oriented relative entropy bound $R(\tilde{P}|_{\mathcal{F}_T} \| P|_{\mathcal{F}_T}) \leq \alpha^2 T/2$ (once again using Girsanov's theorem), we obtain the UQ bound

$$\pm \tilde{P}^0(\tau_a \leq T) \leq \inf_{c>0} \left\{ \frac{1}{c} \log E_{P^0} [\exp(\pm c 1_{\tau_a \leq T})] + \frac{1}{c} \frac{\alpha^2 T}{2} \right\}. \quad (5.7)$$

Note that the goal-oriented bound on the right hand side of equation (5.6) is tighter than the non-goal-oriented bound (5.7), as can be seen by bounding $\frac{1}{2}\alpha^2(\tau_a \wedge T) \leq \frac{1}{2}\alpha^2 T$ in the exponent.

The P^0 -expectations in (5.6) and (5.7) can be computed using the known formula for the distribution of τ_a under the baseline model (see page 196 in [34] and also Appendix F). Assuming that $a > 0$ and $\mu > 0$, the goal-oriented bound is

$$\pm \tilde{P}^0(\tau_a \leq T) \leq \inf_{c>0} \left\{ \frac{1}{c} \log \left[\frac{a}{\sqrt{2\pi}} \int_0^\infty e^{\pm c 1_{z \leq T} + \frac{1}{2} \alpha^2(z \wedge T)} e^{-(a-\mu z)^2/(2z)} z^{-3/2} dz \right] \right\}. \quad (5.8)$$

Similarly, one can compute the non-goal-oriented bound, equation (5.7), and the probability under the baseline model.

Figure 1 shows numerical results comparing these bounds, with parameter values $\mu = 1$, $a = 2$, $\alpha = 0.2$; we compute the integrals numerically using a quadrature method and, as discussed in Remark 3.2, we solve the 1-D optimization problems numerically. The black curve shows the distribution of τ_a under the baseline model. The blue curves show the goal-oriented bounds equation (5.8); these constrain $\tilde{P}^0(\tau_a \leq T)$ to be within the gray region. The red curves show the non-goal-oriented bounds (5.7); we see that the non-goal-oriented bounds significantly overestimate the uncertainty, especially as $P^0(\tau_a \leq T)$ approaches 1.

5.2. UQ for time-averages and invariant measures of non-reversible SDEs

In this example we show how Theorem 3.6 can be combined with the functional inequality methods from [10] to obtain improved UQ bounds for invariant measures of non-reversible SDEs. Specifically, consider the following class of SDEs on \mathbb{R}^d :

Baseline Model:
$$dX_t = (-\nabla V(X_t) + a(X_t))dt + \sqrt{2}dW_t, \quad X_0 \sim \mu^* = \frac{1}{Z} e^{-V(x)} dx, \quad (5.9)$$

where the non-gradient portion of the drift, a , satisfies

$$\operatorname{div}(e^{-V} a) = 0. \quad (5.10)$$

If $a = 0$ then equation (5.9) is an overdamped Langevin equation with potential V . In general, equation (5.9) is non-reversible but equation (5.10) ensures that μ^* remains the invariant measure for $a \neq 0$; such non-reversible terms are used for variance reduction in Langevin samplers [50].




FIGURE 1. *Left:* bound on the CDF of τ_a , the level- a hitting time, for a perturbation of Brownian-motion-with-constant-drift. The solid black line shows the distribution under the baseline model. The red curves (dot-dashed, Eq. (5.7)) are obtained from a non-goal-oriented relative entropy, while the tighter bounds, in blue (dashed, Eq. (5.8)), were obtained by our method with an appropriate goal-oriented relative entropy. *Right:* the (convex) objective function from (3.11) for the goal-oriented upper UQ bound (*i.e.*, corresponding to the upper blue dashed curve on the left).

The alternative model will have a general non-reversible drift:

Alternative Models:
$$d\tilde{X}_t = (-\nabla V(\tilde{X}_t) + b(\tilde{X}_t))dt + \sqrt{2}d\tilde{W}_t, \quad \tilde{X}_0 \sim \tilde{\mu}. \quad (5.11)$$

Denote their respective distributions on path-space, $C([0, \infty), \mathbb{R}^d)$, by P and \tilde{P} . We will obtain UQ bounds for ergodic averages of some observable f :

QoI:
$$F_T \equiv \frac{1}{T} \int_0^T f(x_t)dt. \quad (5.12)$$

Remark 5.2. While the invariant measure of the baseline process is known explicitly (up to a normalization factor; see Eq. 5.9), the invariant measure of the alternative model (5.11) is *not* known for general b ; physically, such non-reversible drifts model external forces. A general drift can be split into $-\nabla V$ and b *via* a Helmholtz decomposition (*i.e.*, b is the divergence-free component). Such a decomposition can inform the choice of V in the baseline model.

Assuming that Girsanov's theorem applies (see Appendix D.3), we have the following expression for the relative entropy

$$R(\tilde{P}|_{\mathcal{F}_T} \| P|_{\mathcal{F}_T}) = E_{\tilde{P}}[G_T], \quad G_T \equiv R(\tilde{\mu} \| \mu^*) + \frac{1}{4} \int_0^T \|b(x_t) - a(x_t)\|^2 dt. \quad (5.13)$$

The quantity G_T is \mathcal{F}_T -measurable and if $R(\tilde{\mu} \| \mu^*)$ is finite and $b - a$ grows sufficiently slowly at infinity then $G_T \in L^1(\tilde{P})$ and hence Theorem 3.6 implies

$$\pm E_{\tilde{P}}[F_T] \leq \inf_{c>0} \left\{ \frac{1}{cT} \log E_P[\exp(\pm cTF_T + G_T)] \right\}, \quad (5.14)$$

where we re-indexed $c \rightarrow cT$.

We can now use the functional inequality approach of [10] to bound equation (5.14). First write

$$\begin{aligned} \pm cTF_T + G_T &= \int_0^T h_c(x_t)dt \equiv TH_{c,T}, \\ h_c(x) &\equiv R(\tilde{\mu}||\mu^*)/T \pm cf(x) + \frac{1}{4}\|b(x) - a(x)\|^2. \end{aligned} \quad (5.15)$$

Utilizing a connection with the Feynman Kac semigroup, the cumulant generating function in (5.14) can be bounded as follows (see Thm. 1 in [53] or Prop. 6 in [10]):

$$\begin{aligned} \frac{1}{T}\Lambda_P^{H_{c,T}}(T) &\leq \kappa(h_c), \\ \kappa(h_c) &\equiv \sup \left\{ \int gA[g]d\mu^* + \int h_c|g|^2d\mu^* : g \in D(A, \mathbb{R}), \|g\|_{L^2(\mu^*)} = 1 \right\}, \end{aligned} \quad (5.16)$$

where A denotes the generator of the baseline SDE (5.9). Note that κ depends only on the symmetric part of A , while the non-reversible component of the drift, a , only contributes to the antisymmetric part of A . Therefore, it suffices to consider the reversible case $a = 0$ when bounding $\kappa(h)$. (The precise form of a will impact the quality of the final UQ bound; see Rem. 5.3).

$\kappa(h_c)$ can be bounded using functional inequalities for A (e.g., Poincaré, log-Sobolev, Liapunov functions). For example, a log-Sobolev inequality for A ,

$$\int g^2 \log(g^2)d\mu^* \leq -\beta \int gA[g]d\mu^* \quad \text{for all } g \in D(A, \mathbb{R}), \quad (5.17)$$

where $\beta > 0$, implies

$$\frac{1}{T}\Lambda_P^{H_{c,T}}(T) \leq \beta^{-1} \log \left(\int \exp(\beta h_c(x))\mu^*(dx) \right) \quad (5.18)$$

(see Cor. 4 in [53] or Prop. 7 in [10]) and hence

$$\pm E_{\tilde{P}}[F_T] \leq \inf_{c>0} \left\{ \frac{1}{c\beta} \log \left(\int \exp(\beta h_c(x))\mu^*(dx) \right) \right\}. \quad (5.19)$$

If the alternative process is started in a known initial distribution, $\tilde{\mu}$, for which $R(\tilde{\mu}||\mu^*)$ is finite and can be computed (or bounded) then the right-hand-side of equation (5.19) only involves the baseline process and other known quantities. It therefore provides (in principle) computable UQ bounds on the expectations of time-averages at finite-times. Alternatively, by taking $T \rightarrow \infty$ the dependence on the initial distribution in (5.14) disappears and we can obtain UQ bounds on the invariant measure of the alternative process. In the remainder of this subsection we explore this second direction.

By starting the alternative process in an invariant measure $\tilde{\mu} = \tilde{\mu}^*$ and taking $T \rightarrow \infty$ in equation (5.19) we obtain the following UQ bounds on the invariant measure of the alternative model:

$$\pm \int f(x)\tilde{\mu}^*(dx) \leq \inf_{c>0} \left\{ \frac{1}{c\beta} \log \left(\int \exp \left(\pm c\beta f(x) + \frac{\beta}{4}\|b(x) - a(x)\|^2 \right) \mu^*(dx) \right) \right\}. \quad (5.20)$$

Note that the bounds (5.20) only involve expectations under the baseline invariant measure, μ^* , and are finite for appropriate unbounded f and b . The ability to treat unbounded perturbations is due to our use of Theorem 3.6, as opposed to equation (2.5), and constitutes a significant improvement on the results in [10, 20], which are

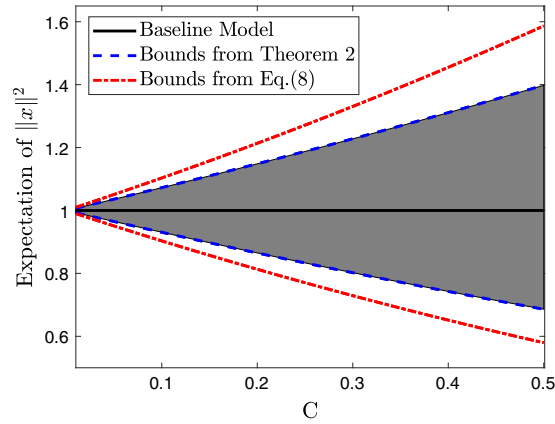


FIGURE 2. Baseline model: $V(x) = \|x\|^2$, $a = 0$ ($d = 2$); $D^2V \geq 2I$ and so a log-Sobolev inequality holds with $\beta = 1$ (see [8]). Perturbation: $b(x) = C(-\sin(x_2), \cos(x_1))$. Observable: $f(x) = \|x\|^2$. The baseline QoI, $\int f d\mu^*$, is shown in the horizontal black line. The bounds developed in Theorem 3.6 constrain the mean under the alternative model to the gray region (see Eq. (5.20)). These bounds are tighter than those obtained from equation (2.5) (see [20]) together with equation (5.21) (shown in red).

tractable only when $b - a$ is bounded. In Figure 2 we compare equation (5.20) with the UQ bound (2.5) from [20], where in the latter we employ the relative entropy rate bound

$$H(\tilde{P}\|P) \leq \|b - a\|_\infty^2/4, \quad (5.21)$$

i.e., we use the strategy described in Remark 3.8. Even in the case where $b - a$ is bounded, the new result (5.20) obtained using Theorem 3.6 is tighter than the bounds obtained from the earlier result (2.5).

Remark 5.3. The left hand side in equation (5.20) is independent of the choice of a in the baseline model. Therefore, one can view a in equation (5.20) as a parameter to be optimized over (minimizing over the set of a 's satisfying Eq. (5.10)).

5.3. Discounted quadratic observable for non-Gaussian perturbations of a Gaussian process

Here we consider a \mathbb{R}^n -valued progressively measurable process, X_t , that has a Gaussian distribution for every t :

Baseline Model:
$$(X_t)_*P = N(\mu_t, \Sigma_t), \quad t \geq 0. \quad (5.22)$$

The mean and variance of the baseline model are considered to be known (*i.e.*, obtained by some approximation or measurement). We will compute UQ bounds on non-Gaussian perturbations of (5.22). We assume Σ_t is (strictly) positive-definite and study the discounted QoI:

QoI:
$$F \equiv \int_{\mathcal{T}} \left(\frac{1}{2} \hat{X}_s^T C_s \hat{X}_s + d_s^T \hat{X}_s \right) \pi(ds), \quad (5.23)$$

where C_s is a positive-definite $n \times n$ -matrix-valued function, d_s is \mathbb{R}^n -valued, and π is a probability measure. Without loss of generality (*i.e.*, neglecting an overall constant factor and after a redefinition of d_s), we have chosen to express F in terms of the centered process $\hat{X}_s = X_s - \mu_s$. The techniques involved here are those of Section 4.3.

At this point, we are being purposefully vague about the nature of the alternative model and even the details of the baseline model, as a large portion of the following computation is independent of these details. We also emphasize that the alternative models do *not* need to be Gaussian. In Section 5.3.1 we will make both base and alternative models concrete, as we study an application to control theory.

Assuming that the conditions of Theorem E.1 of Appendix E are met, we obtain the following UQ bound (note that the appropriate sub sigma-algebras are $\mathcal{G} = \mathcal{F}_s$):

$$\begin{aligned} \pm E_{\tilde{P}}[F] &\leq \int_{\mathcal{T}} \inf_{c>0} \left\{ \frac{1}{c} \Lambda_P^{f_s}(\pm c) + \frac{1}{c} R(\tilde{P}|_{\mathcal{F}_s} \| P|_{\mathcal{F}_s}) \right\} \pi(ds), \\ f_s &\equiv \frac{1}{2} \widehat{X}_s^T C_s \widehat{X}_s + d_s^T \widehat{X}_s. \end{aligned} \quad (5.24)$$

X_t is Gaussian under P , so we can compute

$$\begin{aligned} \Lambda_P^{f_s}(c) &= \log \left(((2\pi)^n |\det(\Sigma_s)|)^{-1/2} \int e^{cd_s^T z - z^T (\Sigma_s^{-1} - cC_s) z / 2} dz \right), \\ E_P[F] &= \int_{\mathcal{T}} \frac{1}{2} \text{tr}(\Sigma_s C_s) \pi(ds). \end{aligned} \quad (5.25)$$

For c small enough, $\Sigma_s^{-1} - cC_s$ is positive-definite. More specifically, C_s is positive-definite, so we can compute a Cholesky factorization $C_s = M_s^T M_s$ with M_s nonsingular. $\Sigma_s^{-1} - cC_s$ is positive-definite if and only if $y^T (M_s^{-1})^T \Sigma_s^{-1} M_s^{-1} y > cy^T y$ for all nonzero y , which is true if and only if $c < 1/\|M_s \Sigma_s M_s^T\|$ (ℓ^2 -matrix norm).

For $c < 1/\|M_s \Sigma_s M_s^T\| \equiv c_s^*$, the integral in equation (5.25) can be computed in terms of the moment generating function of a Gaussian with covariance $(\Sigma_s^{-1} - cC_s)^{-1}$, resulting in the UQ bound

$$- \int_{\mathcal{T}} \inf_{c>0} \left\{ \frac{1}{c} \Lambda_P^{f_s}(-c) + \frac{\eta_s}{c} \right\} \pi(ds) \leq E_{\tilde{P}}[F] \leq \int_{\mathcal{T}} \inf_{0 < c < c_s^*} \left\{ \frac{1}{c} \Lambda_P^{f_s}(c) + \frac{\eta_s}{c} \right\} \pi(ds), \quad (5.26)$$

where $\eta_s \equiv R(\tilde{P}|_{\mathcal{F}_s} \| P|_{\mathcal{F}_s})$ and for the above ranges of c we have

$$\Lambda_P^{f_s}(\pm c) = -\frac{1}{2} \log(|\det(I \mp c \Sigma_s C_s)|) + \frac{1}{2} c^2 d_s^T (I \mp c \Sigma_s C_s)^{-1} \Sigma_s d_s.$$

To make these bounds computable, one needs a bound on the relative entropy; this requires more specificity regarding the nature of the family of alternative models. We study an application to control theory in the following subsection.

5.3.1. Linear-quadratic stochastic control

Here, equation (5.26) will be used to study robustness for linear-quadratic stochastic control (robustness under nonlinear perturbations). Specifically, suppose one is interested in controlling some nonlinear system,

$$d\tilde{X}_t = (B\tilde{X}_t + \sigma\beta(t, \tilde{X}_t) + Du_t)dt + \sigma dW_t, \quad \tilde{X}_0 \sim N(0, \Sigma_0), \quad (5.27)$$

where B, D, σ are $n \times n$ matrices, and W is a \mathbb{R}^n -valued Wiener process. We write the non-linear term with an explicit factor of σ to simplify the use of Girsanov's theorem later on. The control variable is denoted by u_t and \tilde{X}_t is the state; we take \tilde{X}_t to be an observable quantity that can be used for feedback.

The perturbation β may be unknown and even when it is known, optimal control of equation (5.27) is a difficult problem, both analytically and numerically [37]. Therefore, one option is to consider the linear approximation

$$dX_t = (BX_t + Du_t)dt + \sigma dW_t, \quad X_0 \sim N(0, \Sigma_0), \quad (5.28)$$

obtain an explicit formula for the optimal feedback control (under a cost functional to be specified below) for equation (5.28), and use that same feedback function to (suboptimally) control the original system equation (5.27). To evaluate the performance, one must bound the cost functional when the control for the linearized system is used on the nonlinear system.

To make the above precise, we must first specify the cost functional, which we take to be an exponentially discounted quadratic cost:

$$C \equiv E \left[\int_0^\infty \frac{1}{2} (X_s^T Q X_s + u_s^T R u_s) \lambda e^{-\lambda s} ds \right], \quad (5.29)$$

(and similarly with X_t replaced by \tilde{X}_t for the nonlinear system). Here, Q and R are positive-definite $n \times n$ -matrices.

The optimal feedback control for the the linear system equation (5.28) with cost equation (5.29) is given by $u_t = K_\lambda X_t$ where K_λ is obtained as follows (see [3, 9] for details): Define $B_\lambda \equiv B - \frac{\lambda}{2}I$ and solve the following algebraic Riccati equation for Y_λ :

$$B_\lambda^T Y_\lambda + Y_\lambda B_\lambda + Q - Y_\lambda D R^{-1} D^T Y_\lambda = 0. \quad (5.30)$$

The optimal control is then

$$u_t = K_\lambda X_t, \quad K_\lambda \equiv -R^{-1} D^T Y_\lambda. \quad (5.31)$$

Having obtained equation (5.31), we can finally fit the above problem into our UQ framework:

$$\textbf{Baseline Model:} \quad dX_t = A_\lambda X_t dt + \sigma dW_t, \quad A_\lambda \equiv B + D K_\lambda, \quad X_0 \sim N(0, \Sigma_0), \quad (5.32)$$

$$\textbf{Alternative Models:} \quad d\tilde{X}_t = (A_\lambda \tilde{X}_t + \sigma \beta(t, \tilde{X}_t)) dt + \sigma dW_t, \quad \tilde{X}_0 \sim N(0, \Sigma_0), \quad (5.33)$$

$$\textbf{QoI:} \quad F \equiv \int_0^\infty \frac{1}{2} x_s^T C_\lambda x_s \lambda e^{-\lambda s} ds, \quad C_\lambda \equiv Q + K_\lambda^T R K_\lambda, \quad (5.34)$$

where $x \in C([0, \infty), \mathbb{R}^n)$ is the path of the state variable.

Note that, by Girsanov's theorem, the collection of alternative models contains all bounded-drift-perturbations of the baseline model with $\|\beta\|_\infty \leq \alpha$.

The UQ bounds (5.26) applied to the above problem are bounds on the cost of controlling the nonlinear system (5.27) by using the feedback control function that was derived to optimally control the linear system (5.28). We can make the bounds explicit as follows: equation (5.32) has the solution

$$X_t = e^{tA_\lambda} \left(X_0 + \int_0^t e^{-sA_\lambda} \sigma dW_s \right). \quad (5.35)$$

In particular, X_t is Gaussian with mean 0 and covariance

$$\Sigma_t^{\alpha\beta} = (e^{tA_\lambda})_i^\alpha (e^{tA_\lambda})_j^\beta \Sigma_0^{ij} + (e^{tA_\lambda})_i^\alpha (e^{tA_\lambda})_j^\beta \int_0^t (e^{-sA_\lambda} \sigma)_k^i \delta^{k\ell} (e^{-sA_\lambda} \sigma)_\ell^j ds. \quad (5.36)$$

Next, compute a Cholesky factorization of C_λ , $C_\lambda = M_\lambda^T M_\lambda$. The integrand in equation (5.29) is non-negative, hence the hypotheses of Theorem 4.6 are met and we can use equation (5.26) to obtain

$$\begin{aligned} & - \int_{\mathcal{T}} \inf_{c>0} \left\{ -\frac{1}{2c} \log(|\det(I + c\Sigma_s C_\lambda)|) + \frac{1}{2c} \alpha^2 s \right\} \pi(ds) \\ & \leq E_{\tilde{P}}[F] \leq \int_{\mathcal{T}} \inf_{0<c<1/\|M_\lambda \Sigma_s M_\lambda^T\|} \left\{ -\frac{1}{2c} \log(|\det(I - c\Sigma_s C_\lambda)|) + \frac{1}{2c} \alpha^2 s \right\} \pi(ds). \end{aligned} \quad (5.37)$$




FIGURE 3. Bound on the expected quadratic cost, equation (5.34), for a nonlinear stochastic control system (alternative model) equation (5.27), (suboptimally) controlled *via* the optimal control for an approximating linear system (baseline model) equation (5.28). The blue dashed lines, computed from equation (5.37), constrain the QoI to the gray region. The parameter κ (defined in equation (5.38)) governs the strength of the controller.

Figure 3 shows numerical results, corresponding to the following simple example system, with a 2-D state variable and a 1-D control variable:

$$B = \begin{bmatrix} 2 & 0.1 \\ 0.1 & -1 \end{bmatrix}, \quad D = \begin{bmatrix} \kappa \\ 0 \end{bmatrix}, \quad (5.38)$$

$Q = I$, $R = 1$, $\Sigma_0 = 0$, $\lambda = 1/2$, $\sigma = I$, $\alpha = 1/2$. Note that the size of the perturbation, $\sigma\beta$, is not required to be “small” for the method to produce non-trivial bounds, as this is a non-perturbative calculation. Finally, if one instead attempts to use the non-goal-oriented relative entropy to derive UQ bounds, the results are again trivial (infinite), due to the unbounded time-horizon.

5.4. Semi-Markov perturbations of a $M/M/\infty$ queue

Continuous-time jump processes with non-exponential waiting times (*i.e.*, waiting times with memory) are found in many real-world systems; see, for example, the telephone call-center waiting time distribution in Figure 2 of [16]. However, these systems are often approximated by a much simpler Markovian (*i.e.*, exponentially-distributed waiting time) model. Here we derive robustness bounds for such approximations.

More specifically, we take the baseline model to be a $M/M/\infty$ queue (*i.e.*, continuous-time birth-death process) with arrival rate $\alpha > 0$ and service rate $\rho > 0$:

Baseline Model: $P(X_{t+dt} = x + 1 | X_t = x) = \alpha dt$, $P(X_{t+dt} = x - 1 | X_t = x) = \rho x dt$, $x \in \mathbb{Z}_0$. (5.39)

The embedded jump Markov-chain has transition probabilities, $a(x, y)$, given by

$$a(x, x + 1) = \alpha / (\alpha + \rho x), \quad a(x, x - 1) = \rho x / (\alpha + \rho x), \quad (5.40)$$

and the waiting-times are exponentially distributed with jump rates

$$\lambda(x) = \alpha + \rho x. \quad (5.41)$$

We consider the stationary case, *i.e.*, X_0 is Poisson with parameter α/ρ . Also, note that the jump chain with transition matrix (5.40) has the stationary distribution

$$\pi(x) = \frac{(\alpha + \rho x)(\alpha/\rho)^x}{2\alpha x!} e^{-\alpha/\rho}. \quad (5.42)$$

The alternative models will be semi-Markov processes; these describe jump-processes with non-exponential waiting times (*i.e.*, with memory). Mathematically, a semi-Markov model is a piecewise-constant continuous-time process defined by a jump chain, X_n^J , and jump times, J_n , and waiting times (*i.e.*, jump intervals), $\Delta_{n+1} \equiv J_{n+1} - J_n$, that satisfy

$$\begin{aligned} \text{Alternative Models:} \quad \tilde{P}(X_{n+1}^J = y, \Delta_{n+1} \leq t | X_1^J, \dots, X_{n-1}^J, X_n^J, J_1, \dots, J_n) \\ = \tilde{P}(X_{n+1}^J = y, \Delta_{n+1} \leq t | X_n^J) \equiv \tilde{Q}_{X_n^J, y}(t). \end{aligned} \quad (5.43)$$

$\tilde{Q}_{x,y}(t)$ is called the semi-Markov kernel (see [33, 41] for further background and applications). Note that the baseline process (5.39) is described by the semi-Markov kernel

$$Q_{x,y}(t) \equiv \int_0^t q_{x,y}(s) ds = a(x, y) \int_0^t \lambda(x) e^{-\lambda(x)s} ds. \quad (5.44)$$

We emphasize that the alternative models (5.43) are not Markov processes in general; our methods make no such assumption on \tilde{P} .

An important QoI in such a model is the average queue-length,

$$\text{QoI:} \quad F_T \equiv \frac{1}{T} \int_0^T x_t dt. \quad (5.45)$$

More specifically, we will be concerned with the long-time behavior (limit of $E_{\tilde{P}}[F_T]$ as $T \rightarrow \infty$). Time-averages such as equation (5.45) were discussed in Section 4.4, and for these the UQ bound (4.2) implies

$$\pm (E_{\tilde{P}}[F_T] - E_P[F_T]) \leq \inf_{c>0} \left\{ \frac{1}{cT} \Lambda_{\tilde{P}}^{\hat{F}_T}(\pm cT) + \frac{1}{cT} R(\tilde{P}|_{\mathcal{F}_T} \| P|_{\mathcal{F}_T}) \right\}. \quad (5.46)$$

$E_P[F_T] = \alpha/\rho$ is the average queue-length in the (stationary) baseline process. Note that one does not need to first prove $F_T \in L^1(\tilde{P})$ to obtain equation (5.46); non-negativity of the QoI allows one to first work with a truncated QoI and then take limits.

The cumulant generating function of the time-averaged queue-length has the following limit (see Sect. 4.3.4 in [10]):

$$\lim_{T \rightarrow \infty} T^{-1} \Lambda_{\tilde{P}}^{\hat{F}_T}(cT) = \frac{\alpha c^2}{\rho^2(1 - c/\rho)}, \quad c < \rho. \quad (5.47)$$

Combining equations (5.46) and (5.47), we obtain the following UQ bound:

Proposition 5.4. *Let F_T denote the average queue-length (5.45) and \tilde{P} be a semi-Markov perturbation (see Eqs. (5.43) and (5.44)). Then*

$$\begin{aligned} \limsup_{T \rightarrow \infty} (E_{\tilde{P}}[F_T] - \alpha/\rho) &\leq \inf_{0 < c < \rho} \left\{ \frac{1}{c} \frac{\alpha c^2}{\rho^2(1 - c/\rho)} + \frac{1}{c} H(\tilde{P}\|P) \right\} \\ &= \left(2\sqrt{H(\tilde{P}\|P)/\alpha} + H(\tilde{P}\|P)/\alpha \right) \frac{\alpha}{\rho}, \\ \liminf_{T \rightarrow \infty} (E_{\tilde{P}}[F_T] - \alpha/\rho) &\geq - \inf_{c > 0} \left\{ \frac{1}{c} \frac{\alpha c^2}{\rho^2(1 + c/\rho)} + \frac{1}{c} H(\tilde{P}\|P) \right\} \\ &= \begin{cases} - \left(2\sqrt{H(\tilde{P}\|P)/\alpha} - H(\tilde{P}\|P)/\alpha \right) \frac{\alpha}{\rho} & H(\tilde{P}\|P) < \alpha \\ -\frac{\alpha}{\rho} & H(\tilde{P}\|P) \geq \alpha, \end{cases} \end{aligned} \quad (5.48)$$

where $H(\tilde{P}\|P)$ is the relative entropy rate:

$$H(\tilde{P}\|P) \equiv \limsup_{T \rightarrow \infty} \frac{1}{T} R(\tilde{P}|_{\mathcal{F}_T} \| P|_{\mathcal{F}_T}). \quad (5.49)$$

A formula for the relative entropy rate between semi-Markov processes was obtained in [25] under the appropriate ergodicity assumptions:

$$\begin{aligned} H(\tilde{P}\|P) &= \frac{1}{\tilde{m}_{\tilde{\pi}}} \sum_{x,y} \tilde{\pi}(x) \int_0^\infty \tilde{q}_{x,y}(t) \log(\tilde{q}_{x,y}(t)/q_{x,y}(t)) dt, \\ \tilde{m}_{\tilde{\pi}} &\equiv \sum_x \tilde{\pi}(x) \int_0^\infty (1 - \tilde{H}_x(t)) dt, \quad \tilde{H}_x(t) = \sum_y \tilde{Q}_{x,y}(t), \end{aligned} \quad (5.50)$$

where $\tilde{\pi}$ is the invariant distribution for the embedded jump chain of \tilde{P} .

In particular, an alternative process with the same jump-chain as the base process but a different waiting-time distribution, $\tilde{H}_x(t) = \int_0^t \tilde{h}_x(s) ds$, is described by the semi-Markov kernel

$$\tilde{Q}_{x,y}(t) = a(x,y) \tilde{H}_x(t). \quad (5.51)$$

In this case, $\tilde{\pi} = \pi$ and $H(\tilde{P}\|P)$ can be expressed in terms of the relative entropy of the waiting-time distributions:

$$H(\tilde{P}\|P) = \limsup_{T \rightarrow \infty} \frac{1}{T} R(\tilde{P}|_{\mathcal{F}_T} \| P|_{\mathcal{F}_T}) = \frac{1}{\tilde{m}_{\tilde{\pi}}} E_{\pi(\text{d}x)} \left[R(\tilde{H}_x \| H_x) \right], \quad (5.52)$$

$$\tilde{m}_{\tilde{\pi}} = \sum_x \pi(x) \int_0^\infty (1 - \tilde{H}_x(t)) dt, \quad (5.53)$$

where π is the invariant distribution for $a(x,y)$; recall the stationary distribution for the jump chain of the $M/M/\infty$ -queue was given in equation (5.42).

Remark 5.5. The quantity $\tilde{m}_{\tilde{\pi}}$ is the mean sojourn time under the invariant distribution, π , and the expectation $E_{\pi(\text{d}x)} \left[R(\tilde{H}_x \| H_x) \right]$ can be viewed as the mean relative entropy of a single jump (comparing the alternative and baseline model waiting-time distributions). The formula for $H(\tilde{P}\|P)$, equation (5.52), therefore has the clear intuitive meaning of an information loss per unit time.

Other than \tilde{H}_x , the formulas (5.48) and (5.52) include only baseline model quantities. Next we show one can bound \tilde{H}_x for a particular class of semi-Markov processes, and thus obtain computable UQ bounds.

5.4.1. UQ bounds for phase type distributions

Phase-type distributions constitute a semi-parametric description of waiting-time distributions, going beyond the exponential case to describe systems with memory; see [5, 16, 23] for examples and information on fitting such distributions to data. Probabilistically, they can be characterized in terms the time to absorption for a continuous-time Markov chain with a single absorbing state. The density and distribution function of a phase type-distribution, $\text{PH}_k(\nu, T)$, are characterized by a $k \times k$ matrix T , and a probability vector $\nu \in \mathbb{R}^k$ (see [11] for background):

$$\tilde{h}(t) = \nu e^{Tt}(-T\mathbf{1}), \quad \tilde{H}(t) \equiv \int_0^t \tilde{h}(s) ds = \mathbf{1} - \nu e^{Tt}\mathbf{1}, \quad (5.54)$$

where $\mathbf{1}$ is the vector of all 1's and T satisfies:

- (1) $T_{j,j} < 0$ for all j ,
- (2) $T_{j,\ell} \geq 0$ for all $j \neq \ell$,
- (3) $-T\mathbf{1}$ has all non-negative entries.

Combining equation (5.54) with equation (5.52), we obtain a formula for the relative entropy rate of a semi-Markov perturbation:

Proposition 5.6. *Let the baseline model consist of a $M/M/\infty$ queue (5.39) and let \tilde{P} be a semi-Markov perturbation with state- x waiting-time distributed as $\text{PH}_{k(x)}(\nu(x), T(x))$ (and with the same jump-chain, $a(x, y)$, as the base process). Then the relative entropy rate is*

$$\begin{aligned} H(\tilde{P}\|P) &= \frac{1}{\tilde{m}_\pi} \sum_x \pi(x) \int_0^\infty \log \left(\lambda(x)^{-1} \nu(x) e^{(T(x)+\lambda(x))t} (-T(x)\mathbf{1}) \right) \nu(x) e^{T(x)t} (-T(x)\mathbf{1}) dt, \\ \tilde{m}_\pi &= \sum_x \pi(x) (-\nu(x) T^{-1}(x) \mathbf{1}). \end{aligned} \quad (5.55)$$

As a simple example, consider the case where the waiting-time at x is distributed as a convolution of $\exp(\lambda(x))$ and $\exp(\gamma(x))$ -distributions (convolutions of exponential distributions are examples of phase-type distributions; again, see [11]), for some choices of $\gamma(x) > \lambda(x)$. Recall the baseline model rates were given in equation (5.41). The semi-Markov kernel for the alternative process is then

$$\tilde{Q}_{x,y}(t) = a(x, y) \int_0^t \lambda(x) e^{-\lambda(x)s} \frac{1 - e^{-(\gamma(x)-\lambda(x))s}}{1 - \lambda(x)/\gamma(x)} ds. \quad (5.56)$$

Note that the Markovian limit occurs at $\gamma(x) = \infty$, where the integrand is non-analytic. Hence, it appears unlikely that one could perform a perturbative study of this system. In contrast, our method will produce computable UQ bounds.

To proceed, one must assume some partial knowledge of $\gamma(x)$. Here we consider the alternative family consisting of processes of the form equation (5.56) with the following parameter bounds (of course, other choices are possible):

Alternative Models: $0 \leq \delta \leq \lambda(x)/\gamma(x) \leq \epsilon < 1$ for all x , *i.e.*, $\frac{1}{\epsilon} \lambda(x) \leq \gamma(x) \leq \frac{1}{\delta} \lambda(x)$, (5.57)

where the jump-rates for the baseline model are $\lambda(x) = \alpha + \rho x$. As a function of x , equation (5.57) therefore constrains $\gamma(x)$ to lie between two lines. This set describes waiting-times that are approximately $\exp(\lambda(x))$ -distributed when ϵ is small (*i.e.*, when $\gamma(x) \gg \lambda(x)$). The parameter δ is a lower-bound on the perturbation. See the left panel of Figure 4 for an illustrative example of this class of alternative models.




FIGURE 4. *Left*: waiting time densities for a baseline (exponential waiting time) model and a semi-Markov perturbation of the form equation (5.57), corresponding to $\delta = \epsilon = 0.05$. *Right*: upper bound on the relative error for the expected time-averaged-queue-length in the long-time limit, equation (5.60), for a semi-Markov perturbation of a $M/M/\infty$ queue. The parameters δ and ϵ quantify the size of the perturbation; see equation (5.57) for their definitions. As $\epsilon \rightarrow 0$, the semi-Markov waiting times approach exponential waiting times and the relative error approaches 0. Increasing δ reduces the size of the model uncertainty; this is also reflected in the upper bound.

Using equation (5.57) we can bound

$$\begin{aligned}
 H(\tilde{P} \| P) &= \frac{1}{\tilde{m}_\pi} \sum_x \pi(x) \int_0^\infty \lambda(x) e^{-\lambda(x)t} \frac{1 - e^{-(\gamma(x) - \lambda(x))t}}{1 - \lambda(x)/\gamma(x)} \log \left(\frac{1 - e^{-(\gamma(x) - \lambda(x))t}}{1 - \lambda(x)/\gamma(x)} \right) dt \\
 &\leq \alpha r(\delta, \epsilon), \\
 r(\delta, \epsilon) &\equiv \frac{2}{1 + \delta} \int_0^\infty e^{-u} \left(\frac{1 - e^{-(\epsilon^{-1} - 1)u}}{1 - \delta} \mathbf{1}_{u < \frac{\log(1/\epsilon)}{\delta - 1 - 1}} + \frac{1 - e^{-(\delta^{-1} - 1)u}}{1 - \epsilon} \mathbf{1}_{u \geq \frac{\log(1/\epsilon)}{\delta - 1 - 1}} \right) \\
 &\quad \times \log \left(\frac{1 - e^{-(\delta^{-1} - 1)u}}{1 - \epsilon} \right) du.
 \end{aligned} \tag{5.58}$$

To obtain the above, we used the inequality

$$\tilde{m}_\pi = \sum_x \pi(x) \lambda(x)^{-1} (1 + \lambda(x)/\gamma(x)) \geq (1 + \delta) \sum_x \pi(x) \lambda(x)^{-1} = \frac{1 + \delta}{2\alpha}. \tag{5.59}$$

Combining equation (5.58) and equation (5.59) with equation (5.48) gives the relative-error bounds

$$\begin{aligned}
 \limsup_{T \rightarrow \infty} \left(E_{\tilde{P}} [F_T] / \left(\frac{\alpha}{\rho} \right) - 1 \right) &\leq 2\sqrt{r(\delta, \epsilon)} + r(\delta, \epsilon), \\
 \liminf_{T \rightarrow \infty} \left(E_{\tilde{P}} [F_T] / \left(\frac{\alpha}{\rho} \right) - 1 \right) &\geq - \left(2\sqrt{r(\delta, \epsilon)} - r(\delta, \epsilon) \right) \mathbf{1}_{r(\delta, \epsilon) < 1} - \mathbf{1}_{r(\delta, \epsilon) \geq 1}.
 \end{aligned} \tag{5.60}$$

These bounds depend only on the uncertainty parameters ϵ and δ , and not on the base-model parameters. See the right pane in Figure 4 for a contour plot of the logarithm of the upper bound equation (5.60). Note that

the error decreases as δ increases, due to the decreasing uncertainty in the form of the perturbation γ (see Eq. (5.57)). As $\epsilon \rightarrow 0$, the semi-Markov perturbation approaches the baseline model (see Eqs. (5.56) and (5.57)) and so the relative error approaches 0. As previously noted, relative entropy based UQ methods are sub-optimal for risk-sensitive quantities (*i.e.*, rare events); see [7] for a treatment of risk-sensitive UQ for queueing models.

5.5. Value of American put options with variable interest rate

As a final example, we consider the value of an asset in a variable-interest-rate environment, as compared to a constant-interest-rate baseline model. For background see, for example, Chapter 8 of [52] or Chapter 8 of [26]. This example does not fit neatly into any of the problem categories from Section 4, though the technique of Corollary 3.9 will play a critical role. In particular, we will see that the natural QoI for the base and alternative models is different; see equations (5.66) and (5.68).

Specifically, the baseline model in this example is geometric Brownian motion:

Baseline Model:
$$dX_t^r = rX_t^r dt + \sigma X_t^r dW_t, \quad (5.61)$$

where $r, \sigma > 0$ are constants. This has the explicit solution

$$X_t^r = X_0^r \exp(\sigma W(t) + (r - \sigma^2/2)t). \quad (5.62)$$

The variable-interest-rate alternative model has the general form

Alternative Models:
$$d\tilde{X}_t = (r + \Delta r(t, \tilde{X}_t, \tilde{Y}_t))\tilde{X}_t dt + \sigma \tilde{X}_t dW_t. \quad (5.63)$$

The initial values $X_0^r = \tilde{X}_0 \equiv X_0$ will be fixed (here, the superscript on X^r will always denote the constant interest rate, r , that is being used). r will be thought of as a known, fixed value, while Δr will be considered an unknown perturbation that depends on time, on the asset price, as well as on an additional driving process, \tilde{Y}_t , that solves some auxiliary SDE; see Appendix D.3 for more precise assumptions and [52] for discussion of various interest rate models. In Section 5.5.1 below, we will consider the Vasicek model.

The quantity of interest we consider is the value of a perpetual American put option: Let $K > 0$ be the option strike price. The payout if the option is exercised at time $t \geq 0$ is $K - x_t$, where x_t is the asset price at time t . The relevant QoI is then the value, discounted to the present time. In the baseline model, this takes the form $V : [0, \infty] \times C([0, \infty), \mathbb{R}) \rightarrow \mathbb{R}$,

$$V_t[x] = e^{-rt}(K - x_t)\mathbf{1}_{t < \infty}, \quad (5.64)$$

while for the alternative model it is $\tilde{V} : [0, \infty] \times C([0, \infty), \mathbb{R} \times \mathbb{R}^k) \rightarrow \mathbb{R}$,

$$\tilde{V}_t[x, y] = e^{-\int_0^t r + \Delta r(s, x_s, y_s) ds}(K - x_t)\mathbf{1}_{t < \infty}. \quad (5.65)$$

The option-holder's strategy is to exercise the option when the asset price hits some level L , assumed to satisfy $L < X_0$ and $0 < L < K$, *i.e.*, we consider the stopping time $\tau[x, y] = \tau[x] = \inf\{t \geq 0 : x_t \leq L\}$. Therefore the baseline model QoI is

Baseline QoI:
$$E[V_\tau[X^r]] = (K - L)(L/X_0)^{2r/\sigma^2}. \quad (5.66)$$

To evaluate this, one uses the fact that

$$\tau \circ X^r = \inf\{t : X_t^r = L\} = \inf\{t : W_t + (r/\sigma - \sigma/2)t = -\sigma^{-1} \log(X_0/L)\} \quad (5.67)$$

is the level $a \equiv -\sigma^{-1} \log(X_0/L)$ hitting time of a Brownian motion with constant drift $\mu \equiv r/\sigma - \sigma/2$, combined with the formula for the distribution of such hitting times (see [34] and also Appendix F).

The goal of our analysis is now to bound the expected option value in the alternative model:

$$\textbf{Alternative QoI:} \quad E \left[\tilde{V}_{\tau \circ \tilde{X}}[\tilde{X}, \tilde{Y}] \right] = (K - L) E \left[\mathbf{1}_{\tau \circ \tilde{X} < \infty} e^{-\int_0^{\tau \circ \tilde{X}} r + \Delta r(s, \tilde{X}_s, \tilde{Y}_s) ds} \right]. \quad (5.68)$$

Note that both QoIs have unbounded time-horizon; therefore, to obtain non-trivial bounds one must again use a goal-oriented relative entropy.

Remark 5.7. The methods developed in this paper, applied to the goal oriented relative entropy $R(\tilde{P}|_{\mathcal{F}_\tau} \| P|_{\mathcal{F}_\tau})$, are still not capable of comparing the baseline model equation (5.61), with asset-price volatility σ , to an alternative model with perturbed asset-price volatility, $\sigma + \Delta\sigma$, due to the loss of absolute continuity (again, see page 80 in [24]). It is likely that one does have absolute continuity on a smaller sigma algebra than \mathcal{F}_τ , but we do not currently know how to bound the relative entropy on any such smaller sigma algebra. An alternative approach to robust option pricing under uncertainty in σ , utilizing H_∞ -control methods, was developed in [42].

To obtain UQ bounds for equation (5.68), it is useful to define the modified QoI

$$F_t[x, y] = (K - L) \exp \left(- \int_0^t r + \Delta r(s, x_s, y_s) ds \right) \mathbf{1}_{t < \infty}, \quad (5.69)$$

and note that $E \left[\tilde{V}_{\tau \circ \tilde{X}}[\tilde{X}, \tilde{Y}] \right] = E_{\tilde{X}, \tilde{Y}}[F_\tau]$, where $E_{\tilde{X}, \tilde{Y}}$ denotes the expectation with respect to the joint distribution of (\tilde{X}, \tilde{Y}) on path-space. (The notation $E_{X, Y}$ will similarly be used below.) In the next subsection, we will bound $E_{\tilde{X}, \tilde{Y}}[F_\tau]$ for one class of rate perturbations.

5.5.1. Vasicek interest rate model

Here we study a specific type of dynamical interest rate model, known as the Vasicek model (see, *e.g.*, page 150 in [52]):

$$\textbf{Alternative Models:} \quad d\tilde{X}_t = (r + \Delta r_t)\tilde{X}_t dt + \sigma \tilde{X}_t dW_t, \quad (5.70)$$

$$d\Delta r_t = -\gamma \Delta r_t dt + \tilde{\sigma} d\tilde{W}_t, \quad \Delta r_0 = 0, \quad \gamma > 0, \quad \tilde{\sigma} > 0, \quad (5.71)$$

where W and \tilde{W} are independent Brownian motions. The baseline model is still given by equation (5.61). Robustness bounds under such perturbations can be viewed as a model stress test for the type of financial instrument (QoI) studied here; see [22] for a detailed discussion of stress testing.

The SDE (5.71) defines a 2-dimensional parametric family of alternative models, parameterized by $\gamma > 0$ and $\tilde{\sigma} > 0$. Infinite-dimensional alternative families of Δr -models can also be considered; see Appendix I for one such class of examples. Equation (5.71) has the exact solution

$$\Delta r_t = \tilde{\sigma} e^{-\gamma t} \int_0^t e^{\gamma s} d\tilde{W}_s. \quad (5.72)$$

Note that Δr_t is unbounded, and so a comparison-principle bound is not possible here. Under the model (5.70), $r + \Delta r$ can become negative. Here, for both modeling and mathematical reasons, we eliminate this possibility by conditioning on the event $\int_0^{\tau \circ \tilde{X}} r + \Delta r_s ds \geq 0$, *i.e.*, we only consider those paths with non-negative average interest rate up to the stopping time. This amounts to using the modified QoI

$$\tilde{F}_t[x, \Delta r] \equiv (K - L) \exp \left(- \int_0^t r + \Delta r_s ds \right) \mathbf{1}_{t < \infty} \mathbf{1}_{\int_0^t r + \Delta r_s ds \geq 0} \quad (5.73)$$

and bounding

$$\textbf{Alternative QoI:} \quad \tilde{E}[F_\tau] \equiv E_{\tilde{X}, \Delta r}[\tilde{F}_\tau] / P_{\tilde{X}, \Delta r} \left(\int_0^\tau r + \Delta r_s ds \geq 0 \right). \quad (5.74)$$

Other more complex interest rate models exist that enforce a positive rate *via* the dynamics; see, for example, page 275 in [52].

We also restrict to the parameter values $r > \frac{\tilde{\sigma}^2}{2\gamma^2}$, for which a negative average interest rate is sufficiently rare. More precisely, this assumption implies $\lim_{n \rightarrow \infty} \int_0^n r + \Delta r_s ds = \infty$ a.s.; see Appendix H for details. As a consequence, we also obtain $\tilde{F}_n \rightarrow \tilde{F}_\infty = 0$ a.s. (under both the base and alternative models). Therefore, Corollary 4.3 and Lemma C.4 (see Appendix C) can both be used, yielding

$$\pm E_{\tilde{X}, \Delta r}[\tilde{F}_\tau] \leq \inf_{c > 0} \left\{ \frac{1}{c} \Lambda_{P_{X^r, \Delta r}}^{\tilde{F}_\tau}(\pm c) + \frac{1}{c} \liminf_{n \rightarrow \infty} R(P_{\tilde{X}, \Delta r} | \mathcal{F}_{\tau \wedge n} \| P_{X^r, \Delta r} | \mathcal{F}_{\tau \wedge n}) \right\}. \quad (5.75)$$

For this example, we do not utilize Theorem 3.6, as the resulting expectation is difficult to evaluate. Instead, we work with the cumulant generating function and relative entropy terms separately and use Corollary 3.9 together with Theorem 3.1.

First we compute the cumulant generating function. $\tau \circ X^r$ and Δr are independent, hence we can use the result of [1] to find

$$\begin{aligned} \Lambda_{P_{X^r, \Delta r}}^{\tilde{F}_\tau}(\pm c) &= \log \int \int_{-\infty}^{\infty} (2\pi\tilde{\sigma}_t^2)^{-1/2} \exp(\pm c(K - L)e^{-z} 1_{t < \infty} 1_{z \geq 0}) e^{-\frac{(z - rt)^2}{2\tilde{\sigma}_t^2}} dz P_{\tau \circ X^r}(dt), \\ \tilde{\sigma}_t^2 &\equiv \frac{\tilde{\sigma}^2}{2\gamma^3} (2\gamma t + 4e^{-\gamma t} - e^{-2\gamma t} - 3). \end{aligned} \quad (5.76)$$

The integral over t in equation (5.76) can then be computed as discussed in the text surrounding equation (5.67), with $a \equiv -\sigma^{-1} \log(X_0/L)$ and $\mu \equiv r/\sigma - \sigma/2$. (When $t = \infty$, the inner integral should be interpreted as equaling 1.)

Remark 5.8. Note that if the integrated rate (z in Eq. (5.76)) is allowed to be negative (*i.e.*, we remove $1_{z \geq 0}$ from the exponential) then $\Lambda_{P_{X^r, \Delta r}}^{\tilde{F}_\tau}(c)$ is infinite and the upper UQ bound is the trivial bound, $+\infty$. This is the mathematical reason for conditioning on the event $\int_0^\tau r + \Delta r_s ds \geq 0$.

The relative entropy can be computed *via* Girsanov's theorem (Appendix H for a proof that the Girsanov factor is an exponential martingale)

$$R(P_{\tilde{X}, \Delta r} | \mathcal{F}_{\tau \wedge n} \| P_{X^r, \Delta r} | \mathcal{F}_{\tau \wedge n}) \leq E_{\tilde{X}, \Delta r}[G_n], \quad G_n[x, \Delta r] \equiv \frac{1}{2} \int_0^{\tau(x) \wedge n} \sigma^{-2} |\Delta r_s|^2 ds. \quad (5.77)$$

Note that G_n is $\mathcal{F}_{\tau \wedge n}$ -measurable and has finite expectation, as one can bound $\tau \wedge n$ by n and then use the fact that Δr_s is normal with mean 0 and variance $\frac{\tilde{\sigma}^2}{2\gamma}(1 - e^{-2\gamma s})$. Hence, Corollary 3.9 is applicable, resulting in

$$R(P_{\tilde{X}, \Delta r} | \mathcal{F}_{\tau \wedge n} \| P_{X^r, \Delta r} | \mathcal{F}_{\tau \wedge n}) \leq \inf_{\lambda > 1} \left\{ (\lambda - 1)^{-1} \Lambda_{P_{X^r, \Delta r}}^{G_n}(\lambda) \right\}. \quad (5.78)$$

Again using the independence of $\tau \circ X^r$ and Δr , we have

$$\Lambda_{P_{X^r, \Delta r}}^{G_n}(\lambda) = \log \int E \left[e^{\lambda \frac{\sigma^{-2}}{2} \int_0^{t \wedge n} |\Delta r_s|^2 ds} \right] P_{\tau \circ X^r}(dt). \quad (5.79)$$

Therefore

$$\liminf_{n \rightarrow \infty} R(P_{\tilde{X}, \Delta r} | \mathcal{F}_{\tau \wedge n} \| P_{X^r, \Delta r} | \mathcal{F}_{\tau \wedge n}) \leq \inf_{\lambda > 1} \left\{ (\lambda - 1)^{-1} \log \int E \left[e^{\lambda \frac{\sigma^{-2}}{2} \int_0^t |\Delta r_s|^2 ds} \right] P_{\tau \circ X^r}(dt) \right\} \quad (5.80)$$




FIGURE 5. Bound on the expected option value under the Vasicek model. Black curves are for the baseline model (see Eq. (5.66)) and blue dashed curves are the bounds on the alternative model QoI, equation (5.74), that result from our method (see Eqs. (5.75)–(5.83)); the QoI is constrained to the gray regions. *Left*: as a function of the volatility of the interest-rate perturbation, with $\gamma = 2$, $r = 5/4$, $K = 1$, $L = 1/2$, $X_0 = 2$, $\sigma = 4$. *Right*: as a function of the volatility of the asset price, with $\gamma = 2$, $r = 1$, $K = 1$, $L = 1/2$, $X_0 = 2$, $\tilde{\sigma} = 6$.

and the inner expectation can be evaluated using Theorem 1(B) of [18]:

$$E \left[e^{\lambda \frac{\sigma^2}{2} \int_0^t |\Delta r_s|^2 ds} \right] = e^{\gamma t/2} \left[\frac{1}{\sqrt{1 - \frac{\lambda \tilde{\sigma}^2}{\sigma^2 \gamma^2}}} \sinh \left(\gamma t \sqrt{1 - \frac{\lambda \tilde{\sigma}^2}{\sigma^2 \gamma^2}} \right) + \cosh \left(\gamma t \sqrt{1 - \frac{\lambda \tilde{\sigma}^2}{\sigma^2 \gamma^2}} \right) \right]^{-1/2}. \quad (5.81)$$

By analytic continuation, this formula is valid for $\lambda < \sigma^2 \gamma^2 / \tilde{\sigma}^2$. Also, note that for the $n \rightarrow \infty$ limit of equation (5.79) to be finite, we need $\tau \circ X^r < \infty$ P -a.s., *i.e.*, μ and a must have the same signs.

We can use Corollary 3.4 to bound $P_{\tilde{X}, \Delta r} \left(\int_0^\tau r + y_s ds \geq 0 \right)$:

$$\begin{aligned} \pm P_{\tilde{X}, \Delta r} \left(\int_0^\tau r + y_s ds \geq 0 \right) &\leq \inf_{c > 0} \left\{ \frac{1}{c} \log \left(1 + (e^{\pm c} - 1) P_{X^r, \Delta r} \left(\int_0^\tau r + \Delta r_s ds \geq 0 \right) \right) \right. \\ &\quad \left. + \frac{1}{c} \liminf_{n \rightarrow \infty} R(P_{\tilde{X}, \Delta r} |_{\mathcal{F}_{\tau \wedge n}} \| P_{X^r, \Delta r} |_{\mathcal{F}_{\tau \wedge n}}) \right\}, \end{aligned} \quad (5.82)$$

where we computed the MGF as in equation (3.8) and

$$\begin{aligned} P_{X^r, \Delta r} \left(\int_0^\tau r + \Delta r_s ds \geq 0 \right) &= \int \int_0^\infty (2\pi \tilde{\sigma}_t^2)^{-1/2} e^{-\frac{(z - rt)^2}{2\tilde{\sigma}_t^2}} dz P_{\tau \circ X^r}(dt) \\ &= \int \frac{1}{2} \left(\operatorname{erf} \left(rt / \sqrt{2\tilde{\sigma}_t^2} \right) + 1 \right) P_{\tau \circ X^r}(dt). \end{aligned} \quad (5.83)$$

Figure 5 shows numerical results for the bounds on equation (5.74) that result from combining equation (5.75) through equation (5.83) (the full statement of the bound is quite lengthy so, rather than stating it here, we have placed it in Appendix H, Eq. (H.7)). The baseline model is shown in the black solid curves (see Eq. (5.66)) and the blue dashed curves show the bounds (H.7) on the alternative model QoI that result from our method, which constrain the QoI to the gray regions; we emphasize that the alternative model uses the modified QoI, equations (5.73) and (5.74). In the left plot we see that the uncertainty approaches zero as $\tilde{\sigma}$ approaches zero, corresponding to zero rate perturbation. However, $\tilde{\sigma}$ does not need to be small for our method to apply.

Remark 5.9. We are able to perturb the diffusion parameter, $\tilde{\sigma}$, without losing absolute continuity because, technically, the baseline model is the joint distribution of $(X_t^r, \Delta r_t)$. Of course, X^r is not coupled to Δr_t , nor does the baseline QoI depend on it, but regardless, this makes it clear that there is no loss of absolute continuity between the distributions of $(X_t^r, \Delta r_t)$ and $(\tilde{X}_t, \Delta r_t)$ when $\tilde{\sigma}$ is changed.

APPENDIX A. PROOF OF THE UQ INFORMATION INEQUALITY

For completeness, here we provide a proof of the UQ bound (2.4), which can be found in [17, 20]:

Proposition A.1. *Let P, \tilde{P} be probability measures and $F : \Omega \rightarrow \mathbb{R}$, $F \in L^1(\tilde{P})$. Then*

$$\pm E_{\tilde{P}}[F] \leq \inf_{c > 0} \left\{ \frac{1}{c} \Lambda_{\tilde{P}}^F(\pm c) + \frac{1}{c} R(\tilde{P} \| P) \right\}. \quad (\text{A.1})$$

Proof. The key tool is the Gibbs variational principle, which relates the cumulant generating function and relative entropy as follows: If $G : \Omega \rightarrow \mathbb{R}$ is measurable and there exists $c_0 > 0$ such that $E_P[e^{c_0 G}] < \infty$ then

$$\log E_P[e^G] = \sup_{Q: R(Q \| P) < \infty} \{E_Q[G] - R(Q \| P)\}. \quad (\text{A.2})$$

A (slightly weaker) version of this result is found in Proposition 4.5.1 in [19], where it is assumed that G is bounded above. However, the above variant can then easily be obtained by applying equation (A.2) to $G_n \equiv G \wedge n$ and then taking $n \rightarrow \infty$ via the monotone convergence theorem (the assumption $E_P[e^{c_0 G}] < \infty$ is necessary to ensure that $E_Q[G^+] < \infty$ for all Q with $R(Q \| P) < \infty$, where $G^+ \equiv \max\{G, 0\}$ is the positive part of G).

Now fix $c > 0$ and consider the case of the upper signs in equation (A.1) (the proof for the lower signs proceeds similarly). If $R(\tilde{P} \| P) = \infty$ or $\Lambda_{\tilde{P}}^F(c) = \infty$ then the claimed bound (at c) is trivial, so suppose not. In this case, we have $E_P[e^{cF}] < \infty$ and so we can apply equation (A.2) to $G \equiv cF$. Bounding the supremum in equation (A.2) below by the value at $Q = \tilde{P}$, solving for $E_{\tilde{P}}[cF]$, and then dividing by c yields the result. \square

APPENDIX B. TIGHTNESS OF THE BOUNDS FROM THEOREM 3.6

The bounds (3.10) are tight, in the following sense:

Corollary B.1. *Let F and \mathcal{G} be as in Theorem 3.6 and assume that $R(\tilde{P}|_{\mathcal{G}} \| P|_{\mathcal{G}}) < \infty$ and $c \mapsto \log E_{\tilde{P}}[\exp(cF)]$ is finite on a neighborhood of 0. Then there exists a \mathcal{G} -measurable G such that equation (3.10) is an equality.*

Proof. Let $G = \log(d\tilde{P}|_{\mathcal{G}}/dP|_{\mathcal{G}})$. The assumptions, together with equation (3.10), imply $G \in L^1(\tilde{P})$ and we have

$$\begin{aligned} \pm E_{\tilde{P}}[F] &\leq \inf_{c > 0} \left\{ \frac{1}{c} \log E_P[\exp(\pm cF + G)] \right\} = \lim_{c \searrow 0} \frac{1}{c} \log E_{\tilde{P}}[\exp(\pm cF)] \\ &= \frac{d}{dc} \Big|_{c=0} \log E_{\tilde{P}}[\exp(\pm cF)] = \pm E_{\tilde{P}}[F]. \end{aligned} \quad (\text{B.1})$$

Note that in the first line, we used the fact that $c \mapsto \Lambda_{\tilde{P}}^F(\pm c)/c$ are increasing on $c > 0$, as can be proven using Hölder's inequality. \square

Remark B.2. Again, we emphasize that the choice $G = \log(d\tilde{P}|_{\mathcal{G}}/dP|_{\mathcal{G}})$ is rarely tractable and so Corollary B.1 is not useful in practice. Rather, it shows that any loss of tightness in the bound (3.10) is the result of one's choice of \mathcal{G} and G , and is not intrinsic to the use of that bound.

APPENDIX C. TOOLS FOR COMPUTING THE GOAL-ORIENTED RELATIVE ENTROPY

In this appendix, we present several lemmas that are generally useful when computing the goal-oriented relative entropy. First, we address the question of processes with different initial distributions:

Lemma C.1. *Suppose that the base and alternative models are of the form*

$$P = P^\mu = \int P^x(\cdot)\mu(dx), \quad \tilde{P} = \tilde{P}^{\tilde{\mu}} = \int \tilde{P}^x(\cdot)\tilde{\mu}(dx) \quad (\text{C.1})$$

for some probability kernels, P^x and \tilde{P}^x , from $(\mathcal{X}, \mathcal{M})$ to (Ω, \mathcal{F}) , and some probability measures (e.x. initial distributions), μ and $\tilde{\mu}$, on $(\mathcal{X}, \mathcal{M})$.

In this case, one can use the chain-rule for relative entropy (see, e.g., [19]) to bound

$$R(F_*\tilde{P}^{\tilde{\mu}}\|F_*P^\mu) \leq R(\tilde{\mu}\|\mu) + \int R(F_*\tilde{P}^x\|F_*P^x)\tilde{\mu}(dx). \quad (\text{C.2})$$

Remark C.2. Note that without further assumptions, equation (C.2) is not an equality. For example, let ν be a probability measure on \mathbb{R} , $F = \pi_2$ (the projection onto the second component in $\mathbb{R} \times \mathbb{R}$), and for $x \in \mathbb{R}$ define $P^x = \tilde{P}^x = \delta_x \times \nu$. Then for $\mu \neq \tilde{\mu}$ we have $P^\mu = \mu \times \nu$, $\tilde{P}^{\tilde{\mu}} = \tilde{\mu} \times \nu$ and

$$R(F_*\tilde{P}^{\tilde{\mu}}\|F_*P^\mu) = R(\nu\|\nu) = 0 < R(\tilde{\mu}\|\mu) + \int R(F_*\tilde{P}^x\|F_*P^x)\tilde{\mu}(dx). \quad (\text{C.3})$$

Intuitively, inequality in equation (C.2) arises when the QoI F “forgets” some of the information regarding the discrepancy between the initial condition, as the right-hand-side in equation (C.2) incorporates the full discrepancy.

Remark C.3. Recall that $R(F_*\tilde{P}^{\tilde{\mu}}\|F_*P^\mu) = R(\tilde{P}^{\tilde{\mu}}|_{\sigma(F)}\|P^\mu|_{\sigma(F)})$, where $\sigma(F)$ is the sigma algebra generated by F . If F is measurable with respect to a sub sigma algebra, \mathcal{G} , then $\sigma(F) \subset \mathcal{G}$ and one can weaken the bound via $R(F_*\tilde{P}^x\|F_*P^x) \leq R(\tilde{P}^x|_{\mathcal{G}}\|P^x|_{\mathcal{G}})$, as discussed in Theorem 3.1.

The main tool we use to compute the relative entropy up to a stopping time, Lemma C.5 below, requires one to work with bounded stopping times. This limitation can often be overcome by taking limits. For example:

Lemma C.4. *Suppose that one is working within the framework of Assumption 4.1 and one of the following two conditions holds:*

- (1) τ is finite P -a.s. and \tilde{P} -a.s.
- (2) $F_n \rightarrow F_\infty$ as $n \rightarrow \infty$, $n \in \mathbb{Z}^+$, both P and \tilde{P} -a.s.

Then

$$\begin{aligned} R((F_\tau)_*\tilde{P}\|(F_\tau)_*P) &\leq \liminf_{n \rightarrow \infty} R((F_{\tau \wedge n})_*\tilde{P}\|(F_{\tau \wedge n})_*P) \\ &\leq \liminf_{n \rightarrow \infty} R(\tilde{P}|_{\mathcal{F}_{\tau \wedge n}}\|P|_{\mathcal{F}_{\tau \wedge n}}). \end{aligned} \quad (\text{C.4})$$

Proof. Either condition implies that $(F_{\tau \wedge n})_*\tilde{P} \rightarrow (F_\tau)_*\tilde{P}$ and $(F_{\tau \wedge n})_*P \rightarrow (F_\tau)_*P$ weakly as $n \rightarrow \infty$. The result then follows from lower-semicontinuity of relative entropy (see [19]). \square

The general tool we use to compute relative entropy up to a stopping time is the optional sampling theorem (see, e.g., [34]), as summarized by the following lemma:

Lemma C.5. Let P, \tilde{P} be probability measures on the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in \mathcal{T}})$ such that $\tilde{P}|_{\mathcal{F}_t} \ll P|_{\mathcal{F}_t}$ for all t .

Then $\rho_t = \frac{d\tilde{P}|_{\mathcal{F}_t}}{dP|_{\mathcal{F}_t}}$ is an (\mathcal{F}_t, P) -martingale. In the continuous-time case, we also assume that we have a right-continuous version of ρ_t .

Given this, for any stopping time, τ , we have $\frac{d\tilde{P}|_{\mathcal{F}_{\tau \wedge t}}}{dP|_{\mathcal{F}_{\tau \wedge t}}} = E(\rho_t | \mathcal{F}_{\tau \wedge t}) = \rho_{\tau \wedge t}$ and

$$R(\tilde{P}|_{\mathcal{F}_{\tau \wedge t}} \| P|_{\mathcal{F}_{\tau \wedge t}}) = E_{\tilde{P}}[\log(\rho_{\tau \wedge t})]. \quad (\text{C.5})$$

Using Lemma C.5, the relative entropy up to a stopping time for various types of Markov processes can be computed *via* standard techniques; see Appendix D.

APPENDIX D. RELATIVE ENTROPY UP TO A STOPPING TIME FOR MARKOV PROCESSES

Below, we use Lemma C.5 to give explicit formulas for the relative entropy up to a stopping time for several classes of Markov processes in both discrete and continuous time.

D.1. Discrete-time Markov processes

Let \mathcal{X} be a Polish space, $p(x, dy)$ and $\tilde{p}(x, dy)$ be transition probabilities on \mathcal{X} and $\mu, \tilde{\mu}$ be probability measures on \mathcal{X} . Let P^μ and $\tilde{P}^{\tilde{\mu}}$ be the induced probability measures on $\Omega \equiv \prod_{i=0}^{\infty} \mathcal{X}$ with transition probabilities p and \tilde{p} respectively, and initial distribution μ and $\tilde{\mu}$. Let $\pi_n : \Omega \rightarrow \mathcal{X}$ be the natural projections, $\mathcal{F}_n = \sigma(\pi_m : m \leq n)$, and τ be a \mathcal{F}_n -stopping time.

We have the following bound on the relative entropy, *via* the optional sampling theorem:

Lemma D.1. Suppose that $\tilde{p}(x, dy) = g(x, y)p(x, dy)$ for some measurable $g : \mathcal{X} \times \mathcal{X} \rightarrow [0, \infty)$ and let $N \geq 0$. Then

$$R(\tilde{P}^x|_{\mathcal{F}_{\tau \wedge N}} \| P^x|_{\mathcal{F}_{\tau \wedge N}}) = \tilde{E}^x \left[\log \left(\prod_{i=1}^{\tau \wedge N} g(\pi_{i-1}, \pi_i) \right) \right]. \quad (\text{D.1})$$

D.2. Continuous-time Markov chains

Let \mathcal{X} be a countable set, P and \tilde{P} be probability measures on (Ω, \mathcal{F}) , and $X_t : \Omega \rightarrow \mathcal{X}$, $t \in [0, \infty)$, such that $(\Omega, \mathcal{F}, P, X_t)$ and $(\Omega, \mathcal{F}, \tilde{P}, X_t)$ are continuous-time Markov chains (CTMCs) with transition probabilities $a(x, y)$, $\tilde{a}(x, y)$, jump rates $\lambda(x)$, $\tilde{\lambda}(x)$, and initial distributions $\mu, \tilde{\mu}$ respectively. Let \mathcal{F}_t be the natural filtration for X_t , and X_n^J be the embedded jump chain with jump times $J_n \in [0, \infty]$. We assume $J_n(\omega) \rightarrow \infty$ as $n \rightarrow \infty$ for all ω .

Suppose $\tilde{\mu} \ll \mu$ and that whenever $a(x, y) = 0$ we also have $\tilde{a}(x, y) = 0$ (note that this also implies that $\tilde{\lambda}(x) = 0$ whenever $\lambda(x) = 0$). Then $\tilde{P}|_{\mathcal{F}_t} \ll P|_{\mathcal{F}_t}$ for all $t \geq 0$ and

$$\frac{d\tilde{P}|_{\mathcal{F}_t}}{dP|_{\mathcal{F}_t}} = \frac{\tilde{\mu}(X_0)}{\mu(X_0)} \exp \left(\int_0^t \log \left(\frac{\tilde{\lambda}(X_{s-})\tilde{a}(X_{s-}, X_s)}{\lambda(X_{s-})a(X_{s-}, X_s)} \right) N(ds) - \int_0^t \tilde{\lambda}(X_s) - \lambda(X_s) ds \right), \quad (\text{D.2})$$

where $N_t = \max\{n : J_n \leq t\}$ is the number of jumps up to time t (see, *e.g.*, [32]). Note that in defining the functions $\tilde{\mu}(x)/\mu(x)$ and $\tilde{\lambda}(x)\tilde{a}(x, y)/(\lambda(x)a(x, y))$, we use the convention $0/0 \equiv 1$.

Lemma C.5 then yields the following:

Lemma D.2.

$$R(\tilde{P}|_{\mathcal{F}_{\tau \wedge t}} \| P|_{\mathcal{F}_{\tau \wedge t}}) = R(\tilde{\mu} \| \mu) + E_{\tilde{P}} \left[\int_0^{\tau \wedge t} \log \left(\frac{\tilde{\lambda}(X_{s-})\tilde{a}(X_{s-}, X_s)}{\lambda(X_{s-})a(X_{s-}, X_s)} \right) N(ds) - \int_0^{\tau \wedge t} \tilde{\lambda}(X_s) - \lambda(X_s) ds \right]. \quad (\text{D.3})$$

D.3. Change of drift for SDEs

Here we consider the case where P^x and \tilde{P}^x are the distributions on path-space, $C([0, \infty), \mathbb{R}^n)$, of the solution flows X_t^x and \tilde{X}_t^x of a pair of SDEs starting at $x \in \mathbb{R}^n$. More precisely:

Assumption D.3. *Assume:*

- (1) We have filtered probability spaces, $(M, \mathcal{G}_\infty, \mathcal{G}_t, P)$, $(\tilde{M}, \tilde{\mathcal{G}}_\infty, \tilde{\mathcal{G}}_t, \tilde{P})$ that satisfy the usual conditions [34] and are equipped with \mathbb{R}^m -valued Wiener processes W_t and \tilde{W}_t respectively (Wiener processes with respect to the respective filtrations).
- (2) We have another probability space (N, \mathcal{N}, ν) , a \mathcal{G}_0 -measurable random quantity $Z : M \rightarrow N$, and a $\tilde{\mathcal{G}}_0$ -measurable random quantity $\tilde{Z} : \tilde{M} \rightarrow N$, and they are both distributed as ν .

Remark D.4. Intuitively, we are thinking of Z as representing some outside data, independent of $\{W_t\}_{t \geq 0}$, that is informing the model. \mathcal{G}_t can then be taken to be the completed sigma algebra generated by Z and the Wiener process up to time t . Since Z and the Wiener process are independent, W_t is still a Wiener process with respect to the resulting filtration.)

- (3) We have $x \in \mathbb{R}^n$ and \mathbb{R}^n -valued processes X_t^x and \tilde{X}_t^x that are adapted to \mathcal{G}_t and $\tilde{\mathcal{G}}_t$ respectively, and satisfy the SDEs

$$dX_t^x = b(t, X_t^x, Z)dt + \sigma(t, X_t^x, Z)dW_t, \quad X_0^x = x, \quad (\text{D.4})$$

$$d\tilde{X}_t^x = \tilde{b}(t, \tilde{X}_t^x, \tilde{Z})dt + \sigma(t, \tilde{X}_t^x, \tilde{Z})d\tilde{W}_t, \quad \tilde{X}_0^x = x, \quad (\text{D.5})$$

where $b : [0, \infty) \times \mathbb{R}^n \times N \rightarrow \mathbb{R}^n$, $\sigma : [0, \infty) \times \mathbb{R}^n \times N \rightarrow \mathbb{R}^{n \times m}$, and $\beta : [0, \infty) \times \mathbb{R}^n \times N \rightarrow \mathbb{R}^m$ are measurable and the modified drift is

$$\tilde{b} = b + \sigma\beta. \quad (\text{D.6})$$

- (4) $\beta(s, X_s^x) \in L_{\text{loc}}^2(W)$ and the following process is a martingale under P :

$$\rho_T^x \equiv \exp \left(\int_0^T \beta(s, X_s^x, Z) \cdot dW_s - \frac{1}{2} \int_0^T \|\beta(s, X_s^x, Z)\|^2 ds \right). \quad (\text{D.7})$$

This holds if the Novikov condition is satisfied; see, for example, page 199 of [34].

- (5) The SDE equation (D.5) satisfies the following uniqueness in law property for all $T > 0$: Suppose X_t^t is a continuous, adapted process on another filtered probability space satisfying the usual conditions, $(M', \mathcal{G}'_\infty, \mathcal{G}'_t, P')$, that is also equipped with \mathcal{G}'_t -Wiener process W'_t and a N -valued \mathcal{G}'_0 -measurable random quantity, Z' , distributed as ν . Finally, suppose also that X'_t satisfies

$$X'_t = x + \int_0^t \tilde{b}(s, X'_s, Z')ds + \int_0^t \sigma(s, X'_s, Z')dW'_t \quad (\text{D.8})$$

for $0 \leq t \leq T$. Then the joint distribution of $(X'|_{[0, T]}, Z')$ equals the joint distribution of $(\tilde{X}^x|_{[0, T]}, \tilde{Z})$.

Given this, we define $P^x = (X^x)_*P$ and $\tilde{P}^x = (\tilde{X}^x)_*\tilde{P}$, i.e., the distributions on path-space,

$$(\Omega, \mathcal{F}_\infty, \mathcal{F}_t) \equiv (C([0, \infty), \mathbb{R}^n), \mathcal{B}(C([0, \infty), \mathbb{R}^n)), \sigma(\pi_s, s \leq t)), \quad (\text{D.9})$$

where π_t denotes evaluation at time t and $C([0, \infty), \mathbb{R}^n)$ is given the topology of uniform convergence on compact sets. For $T \geq 0$, we let $P_T^x \equiv P^x|_{\mathcal{F}_T}$ and $\tilde{P}_T^x \equiv \tilde{P}^x|_{\mathcal{F}_T}$.

The relative entropy can be computed *via* Girsanov's theorem, together with Lemma C.5:

Lemma D.5. *Let $x \in \mathbb{R}^n$ and suppose Assumption D.3 holds. For all $T > 0$ we have*

$$R(\tilde{P}^x|_{\mathcal{F}_{\tau \wedge T}} \| P^x|_{\mathcal{F}_{\tau \wedge T}}) \leq \frac{1}{2} E_{\tilde{P}} \left[\int_0^{\tilde{\sigma}_T^x} \|\beta(s, \tilde{X}_s^x, \tilde{Z})\|^2 ds \right], \quad (\text{D.10})$$

where $\tilde{\sigma}_T^x = (\tau \circ \tilde{X}^x) \wedge T$. In terms of the base process, we have

$$R(\tilde{P}^x|_{\mathcal{F}_{\tau \wedge T}} \| P^x|_{\mathcal{F}_{\tau \wedge T}}) \leq \frac{1}{2} E_P \left[\int_0^{\sigma_T^x} \|\beta(s, X_s^x, Z)\|^2 ds \rho_{\sigma_T^x}^x \right], \quad (\text{D.11})$$

where $\sigma_T^x = (\tau \circ X^x) \wedge T$ and

$$\rho_{\sigma_T^x}^x = \exp \left(\int_0^{\sigma_T^x} \beta(s, X_s^x, Z) \cdot dW_s - \frac{1}{2} \int_0^{\sigma_T^x} \|\beta(s, X_s^x, Z)\|^2 ds \right). \quad (\text{D.12})$$

A class of alternative models that are of particular case of interest is covered by the following corollary:

Corollary D.6. *Suppose the baseline model is the solution to the following SDE on \mathbb{R}^n ,*

$$dX_t^x = b(t, X_t^x)dt + \sigma(t, X_t^x)dW_t, \quad X_0^x = x, \quad (\text{D.13})$$

and the alternative model is given by

$$d\tilde{X}_t^x = b(t, \tilde{X}_t^x)dt + \sigma(t, \tilde{X}_t^x)\beta(t, \tilde{X}_t^x, \tilde{Y}_t^y, \tilde{Z})dt + \sigma(t, \tilde{X}_t^x)d\tilde{W}_t^1, \quad \tilde{X}_0^x = x, \quad (\text{D.14})$$

which includes a modified drift that depends on external data, \tilde{Z} , and is also coupled to another SDE

$$d\tilde{Y}_t^y = f(t, \tilde{X}_t^x, \tilde{Y}_t^y, \tilde{Z})dt + \eta(t, \tilde{X}_t^x, \tilde{Y}_t^y, \tilde{Z})d\tilde{W}_t^2, \quad (\text{D.15})$$

on \mathbb{R}^k , where \tilde{W}_t^1 and \tilde{W}_t^2 are independent Wiener processes.

Suppose also we have Y_t^y such that

$$dY_t^y = f(t, X_t^x, Y_t^y, Z)dt + \eta(t, X_t^x, Y_t^y, Z)dW_t', \quad Y_0^y = y, \quad (\text{D.16})$$

where W_t' is a Wiener process independent from W .

If the base, (X_t^x, Y_t^y) , and alternative, $(\tilde{X}_t^x, \tilde{Y}_t^y)$, systems satisfy Assumption D.3, we have

$$\begin{aligned} R\left(\left(\tilde{X}_*^x \tilde{P}\right)|_{\mathcal{F}_{\tau \wedge T}} \left\| \left(X_*^x P\right)|_{\mathcal{F}_{\tau \wedge T}}\right.\right) &\leq R\left(\left(\left(\tilde{X}^x, \tilde{Y}^y\right)_* \tilde{P}\right)|_{\mathcal{F}_{(\tau \circ \pi_1) \wedge T}} \left\| \left(\left(X^x, Y^y\right)_* P\right)|_{\mathcal{F}_{(\tau \circ \pi_1) \wedge T}}\right.\right) \\ &\leq \frac{1}{2} E_{\tilde{P}} \left[\int_0^{(\tau \circ \tilde{X}^x) \wedge T} \|\beta(s, \tilde{X}_s^x, \tilde{Y}_s^y, \tilde{Z})\|^2 ds \right] \end{aligned} \quad (\text{D.17})$$

for any stopping time τ on $C([0, \infty), \mathbb{R}^n)$. Here π_1 denotes the map that takes a path in $C([0, \infty), \mathbb{R}^n \times \mathbb{R}^k)$ and returns the path of the \mathbb{R}^n -valued component.

APPENDIX E. UQ FOR GENERAL DISCOUNTED OBSERVABLES

Theorem 4.6 can be generalized beyond the case of exponential discounting. The proof is very similar and so it is omitted.

Theorem E.1. *Let π be a sigma-finite positive measure on \mathcal{T} , $\tau : \Omega \rightarrow \bar{\mathcal{T}}$ be a \mathcal{F}_t -stopping time, $f : \mathcal{T} \times \Omega \rightarrow \bar{\mathbb{R}}$ be progressively measurable, and suppose we also have one of the following two conditions:*

(1) π is a finite measure and $f_s \geq 0$ for π -a.e. s .

(2) $E_{\tilde{P}} \left[\int_{[0, \tau]} |f_s| d\pi \right] < \infty$.

For $t \in \overline{T}$ define the progressive process $F_t = \int_{[0, t]} f_s \pi(ds)$. Then

$$\pm E_{\tilde{P}} [F_\tau] \leq \int_T \inf_{c > 0} \left\{ \frac{1}{c} \Lambda_P^{1_{s \leq \tau} f_s} (\pm c) + \frac{1}{c} R(\tilde{P}|_{\mathcal{F}_{s \wedge \tau}} \| P|_{\mathcal{F}_{s \wedge \tau}}) \right\} \pi(ds).$$

APPENDIX F. HITTING TIMES OF BROWNIAN MOTION WITH CONSTANT DRIFT

For the reader's convenience, here we recall the moment generating function and distribution of hitting times for Brownian motion with constant drift; see, for example, Chapter 8 of [52] and page 196 in [34].

Lemma F.1. *Let W_t be a \mathbb{R} -valued Wiener process (i.e., Brownian motion) on (Ω, \mathcal{F}, P) , $\mu \in \mathbb{R}$, $a \in \mathbb{R} \setminus \{0\}$. Let $X_t = \mu t + W_t$ be Brownian motion with constant drift μ . Define the level- a hitting time $\tau_a = \inf\{t \geq 0 : X_t = a\}$. For any measurable $A \subset (0, \infty]$ we have*

$$P(\tau_a \in A) = \frac{|a|}{\sqrt{2\pi}} \int_{A \cap (0, \infty)} t^{-3/2} e^{-\frac{(a-\mu t)^2}{2t}} dt + \left(1 - e^{\mu a - |\mu a|}\right) 1_{\infty \in A}. \quad (\text{F.1})$$

In addition, for $\lambda > 0$ the MGF is

$$E[e^{-\lambda \tau_a}] = \exp\left(a\mu - |a|\sqrt{\mu^2 + 2\lambda}\right). \quad (\text{F.2})$$

APPENDIX G. EXPECTED HITTING TIMES OF PERTURBED BROWNIAN MOTION

In this appendix we bound the expected hitting times (5.4) for perturbations of Brownian motion of the form (5.2), using the techniques of Section 4.2. We assume that μ and a have the same sign in order to ensure the hitting time is a.s.-finite under the baseline model. For a given $\alpha > 0$ we specialize equation (5.2) to the following:

Alternative Models: perturbations by drifts, β , with $\|\beta\|_\infty \leq \alpha$. (G.1)

This example uses the technique of Section 4.2.

The optimal bounds can be obtained directly (i.e., without using the UQ methods developed above) by combining the known distribution of τ_a [34] with the comparison principle, i.e., with

$$(\mu - \|\beta\|_\infty)t + \widetilde{W}_t \leq \widetilde{X}_t \leq (\mu + \|\beta\|_\infty)t + \widetilde{W}_t. \quad (\text{G.2})$$

We assume that $\alpha < \mu$ so that the constant-drift lower bound, and hence also the perturbed model, have a.s.-finite hitting times. This yields

$$\frac{a}{\mu + \alpha} \leq E_{\tilde{P}^0}[\tau_a] \leq \frac{a}{\mu - \alpha}, \quad (\text{G.3})$$

where upper and lower bounds are achieved in the cases of constant drift perturbations α and $-\alpha$ respectively.

The optimal bounds equation (G.3) provide a useful test-case for our goal-oriented UQ bounds. Equation (G.1) together with Girsanov's theorem gives a bound on the relative entropy:

$$R(\tilde{P}^0|_{\mathcal{F}_{\tau_a \wedge n}} \| P^0|_{\mathcal{F}_{\tau_a \wedge n}}) \leq \frac{\alpha^2}{2} E_{\tilde{P}^0}[\tau_a \wedge n] \quad (\text{G.4})$$

for all n ; in the language of Theorem 3.6, we can take $\mathcal{G} = \mathcal{F}_{\tau_a \wedge n}$ and $G = \frac{\alpha^2}{2}(\tau_a \wedge n)$.

Using the known formula for the cumulant generating function of τ_a (see Chap. 8 of [52]) and analytic continuation, the cumulant generating function in the baseline model can be computed:

$$\Lambda_{P^0}^{\tau_a}(c) = a\mu - a\sqrt{\mu^2 - 2c}, \quad c < \mu^2/2. \quad (\text{G.5})$$

Therefore Corollary 4.5 yields

$$\begin{aligned} & - \inf_{c>0} \left\{ (c+K)^{-1} \left(a\mu - a\sqrt{\mu^2 + 2c} \right) \right\} \\ & \leq E_{\tilde{P}^0}[\tau_a] \leq \inf_{K < \lambda < \mu^2/2} \left\{ (\lambda - K)^{-1} \left(a\mu - a\sqrt{\mu^2 - 2\lambda} \right) \right\}, \end{aligned} \quad (\text{G.6})$$

where $K \equiv \frac{1}{2}\alpha^2$. The above optimization problems can be solved explicitly, with minimizers $c^* = K + \sqrt{2K}\mu$ and $\lambda^* = -K + \sqrt{2K}\mu$ respectively. Computing the corresponding minimum values, one finds that the UQ bounds resulting from our method, equation (G.6), are the same as the optimal bounds, equation (G.3), that were obtained from the comparison principle. We also note that, while the comparison principle is only available in very specialized circumstances, our UQ method is quite general; in particular, our method is not restricted to 1-D systems. Finally, the non-goal-oriented UQ bounds, equation (2.4), are not effective here; the non-goal-oriented relative entropy over an infinite time-horizon is $R(\tilde{P}^0 \| P^0)$, which is infinite for the constant drift perturbations $\beta = \pm\alpha$.

APPENDIX H. VASICEK MODEL UQ BOUNDS

Here we record a pair of lemmas that are needed for the analysis of option values under the Vasicek interest rate model, as well as the final UQ bound obtained by our method. See Section 5.5.1 for definitions of the notation.

Lemma H.1. *Suppose $r > \frac{\tilde{\sigma}^2}{2\gamma^2}$. Then $\lim_{n \rightarrow \infty} \int_0^n r + \Delta r_s ds = \infty$ a.s. under both the base and the alternative models.*

Proof. The distribution of Δr is the same under both the base and alternative models, so it does not matter which one we consider. As shown in [1], $\int_0^t r + \Delta r_s ds$ is normally distributed for all $t \geq 0$, with mean $m_t \equiv rt$ and variance

$$\tilde{\sigma}_t^2 = \frac{\tilde{\sigma}^2}{2\gamma^3} (2\gamma t + 4e^{-\gamma t} - e^{-2\gamma t} - 3). \quad (\text{H.1})$$

Using this, for any $R > 0$ we can compute

$$\begin{aligned} P \left(\liminf_{n \rightarrow \infty} \int_0^n r + \Delta r_s ds < R \right) & \leq \lim_{N \rightarrow \infty} \sum_{n \geq N} P \left(\int_0^n r + \Delta r_s ds < R \right) \\ & \leq \lim_{N \rightarrow \infty} \sum_{n \geq N} P \left(\exp \left(- \int_0^n r + \Delta r_s ds \right) > \exp(-R) \right) \\ & \leq e^R \lim_{N \rightarrow \infty} \sum_{n \geq N} e^{-rn + \frac{1}{2}\tilde{\sigma}_n^2} \\ & \leq e^{R + \tilde{\sigma}^2/\gamma^3} \lim_{N \rightarrow \infty} \sum_{n \geq N} e^{-\left(r - \frac{\tilde{\sigma}^2}{2\gamma^2}\right)n} = 0. \end{aligned} \quad (\text{H.2})$$

Therefore

$$P \left(\liminf_{n \rightarrow \infty} \int_0^n r + \Delta r_s ds < \infty \right) = \lim_{R \rightarrow \infty} P \left(\liminf_{n \rightarrow \infty} \int_0^n r + \Delta r_s ds < R \right) = 0. \quad (\text{H.3})$$

□

Lemma H.2.

$$\rho_t \equiv \exp \left(\int_0^t \sigma^{-1} \Delta r_s dW_s - \frac{1}{2} \int_0^t \sigma^{-2} |\Delta r_s|^2 ds \right). \quad (\text{H.4})$$

is a P -martingale.

Proof. As shown in [34] (see Cor. 5.14), it suffices to prove that there exists $\{t_k\}_{k=0}^\infty$ with $0 = t_0 < t_1 < \dots < t_n \nearrow \infty$ such that

$$E \left[\exp \left(\frac{1}{2} \int_{t_{n-1}}^{t_n} \sigma^{-2} |\Delta r_s|^2 ds \right) \right] < \infty \text{ for all } n. \quad (\text{H.5})$$

Letting $\Delta t_n = t_n - t_{n-1}$ and using Jensen's inequality, we have

$$\begin{aligned} E \left[\exp \left(\frac{1}{2} \int_{t_{n-1}}^{t_n} \sigma^{-2} |\Delta r_s|^2 ds \right) \right] &\leq E \left[\Delta t_n^{-1} \int_{t_{n-1}}^{t_n} \exp \left(\frac{\Delta t_n}{2} \sigma^{-2} |\Delta r_s|^2 \right) ds \right] \\ &\leq \sup_{t_{n-1} \leq s \leq t_n} E \left[\exp \left(\frac{\Delta t_n}{2} \sigma^{-2} |\Delta r_s|^2 \right) \right]. \end{aligned} \quad (\text{H.6})$$

Δr_t is normal with mean 0 and variance $\tilde{\sigma}^2(1 - e^{-2\gamma t})/(2\gamma)$. Hence, if we fix $\Delta t_n = \epsilon$ small enough then the upper bound in equation (H.6) is finite for all n . \square

The final UQ bound on the Vasicek model obtained by the calculations in Section 5.5.1 is the following:

$$\pm \tilde{E}[F_\tau] \leq \inf_{c>0} \left\{ \frac{1}{c} \log \int \int_{-\infty}^{\infty} (2\pi\tilde{\sigma}_t^2)^{-1/2} \exp(\pm c(K - L)e^{-z} 1_{t<\infty} 1_{z \geq 0}) e^{-\frac{(z-rt)^2}{2\tilde{\sigma}_t^2}} dz P_{\tau \circ X_r}(dt) + \frac{D}{c} \right\} / K_{\mp}, \quad (\text{H.7})$$

where

$$\begin{aligned} K_{\pm} &= \pm \inf_{c>0} \left\{ \frac{1}{c} \log \left(1 + (e^{\pm c} - 1) \int \frac{1}{2} \left(\operatorname{erf} \left(rt / \sqrt{2\tilde{\sigma}_t^2} \right) + 1 \right) P_{\tau \circ X_r}(dt) \right) + \frac{D}{c} \right\}, \\ D &= \inf_{\lambda>1} \left\{ (\lambda - 1)^{-1} \log \int e^{\gamma t/2} \left[\frac{1}{\sqrt{1 - \frac{\lambda\tilde{\sigma}^2}{\sigma^2\gamma^2}}} \sinh \left(\gamma t \sqrt{1 - \frac{\lambda\tilde{\sigma}^2}{\sigma^2\gamma^2}} \right) \right. \right. \\ &\quad \left. \left. + \cosh \left(\gamma t \sqrt{1 - \frac{\lambda\tilde{\sigma}^2}{\sigma^2\gamma^2}} \right) \right]^{-1/2} P_{\tau \circ X_r}(dt) \right\}, \end{aligned} \quad (\text{H.8})$$

$$\begin{aligned} P_{\tau \circ X_r}(dt) &= \frac{|a|}{\sqrt{2\pi}} t^{-3/2} e^{-(a-\mu t)^2/(2t)} 1_{(0,\infty)} dt, \quad a = -\sigma^{-1} \log(X_0/L), \quad \mu = r/\sigma - \sigma/2, \\ \tilde{\sigma}_t^2 &= \frac{\tilde{\sigma}^2}{2\gamma^3} (2\gamma t + 4e^{-\gamma t} - e^{-2\gamma t} - 3). \end{aligned}$$

In creating Figure 5, the integrals over t and z were computed numerically *via* quadrature methods and the optimization over c was also performed numerically; see Remark 3.2. Finally, note that K_+ and K_- are, respectively, upper and lower bounds on a probability and they reflect that fact:

$$\begin{aligned} K_+ &\leq \lim_{c \rightarrow \infty} \left(\frac{1}{c} \log \left(1 + (e^c - 1) \int \frac{1}{2} \left(\operatorname{erf} \left(rt / \sqrt{2\tilde{\sigma}_t^2} \right) + 1 \right) P_{\tau \circ X_r}(dt) \right) + \frac{D}{c} \right) = 1, \\ K_- &\geq - \lim_{c \rightarrow \infty} \left(\frac{1}{c} \log \left(1 + (e^{-c} - 1) \int \frac{1}{2} \left(\operatorname{erf} \left(rt / \sqrt{2\tilde{\sigma}_t^2} \right) + 1 \right) P_{\tau \circ X_r}(dt) \right) + \frac{D}{c} \right) = 0. \end{aligned} \quad (\text{H.9})$$

APPENDIX I. AMERICAN PUT OPTIONS: BOUNDED RATE PERTURBATIONS

In this appendix, we derive UQ bounds on the value of American put options for another class of rate perturbations. This is a simpler example to analyze than the Vasicek model from Section 5.5.1 and provides a good benchmark case, as we will also have access to comparison principle bounds. We use the terminology and notation introduced in Section 5.5.

Specifically, here we consider interest rate perturbations of the following form:

$$\mathbf{Alternative\ Models:} \quad \Delta r_- \leq h_-(t) \leq \Delta r(t, \cdot, \cdot) \leq h_+(t) \leq \Delta r_+, \quad t \geq 0, \quad (\text{I.1})$$

where $\Delta r_{\pm} \in \mathbb{R}$, $r + \Delta r_- > 0$, and $h_-, h_+ : [0, \infty) \rightarrow \mathbb{R}$. The parameters Δr_- and Δr_+ define a fluctuation range, with h_{\pm} allowing for specification of a time-dependent envelope on the fluctuation.

An elementary bound on \tilde{X}_t follows from the uniform bounds $\Delta r_- \leq \Delta r(t, x, y) \leq \Delta r_+$ together with the comparison principle:

$$X_t^{r+\Delta r_-} \leq \tilde{X}_t \leq X_t^{r+\Delta r_+}, \quad \text{and hence } \tau \circ X^{r+\Delta r_-} \leq \tau \circ \tilde{X} \leq \tau \circ X^{r+\Delta r_+}. \quad (\text{I.2})$$

Using only the information $\Delta r_- \leq \Delta r(t) \leq \Delta r_+$, the optimal UQ bounds can be obtained from equation (I.1) and equation (I.2), together with the exact value for the constant-rate processes, equation (5.66):

$$(K - L) (L/X_0)^{2(r+\Delta r_+)/\sigma^2} \leq E \left[\tilde{V}_{\tau \circ \tilde{X}}[\tilde{X}, \tilde{Y}] \right] \leq (K - L) (L/X_0)^{2(r+\Delta r_-)/\sigma^2}. \quad (\text{I.3})$$

This provides a useful comparison for the UQ bounds derived *via* the methods developed above.

To use the goal-oriented method, first note the relative entropy bounds (obtained from Girsanov's theorem):

$$R(P_{\tilde{X}, \tilde{Y}} |_{\mathcal{F}_{\tau \wedge n}} \| P_{X^r, Y} |_{\mathcal{F}_{\tau \wedge n}}) \leq E_{\tilde{X}, \tilde{Y}} [G_n], \quad G_n[x, y] \equiv \frac{\sigma^{-2}}{2} \int_0^{\tau(x) \wedge n} |\Delta r(s, x_s, y_s)|^2 ds. \quad (\text{I.4})$$

G_n is $\mathcal{F}_{\tau \wedge n}$ -measurable, hence we can use Theorem 3.6 to obtain

$$\pm E_{\tilde{X}, \tilde{Y}} [F_{\tau \wedge n}] \leq \inf_{c > 0} \left\{ \frac{1}{c} \log E_{X^r, Y} \left[\exp \left(\pm c F_{\tau(x) \wedge n} + \frac{\sigma^{-2}}{2} \int_0^{\tau(x) \wedge n} |\Delta r(s, x_s, y_s)|^2 ds \right) \right] \right\}. \quad (\text{I.5})$$

The above steps are justified by the bounds on Δr .

The bounds on Δr also allow us to take $n \rightarrow \infty$ in equation (I.5) to obtain

$$\begin{aligned} \pm E \left[\tilde{V}_{\tau \circ \tilde{X}}[\tilde{X}, \tilde{Y}] \right] &= \pm E_{\tilde{X}, \tilde{Y}} [F_{\tau}] = \lim_{n \rightarrow \infty} \pm E_{\tilde{X}, \tilde{Y}} [F_{\tau \wedge n}] \\ &\leq \inf_{c > 0} \left\{ \frac{1}{c} \log E_{X^r, Y} \left[\exp \left(\pm c F_{\tau(x)} + \frac{\sigma^{-2}}{2} \int_0^{\tau(x)} |\Delta r(s, x_s, y_s)|^2 ds \right) \right] \right\}. \end{aligned} \quad (\text{I.6})$$

Bounding Δr in terms of h_{\pm} then yields

$$\begin{aligned} \pm E \left[\tilde{V}_{\tau \circ \tilde{X}}[\tilde{X}, \tilde{Y}] \right] &\leq \inf_{c > 0} \left\{ \frac{1}{c} \log E_{X^r} [M_c^{\pm}] \right\}, \\ M_c^{\pm}[x] &\equiv \exp \left(\frac{\sigma^{-2}}{2} \int_0^{\tau(x)} |h(s)|^2 ds \pm c(K - L) e^{-\int_0^{\tau(x)} r + h_{\mp}(s) ds} 1_{\tau(x) < \infty} \right), \end{aligned} \quad (\text{I.7})$$

where $h(t) \equiv \max\{h_+(t), -h_-(t)\}$. The expectation can be evaluated using the known distribution of $\tau \circ X^r$ (see the discussion surrounding Eq. (5.67)):

$$\begin{aligned} E_{X^r} [M_c^{\pm}] &= \frac{|a|}{\sqrt{2\pi}} \int_0^{\infty} \exp \left(\frac{\sigma^{-2}}{2} \int_0^t h_s^2 ds \pm c(K - L) e^{-\int_0^t r + h_{\mp}(s) ds} - \frac{(a - \mu t)^2}{2t} \right) t^{-3/2} dt \\ &\quad + \exp \left(\frac{\sigma^{-2}}{2} \int_0^{\infty} h_s^2 ds \right) (1 - e^{\mu a - |\mu a|}). \end{aligned} \quad (\text{I.8})$$

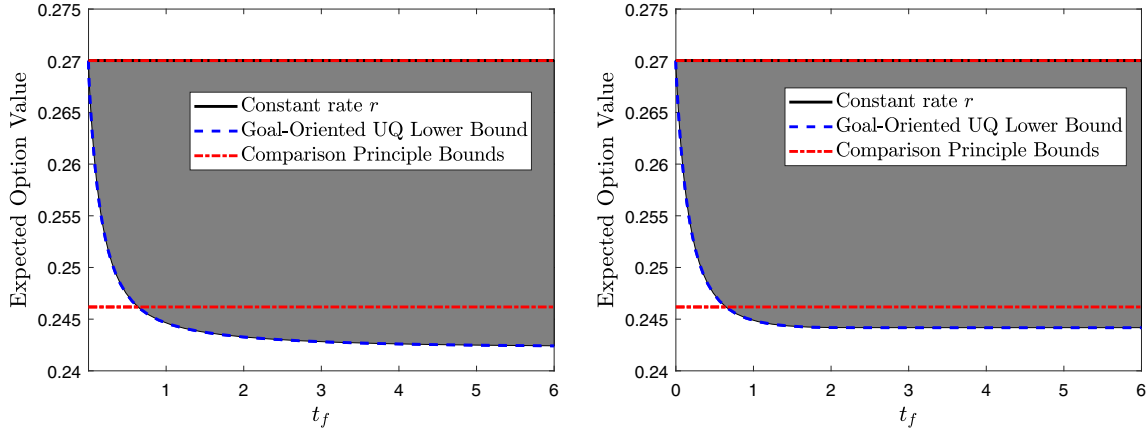


FIGURE I.1. Bounds on the expectation of the option value as a function of t_f , the time by which the interest rate is assumed to have completed the drop from $r + \Delta r_+$ to r ; parameter values are $r = 2$, $K = 1$, $L = 1/2$, $\sigma = 3$, $X_0 = 2$, $\Delta r_+ = 0.3$. The dashed blue curve is the lower bound obtained from our method, equation (I.7). The upper and lower red lines are the comparison-principle bounds based solely on knowledge of $\|\Delta r\|_\infty$; see equation (I.3). The expected value in the constant-rate, r , model (black line, Eq. (5.66)) is the same as the optimal upper bound. The optimal upper bound together with the goal-oriented lower bound constrain the QoI to the gray region. The left plot was obtained using the base-model rate fixed at r and the right plot was obtained by minimizing over the rate parameter that is assigned to the baseline model, as described in Remark I.1 (*i.e.*, replacing r with κ on the right-hand-side of equation (I.7), replacing h_\pm with $h_-^\kappa = r - \kappa$, $h_+^\kappa = r - \kappa + \Delta r_+ 1_{[0, t_f]}$, and minimizing over κ). Note that here, the best performance is achieved by a combination the comparison-principle bounds with our goal-oriented bounds.

Note that all reference to the driving process Y , which is not coupled to X^r , has been eliminated.

Remark I.1. One can obtain somewhat tighter bounds in equation (I.7) by optimizing over the parameter that defines the baseline model, similarly to equation (3.16); in other words, replace r with a parameter, κ , on the right-hand side, find the corresponding h_\pm^κ satisfying $h_-^\kappa \leq r - \kappa + \Delta r \leq h_+^\kappa$, and minimize over κ ; see Figure I.1.

We show numerical results for the following scenario: Suppose the rate drops from $r + \Delta r_+$ to r , and is certain to have completed this drop by time t_f (one could also study a similar scenario where r increases), but the timing and profile of the drop is otherwise unknown, *i.e.*, we consider Δr_+ and t_f to be known, and the only constraint on Δr is $0 \leq \Delta r(t, \cdot, \cdot) \leq \Delta r_+ 1_{[0, t_f]}(t)$. Using $h_- = 0$ and $h_+(t) = \Delta r_+ 1_{[0, t_f]}(t) = h(t)$, UQ bounds are obtained by combining equation (I.7) with equation (I.8).

In Figure I.1 we show the resulting bounds on the option value, as a function of t_f . Note that the bounds resulting from our UQ method are an improvement over the comparison-principle bound equation (I.3) for an initial range of t_f , and stabilizes at a value near the comparison-principle bound for large t_f . For large t_f our method is sub-optimal, but still competitive. Minimizing over the parameters assigned to the baseline model (right plot) improves the bound. We find similar behavior when assuming other time-dependent envelopes of Δr . It should again be noted that the comparison principle is available only in very special circumstances and is used here for benchmarking purposes; see Section 5.5.1 for a more realistic model where the comparison principle is not an effective tool.

Acknowledgements. The research of J.B., M.K., and L. R.-B. was partially supported by NSF TRIPODS CISE-1934846. The research of M.K. and L. R.-B. was partially supported by the National Science Foundation (NSF) under the grant DMS-1515712 and the Air Force Office of Scientific Research (AFOSR) under the grant FA-9550-18-1-0214.

REFERENCES

- [1] M. Abundo, On the representation of an integrated Gauss–Markov process. *Sci. Math. Jpn.* **77** (2015) 357–361.
- [2] D. Anderson, An efficient finite difference method for parameter sensitivities of continuous time Markov chains. *SIAM J. Numer. Anal.* **50** (2012) 2237–2258.
- [3] B.D.O. Anderson and J.B. Moore, *Optimal Control: Linear Quadratic Methods*. Dover Books on Engineering. Dover Publications (2007).
- [4] G. Arampatzis and M.A. Katsoulakis, Goal-oriented sensitivity analysis for lattice kinetic Monte Carlo simulations. *J. Chem. Phys.* **140** (2014) 124108.
- [5] S. Asmussen, O. Nerman and M. Olsson, Fitting phase-type distributions via the EM algorithm. *Scand. J. Stat.* **23** (1996) 419–441.
- [6] R. Atar, K. Chowdhary and P. Dupuis, Robust bounds on risk-sensitive functionals via Rényi divergence. *SIAM/ASA J. Uncertainty Quantif.* **3** (2015) 18–33.
- [7] R. Atar, A. Budhiraja, P. Dupuis and R. Wu, Robust bounds and optimization at the large deviations scale for queueing models via Rényi divergence. Preprint [arXiv:2001.02110](https://arxiv.org/abs/2001.02110) (2020).
- [8] D. Bakry and M. Emery, Hypercontractivité do semi-groupes de diffusion. *C.R. Acad. Sci. Paris Sér I Math.* **299** (1984) 775–778.
- [9] H. Bijl and T.B. Schön, Optimal controller/observer gains of discounted-cost LQG systems. *Automatica* **101** (2019) 471–474.
- [10] J. Birrell and L. Rey-Bellet, Uncertainty quantification for markov processes via variational principles and functional inequalities. *SIAM/ASA J. Uncertainty Quantif.* **8** (2020) 539–572.
- [11] M. Bladt and B.F. Nielsen, Matrix-exponential distributions in applied probability. In: *Probability Theory and Stochastic Modelling*. Springer, New York (2017).
- [12] S. Boucheron, G. Lugosi and P. Massart, *Concentration Inequalities*. Oxford University Press, Oxford (2013).
- [13] S. Boyd, S.P. Boyd and L. Vandenberghe, *Convex optimization*. *Berichte über verteilte messsysteme*, no. pt. 1. Cambridge University Press (2004).
- [14] T. Breuer and I. Csiszár, Measuring distribution model risk. *Math. Finance* **26** (2013) 395–411.
- [15] T. Breuer and I. Csiszár, Systematic stress tests with entropic plausibility constraints. *J. Banking Finance* **37** (2013) 1552–1559.
- [16] L. Brown, N. Gans, A. Mandelbaum, A. Sakov, H. Shen, S. Zeltyn and L. Zhao, Statistical analysis of a telephone call center. *J. Am. Stat. Assoc.* **100** (2005) 36–50.
- [17] K. Chowdhary and P. Dupuis, Distinguishing and integrating aleatoric and epistemic variation in uncertainty quantification. *ESAIM: M2AN* **47** (2013) 635–662.
- [18] T. Dankel, On the distribution of the integrated square of the Ornstein-Uhlenbeck process. *SIAM J. Appl. Math.* **51** (1991) 568–574.
- [19] P. Dupuis and R.S. Ellis, A weak convergence approach to the theory of large deviations. In: *Wiley Series in Probability and Statistics*. John Wiley & Sons, New York (2011).
- [20] P. Dupuis, M.A. Katsoulakis, Y. Pantazis and P. Plecháč, Path-space information bounds for uncertainty quantification and sensitivity analysis of stochastic dynamics. *SIAM/ASA J. Uncertainty Quantif.* **4** (2016) 80–111.
- [21] P. Dupuis, M.A. Katsoulakis, Y. Pantazis and L. Rey-Bellet, Sensitivity analysis for rare events based on Rényi divergence. *Ann. Appl. Probab.* **30** (2020) 1507–1533.
- [22] B. Engelmann and R. Rauhmeier, *The Basel II Risk Parameters: Estimation, Validation, Stress Testing – With Applications to Loan Risk Management*. Springer, Berlin-Heidelberg (2011).
- [23] M.J. Faddy, Examples of fitting structured phase-type distributions. *Appl. Stochastic Models Data Anal.* **10** (1994) 247–255.
- [24] M.I. Freidlin, Functional integration and partial differential equations. In: Vol. 109 of *Annals of Mathematics Studies*. Princeton University Press (2016).
- [25] V. Girardin and N. Limnios, On the entropy for semi-Markov processes. *J. Appl. Probab.* **40** (2003) 1060–1068.
- [26] P. Glasserman, Monte Carlo methods in financial engineering. In: *Stochastic Modelling and Applied Probability*. Springer, New York (2013).
- [27] P. Glasserman and X. Xu, Robust risk measurement and model risk. *Quant. Finance* **14** (2014) 29–58.
- [28] P.W. Glynn, Likelihood ratio gradient estimation for stochastic systems. *Commun. ACM* **33** (1990) 75–84.
- [29] K. Gourgoulis, M.A. Katsoulakis and L. Rey-Bellet, Information metrics for long-time errors in splitting schemes for stochastic dynamics and parallel kinetic Monte Carlo. *SIAM J. Sci. Comput.* **38** (2016) A3808–A3832.
- [30] K. Gourgoulis, M.A. Katsoulakis, L. Rey-Bellet and J. Wang, How biased is your model? Concentration inequalities, information and model bias. *IEEE Trans. Inf. Theory* **66** (2020) 3079–3097.
- [31] M. Hairer and A.J. Majda, A simple framework to justify linear response theory. *Nonlinearity* **23** (2010) 909–922.
- [32] S. Heinz and H. Bessaih, Stochastic equations for complex systems: theoretical and computational topics. In: *Mathematical Engineering*. Springer International Publishing (2015).

- [33] J. Janssen and R. Manca, *Applied semi-Markov Processes*. Springer, New York (2006).
- [34] I. Karatzas and S. Shreve, *Brownian motion and stochastic calculus*. In: *Graduate Texts in Mathematics*. Springer, New York (2014).
- [35] M.A. Katsoulakis, L. Rey-Bellet and J. Wang, Scalable information inequalities for uncertainty quantification. *J. Comput. Phys.* **336** (2017) 513–545.
- [36] D. Kim, B.J. Debusschere and H.N. Najm, Spectral methods for parametric sensitivity in stochastic dynamical systems. *Biophys. J.* **92** (2007) 379–393.
- [37] H. Kushner and P.G. Dupuis, *Numerical methods for stochastic control problems in continuous time*. In: *Stochastic Modelling and Applied Probability*. Springer, New York (2013).
- [38] J.C. Lagarias, J.A. Reeds, M.H. Wright and P.E. Wright, Convergence properties of the nelder-mead simplex method in low dimensions. *SIAM J. Optim.* **9** (1998) 112–147.
- [39] H. Lam, Robust sensitivity analysis for stochastic systems. *Math. Oper. Res.* **41** (2016) 1248–1275.
- [40] F. Liese and I. Vajda, On divergences and informations in statistics and information theory. *IEEE Trans. Inf. Theory* **52** (2006) 4394–4412.
- [41] N. Limnios and G. Oprisan, *Semi-Markov processes and reliability*. In: *Statistics for Industry and Technology*. Birkhäuser Boston (2012).
- [42] W.M. McEneaney, A robust control framework for option pricing. *Math. Oper. Res.* **22** (1997) 202–221.
- [43] J.A. Nelder and R. Mead, A simplex method for function minimization. *Comput. J.* **7** (1965) 308–313.
- [44] H. Owhadi, C. Scovel, T.J. Sullivan, M. McKerns and M. Ortiz, Optimal uncertainty quantification. *SIAM Rev.* **55** (2013) 271–345.
- [45] W. Page, *Applications of mathematics in economics*. MAA notes, Mathematical Association of America (2013).
- [46] Y. Pantazis and M.A. Katsoulakis, A relative entropy rate method for path space sensitivity analysis of stationary complex stochastic dynamics. *J. Chem. Phys.* **138** (2013) 054115.
- [47] H.L. Peter and S.T. J., *Uncertainty within economic models*. In: *World Scientific Series In Economic Theory*. World Scientific Publishing Company (2014).
- [48] S. Plyasunov and A.P. Arkin, Efficient stochastic sensitivity analysis of discrete event systems. *J. Comput. Phys.* **221** (2007) 724–738.
- [49] H. Qian, Open-system nonequilibrium steady state: statistical thermodynamics, fluctuations, and chemical oscillations. *J. Phys. Chem. B* **110** (2006) 15063–15074.
- [50] L. Rey-Bellet and K. Spiliopoulos, Irreversible Langevin samplers and variance reduction: a large deviations approach. *Nonlinearity* **28** (2015) 2081–2103.
- [51] P.W. Sheppard, M. Rathinam and M. Khammash, A pathwise derivative approach to the computation of parameter sensitivities in discrete stochastic chemical systems. *J. Chem. Phys.* **136** (2012) 034115.
- [52] S.E. Shreve, *Stochastic Calculus for Finance II: Continuous-time Models*. In: Vol. 11 of *Springer Finance Textbooks*. Springer, New York (2004).
- [53] L. Wu, A deviation inequality for non-reversible Markov processes. *Ann. Inst. Henri Poincaré (B) Probab. Statistics* **36** (2000) 435–445.