Monte Carlo Methods for Adaptive Disorder Problems

Michael Ludkovski

Abstract We develop a Monte Carlo method to solve continuous-time adaptive disorder problems. An unobserved signal X undergoes a disorder at an unknown time to a new unknown level. The controller's aim is to detect and identify this disorder as quickly as possible by sequentially monitoring a given observation process Y. We adopt a Bayesian setup that translates the problem into a two-step procedure of (i) stochastic filtering followed by (ii) an optimal stopping objective. We consider joint Wiener and Poisson observation processes Y and a variety of Bayes risk criteria. Due to the general setting, the state of our model is the full infinite-dimensional posterior distribution of X. Our computational procedure is based on combining sequential Monte Carlo filtering procedures with the regression Monte Carlo method for high-dimensional optimal stopping problems. Results are illustrated with several numerical examples.

1 Introduction

Disorder detection and isolation is a classical problem in statistical signal processing. In its adaptive or robust form, a signal X changes at a random time θ to a random level χ from its original value μ. The signal is not observed directly, and neither are θ or χ. Instead, partial observations are available in the form of an observation process whose dynamics are driven by X. In this paper we consider such an adaptive disorder problem in continuous time with an additive jump-diffusion observation process whose drift and jump rate depend on X. More precisely, the observations consist of two independent channels, with channel I observing X in Gaussian white noise and channel II observing a counting process with X-dependent intensity.

We focus on a Bayesian risk-minimization problem whereby a controller, based on her up-to-date observations, is asked to sound an alarm, followed immediately with an announcement. The timing of the alarm should be as soon as possible after the disorder time θ and the announcement should match the new level χ. Hence, the controller faces a dynamic control problem of choosing an alarm time τ and announcement d to minimize her risk criterion. The risk is measured through a Bayesian expected cost function based on the posteriors of θ and χ. Namely, starting from the given Bayesian priors for

Michael Ludkovski
Department of Statistics and Applied Probability University of California, Santa Barbara, CA 93106-3110 USA e-mail: ludkovski@pstat.ucsb.edu

(θ, χ), the posterior distribution π_t of X is derived sequentially given observations up to date t, yielding an optimal stopping problem for π.

The Bayesian formulation is relatively under-used in the sequential detection literature because in the absence of finite-dimensional statistics for π, we must work directly with the nonlinear filtering equations. Here we resolve this challenge by providing an efficient computational approximation algorithm for minimizing Bayes risk that preserves the key features of tracking the full posterior and solving a dynamic control problem. Moreover, the Bayesian setup allows us to rigorously quantify the trade-off between multiple risk criteria, as well as input prior information about the disorder distribution that is often available in practice.

Filtering in continuous-time semimartingale models is by now a classical topic rooted in the seminal works of Jacod, Kushner, Zakai and others in 1970s. Unified presentation of additive jump-diffusion models is available in [33] and the recent articles [17, 10]. In our case the signal X is of particularly simple form, consisting of a single-arrival point process. Nevertheless, while a variety of observation models have been treated, the vast majority of literature assumes a parametric setup, namely taking the post-disorder level χ as known, whereby only estimation of the disorder time θ is required. An exception are the inspiring works by [3] for adaptive Poisson disorder and [31] for adaptive Wiener disorder case. A different strand in sequential analysis has analyzed change point detection problems (estimating θ only) with unknown post-disorder levels in discrete-time [21, 26].

The need to estimate the constant post-disorder level χ is related to the problem of parameter estimation in hidden Markov models (HMM). Maximum likelihood treatment of this problem is possible using the EM algorithm, see [17]. However, for Bayes risk objectives the point estimates such maximum likelihood procedures provide are inappropriate, lacking the estimate uncertainty and sequential updating that is required. More relevant for us is therefore the class of sequential Monte Carlo (SMC) methods [15] for nonlinear filtering. These methods, especially common for discrete-time models, approximate π with an empirical particle cloud that is sequentially propagated according to a mutation-selection procedure. Similar to our disorder setting, Chopin [6, 7] considered SMC filtering of jump Markov processes, while [19, 32] studied tracking applications in engineering that require SMC filtering of piecewise deterministic processes observed in Gaussian noise. The reference volumes [5] and [1] summarize the current state-of-the-art of SMC.

From the control perspective, our setup is closely aligned with optimal stopping of continuous-time HMMs which goes back to the original disorder problem posed by Shiryaev in 1960s [30]. Historically, the focus has been on analytic methods that assume known χ and exponentially distributed θ, see [28, 10, 12]. In the recent paper [25] the author together with S. Sezer considered a general optimal stopping problem with similar risk objectives and Poissonian observations for a generic *finite-dimensional* hidden Markov chain X.

In this paper we generalize the previous models for Poisson and Wiener disorders in [3] and [31] respectively, to provide a computational framework for minimizing Bayesian risk in adaptive disorder problems. Thus, we allow for simultaneous observations of a point-process and diffusion channels, as well as arbitrary disorder time and post-disorder rate distributions. Finally, we permit a wide variety of Bayes risk functions, including most of those proposed in the literature, such as probability of false alarm, detection delay, and penalty for mis-identifying the post-disorder rate. In our general setting the resulting filtering problem is infinite-dimensional which leads to a stopping problem with measure-valued state variables. We resolve these challenges by linking together the aforementioned literatures on SMC methods for change-point detection and the recent simulation-based algorithms for optimal stopping, see e.g. [16, 23]. More precisely, we propose to use particle filters for the filtering step

and regression Monte Carlo for the optimization step. Thanks to these choices, our numerical scheme is fully Monte Carlo based, and generates computational efficiencies by integrating the filtering and Snell envelope calculations. Moreover, the algorithm is robust to model specification, requiring only the ability to simulate the underlying stochastic processes and likelihood ratios. This approach was first proposed in [24] in the context of classical nonlinear filtering of diffusions. Herein we adapt it to the robust disorder problem which requires significant adjustments both in the model setup and computational implementation.

Besides theoretical interest, the (adaptive) disorder problem has numerous applications in reliability theory, threat detection, and finance and insurance investment decisions. We refer to [25] for a catalogue of problems that correspond to optimal stopping of a partially observed Markov chain which is a finite-dimensional version of our setup. Let us highlight two typical examples. In cyber-security, a controller must detect and identify unusual network traffic that might be a sign of a security intrusion. Depending on network volume, packet traffic Y could be modeled either as a point process or as a Brownian motion (or a multi-scale combination of the two). The main trade-off is then between detection delay and frequency of false alarms.

In real options, project manager must decide on her start-up investments to match the demand level. Consider a novel technology product for which the demand level rises from μ to χ at some unknown date θ. The manager must identify χ and θ as close as possible in order to minimize lost profit opportunities while avoiding over-capacity. Managers make their decisions by monitoring market conditions that are observed through frequency of positive economic events, as well as a related index that is modeled as a Brownian motion with drift X_t.

The rest of the paper is organized as follows. In Section 2 we provide a rigorous formulation of our Bayesian adaptive disorder problem, in particular following the reference measure approach of [3]. Section 3 is devoted to the filtering sub-problem, while Section 4 describes the optimal stopping sub-problem. Taken together, these sections provide a complete description of our numerical algorithm, which is summarized in pseudo-code in the Appendix. Section 5 illustrates our approach with several numerical examples. Finally, Section 6 points out possible generalizations and directions for future research.

2 Problem Formulation

2.1 Canonical setup

Let $(\Omega,\mathscr{F},\mathbb{P}_0)$ be a probability space supporting a Poisson process $N=\{N_t;t\geq 0\}$ with given intensity $\mu>0$ and an independent Wiener process $Y=\{Y_t;t\geq 0\}$. The arrival times of N are denoted as $\sigma_1,\sigma_2,\ldots$. We assume that this space in addition supports two positive random variables θ and χ, independent of N and Y. The distribution of χ is $F(\cdot)$ and the distribution of θ is $G(\cdot)$. We assume that G is absolutely continuous except possibly for a point mass of weight p_0 at zero, so that $\mathbb{P}_0(\theta=0)=p_0$. All of μ,F,G are assumed known.

The quantities θ and χ are used to define the signal process X. Namely, X is given by the right-continuous, piecewise-constant process

$$X_t=\mu 1_{t<\theta}+\chi 1_{t\geq\theta},\quad t\geq 0. \tag{1}$$

Because χ and θ are independent, the pair (X_t,t) form a Markov process. More abstractly, X can be viewed as a simple point process with a single arrival time θ and corresponding mark χ.

Thanks to absolute continuity of G, we may define the hazard rate λ_t of θ via $\lambda_t \triangleq \frac{G(dt)}{1-G(t)}$. In the special case where θ is a mixture of the point mass at zero and an exponential $Exp(\lambda)$ distribution, $\lambda_t \equiv \lambda$ is constant, making X a time-homogenous continuous-time Markov chain.

We denote by $\mathbb{Y} = \{\mathscr{Y}_t\}_{t\geq 0}$ the right-continuous augmentation of the natural filtration $\sigma\,(N_s, Y_s : 0 \leq s \leq t)$ of (N,Y) and define the extended filtration $\mathbb{F} = \{\mathscr{F}_t\}_{t\geq 0}$ where

$$\mathscr{F}_t \triangleq \mathscr{Y}_t \vee \sigma\{\theta,\chi\}, \qquad t \geq 0.$$

The signal X is a $\mathscr{F}$-semimartingale. Let us define the $(\mathbb{P}_0,\mathscr{F})$ square-integrable martingale

$$\begin{aligned} M_t &= 1_{\theta\leq t} - \int_0^{t\wedge\theta} \lambda_s\, ds \\ &= 1_{\theta\leq t} + \log\left[1 - G\left(\int_0^t 1_{\theta>u}\, du\right)\right]. \end{aligned}$$

Then the indicator $E_t = 1_{\theta\leq t}$ admits the decomposition (see [31])

$$dE_t = \lambda_t(1-E_t)\,dt + dM_t,$$

and comparing with (1) we obtain the semi-martingale representation of X as

$$dX_t = \lambda_t(\chi - X_t)\,dt + \chi\, dM_t. \tag{2}$$

We remark that in (2) M is independent of N and Y.

The generator A of the Markov process (X_t,t) [9] is given through its action on a smooth function $f(x,t)$ as

$$(Af)(x,t) = \frac{\partial}{\partial t} f(x,t) + 1_{x=\mu}\lambda_t \int_{\mathbb{R}_+} \{f(y,t) - f(x,t)\}\, F(dy).$$

In particular, $Af(X_t,t) = \lambda_t[f(\chi,t) - f(X_t,t)]$. In the special case where $F(\cdot)$ is discrete, placing mass at $\ell \geq 1$ points, X is a finite-state inhomogeneous Markov chain taking on $\ell+1$ values and its generator A_t (also known as the transition matrix) is an $(\ell+1)\times(\ell+1)$ matrix. Under the additional assumption of exponential disorder time θ, this classical case corresponds to a continuous-time HMM and is treated in detail in e.g. [3, 25, 17].

2.2 Physical probability $\mathbb{P}$

The reference probability measure $\mathbb{P}_0$ is a theoretical device and our problem is in fact under the physical measure $\mathbb{P}$. In this section we construct $\mathbb{P}$ from $\mathbb{P}_0$ through a change of measure.

Consider the Doleans-Dade exponential martingale

$$\begin{aligned} L_t &= \mathscr{E}_t(h\cdot Y)\mathscr{E}_t(\Lambda\cdot N) \\ &= \exp\left\{\int_0^t h_s\, dY_s - \frac{1}{2}\int_0^t h_s^2\, ds\right\} \exp\left\{\int_0^t \log\Big(\frac{\Lambda_{s-}}{\mu}\Big)\, dN_s - \int_0^t (\Lambda_s - \mu)\, ds\right\}, \end{aligned} \tag{3}$$

where h and Λ are two $\mathscr{F}^X$-measurable functions. Let $T < \infty$ be a given problem horizon. Assuming that $\mathbb{E}_0[L_T] = 1$, i.e. L is a true $\mathbb{P}_0$-martingale on $[0,T]$ we define the $\mathbb{P}_0$-equivalent measure $\mathbb{P}$ via its Radon Nikodym derivative

$$\frac{d\mathbb{P}}{d\mathbb{P}_0}\Big|_{\mathscr{F}_t} = L_t, \qquad 0 \le t \le T. \tag{4}$$

The $\mathbb{P}_0$-martingale L solves the stochastic differential equation

$$\frac{dL_s}{L_{s-}} = h_{s-}\, dY_s + \frac{\Lambda_{s-} - \mu}{\mu}\,(dN_s - \mu\, ds), \tag{5}$$

and is closely related to the likelihood processes from statistical signal processing. It will also play a crucial role in our particle filtering algorithms in Section 3.2. The following proposition (compare [1, Prop 3.13]) is classical and motivates our construction.

Proposition 1. *The following hold under* $\mathbb{P}$*:*

- *The process* $N_t - \int_0^t \Lambda_s\, ds$ *is a* $(\mathbb{P}, \mathscr{Y})$*-martingale.*
- *The process* $W_t \triangleq Y_t - \int_0^t h_s\, ds$ *is a* $(\mathbb{P}, \mathscr{Y})$*-Wiener process, independent of* N*.*
- *The signal* X *is not affected, i.e.* (1) *and* (2) *continue to hold.*

To summarize, under $\mathbb{P}$, the counting process N has stochastic intensity Λ_t (also known as a Cox process) and the continuous process Y (no longer independent of N) satisfies the additive Îto stochastic differential equation $Y_t = \int_0^t h_s\, ds + W_t$.

We will work with the Markov case, $h_s = h(X_s)$ and $\Lambda_s = \Lambda(X_s)$. By relabeling the state-space of X, we may assume without loss of generality that $\Lambda(x) = x = \chi$ is the identity map, an assumption already reflected in (1). As a concrete example, $h(x) = \alpha x$ matches the classical Kalman-Bucy models in the filtering literatures. Thus, we interpret $h(X)$ as the drift of Y, and X as the intensity of N under $\mathbb{P}$. The controller operates under $\mathbb{P}$ and has access only to the observable filtration $\mathscr{Y}_t = \sigma(N,Y)$. In contrast, the full information in the model is conveyed by $\mathscr{F}_t = \mathscr{Y}_t \vee \sigma(X_s : s \le t)$. Observe that under $\mathbb{P}$, it is impossible to fully detect the drift of Y or the intensity of N so the controller has *partial* information.

In the above Markovian case, to guarantee that L is a true $\mathbb{P}_0$-martingale we need the same property for each of the two terms in (3). For the diffusion term, a typical sufficient condition is Novikov's,

$$\mathbb{E}_0\left[\exp\Big\{\frac{1}{2}\int_0^T h^2(X_s)\, ds\Big\}\right] < \infty,$$

and reduces to $\mathbb{E}_0[\exp\{\frac{1}{2}(T-\theta)h^2(\chi)\}] < \infty$. When $h(x) = \alpha x$ is linear, it is equivalent to existence of exponential moments for χ^2. For the jump term, a sufficient integrability condition is

$$\mathbb{E}_0\left[\prod_{k=1}^{N_T} \frac{\Lambda(X_{\sigma_k})}{\mu}\right] = \mathbb{E}_0\left[e^{\mu(T-\theta)_+ \cdot \frac{\chi}{\mu}}\right] \le \mathbb{E}_0\left[e^{T\chi}\right] < \infty,$$

where the first equality used the fact $N_T \sim Poisson(\mu T)$. Thus, it suffices to have a moment generating function for χ with radius of convergence bigger than T.

Remark 1. We note that the above observation scheme is equivalent to observing a single jump-diffusion channel $\bar{Y}$ with the dynamics $d\bar{Y}_t = h_t\, dt + dW_t + dN_t$. Since the paths of $\int h_s\, ds + W$ are

continuous a.s., while the paths of N are discontinuous a.s., the controller can then decompose back her observations into $\bar{Y} = Y^c + Y^d$, $Y^c = \int h_s\,ds + W = Y$ and $Y^d = N$ as above.

2.3 Bayes Risk

The aim of the controller is to minimize *risk* related to the disorder. She does so by stopping the observations at a certain alarm time τ coupled with an announcement d about the post-disorder level. Both τ and d are selected dynamically based on observed data, leading to a stochastic control setup.

We focus on three types of risk: (i) risk of false alarms that occur by early announcements before actual disorder time θ; (ii) risk of detection delay that occurs when the announcement is after θ; (iii) identification risk due to wrong announcement about the post-disorder rate Λ. The corresponding cost functions are frequency of false alarms, detection delay penalty, and penalties for mis-identification of Λ. The controller will select the pair (τ, d) to minimize the expected Bayes risk of these competing objectives.

Let the control variables be the alarm time τ and the decision d, where $\tau \in [0,T]$ and $d \in D$ for some given subset $D \in \mathbb{R}_+$. Let $\mathscr{S}$ be the set of all $\mathscr{Y}$-stopping times smaller than T. Since the controller has access only to the $\mathscr{Y}$-filtration, it is required that $\tau \in \mathscr{S}$ and $d \in \mathscr{Y}_\tau$. Our Bayes risk objective is then to minimize

$$R(\tau,d) \triangleq \mathbb{E}^{\pi_0}\left[1_{\{\tau<\theta\}} + c_1(\tau-\theta)_+ + H(d,X_\tau)\right], \tag{6}$$

where $(x)_+ \equiv \max(x,0)$. Above $c_1 \geq 0$ is a constant related to the penalty for detection delay, and $H(d,x)$ is the penalty for making the announcement d when the true state is x. Without loss of generality, we normalize the cost of false alarms to 1.

The decision variable d corresponds to the decision-maker's best *guess* about the value of χ. However, rather than being simply the maximum likelihood estimate (or the conditional mean) of χ, this announcement is chosen to minimize the given risk criterion $H(d,x)$. Three representative criteria we consider are:

1. The mean squared error between χ and d, $H(d,X_\tau) = 1_{\{\tau\geq\theta\}}(X_\tau - d)^2$, leading to candidate announcement $d^*(\tau) = \mathbb{E}[X_\tau|\mathscr{Y}_\tau, \theta \leq \tau]$;
2. Directional discrepancy between d and χ relative to μ, $H(d,X_\tau) = 1_{\{\tau\geq\theta\}}|1_{X_\tau\geq\mu} - 1_{d\geq\mu}|$, leading to $d^* > \mu$ if and only if $\mathbb{P}(X \geq \mu|\mathscr{Y}_t) > 1/2$, and piecewise linear stopping costs of $\mathbb{P}(\chi \geq \mu|\mathscr{Y}_\tau) \wedge (1 - \mathbb{P}(\chi \geq \mu|\mathscr{Y}_\tau))$;
3. Signed absolute difference between d and χ, $H(d,x) = c_1 1_{\{x<d\}}(d-x) + c_2 1_{\{x\geq d\}}(x-d)$ with candidate announcement d^* as the quantile $\mathbb{P}(X_\tau \leq d^*|\mathscr{Y}_\tau) = \frac{c_2}{c_1+c_2}$.

We emphasize that when a mixture of the above stopping criteria are used, the optimal announcement d^* will be a weighted average of the candidate d's, but can always be straightforwardly computed in any given model. To fix ideas, for the remainder of this section we will consider the mixed weighted penalty,

$$H(d,X_\tau) = c_2 1_{\{X_\tau\neq\mu\}}(X_\tau - d)^2 + c_3 1_{\{X_\tau\neq\mu\}}|1_{X_\tau\geq\mu} - 1_{d\geq\mu}|, \qquad c_2, c_3 \geq 0.$$

When the decision set D is finite we can also view the decision variable d in the context of hypothesis testing. Namely, $d \in D$ is identified with picking the d-th hypothesis about χ and $H(d,x)$ represents the

penalty for mis-identification. We refer to [11, 12] for related hypothesis testing problems of finite-state HMMs.

The Bayes risk R in (6) is affected by the distribution of the random variable X_0. Recall that the prior distribution π_0 of the disorder time is given by

$$\pi_0(A) \triangleq \mathbb{P}(X_0 \in A) = p_0 1_{\mu \in A} + (1-p_0)F(A), \qquad A \in \mathscr{B}(\mathbb{R}_+),$$

for any Borel subset A. Summarizing, the controller's aim is to minimize the Bayes risk, by computing the value function

$$U(\pi_0) \triangleq \inf_{(\tau,d)\in\mathscr{S}\times D} R(\tau,d). \tag{7}$$

In the remainder of the paper we will be concerned with computing U, as well as studying the optimal τ^* and d^*.

Remark 2. In much of the sequential detection literature, one is concerned with estimating θ and χ without any explicit reference to Bayes risk criteria. For example, a popular choice is to announce disorder as soon as a critical threshold b is crossed by a (running) summary statistic of π_t. Thus, a heuristic solution is to e.g. take $\tau := \inf\{t : \mathbb{E}[\chi|\mathscr{Y}_t] > b\}$ and $d := \mathbb{E}[\chi|\mathscr{Y}_\tau]$ for appropriately chosen b. Such threshold rules can be asymptotically justified through e.g. a quickest detection problem with a constraint on false alarm probability [2]. In contrast, our formulation of $R(\tau,d)$ incorporates *multiple* simultaneous risk objectives, making it difficult to justify any particular threshold strategy or to come up with a good statistic that incorporates all the risk criteria.

3 Filtering

Because X is not observed, (6) is not in standard form, as the rewards are not adapted to the controller's filtration. Applying iterated expectations and noting $(\tau-\theta)_+ = \int_0^\tau 1_{\theta\le t}\,dt$, we obtain

$$R(\tau,d) = \mathbb{E}^{\pi_0}\left[\int_0^\tau c_1\mathbb{E}[1_{X_t\neq\mu}|\mathscr{Y}_t]\,dt + \mathbb{E}\left[1_{X_\tau=\mu} + c_2 1_{X_\tau\neq\mu}(X_\tau-d)^2 + c_3 1_{X_\tau\neq\mu}|1_{X_\tau\ge\mu} - 1_{d\ge\mu}|\,\middle|\,\mathscr{Y}_\tau\right]\right].$$

Furthermore, to work under the more convenient reference measure $\mathbb{P}_0$ we use the Radon-Nikodym process L to re-write the performance criterion as

$$R(\tau,d) = \mathbb{E}_0^{\pi_0}\left[\mathbb{E}_0[L_\tau 1_{X_\tau=\mu} + c_1\int_0^\tau L_s 1_{X_s\neq\mu}\,ds + L_\tau H(d,X_\tau)\big|\mathscr{Y}_\tau]\right]. \tag{8}$$

This leads us to define for any smooth bounded function $f:\mathbb{R}_+\to\mathbb{R}_+$

$$\rho_t f \triangleq \mathbb{E}_0[L_t f(X_t)\big|\,\mathscr{Y}_t]. \tag{9}$$

Let $\mathscr{M}(\mathbb{R}_+)$ be the space of all σ-finite positive measures on $\mathbb{R}_+$. Then $\rho_t \in \mathscr{M}(\mathbb{R}_+)$ is defined implicitly through (9) and, as a function of t, ρ is a measure-valued process known as the unnormalized density of X given observed $\mathscr{Y}_t$. Then we can re-write (8) as

$$U(\pi_0) = \inf_{\tau\in\mathscr{S}} \mathbb{E}_0^{\pi_0}\left[\int_0^\tau \rho_s H_1\, ds + \rho_\tau H_2\right], \tag{10}$$

with $H_1(x) \triangleq c_1 1_{\{x\neq\mu\}}$, and

$$H_2(x) \triangleq 1_{\{x=\mu\}} + \inf_{d\in D} H(d,x).$$

Above we have made use of the fact that conditional on the alarm time τ^*, the risk-minimizing announcement d^* can be found by optimizing the penalty function $d \mapsto H(d, X_{\tau^*})$, effectively removing d from (10).

The new problem (10) has considerably simplified the original (6) by switching to the reference measure $\mathbb{P}_0$. Indeed, under $\mathbb{P}_0$ the driving processes N, Y are of particularly simple form, decoupling observations and filtering. Moreover, as we will see below, the unnormalized filter ρ possesses linear dynamics as well. To analyze (10) it is now necessary to understand the dynamics of $\rho_t H_1$ and $\rho_t H_2$. The following lemma gives their explicit description through the well-known Zakai equation, see e.g. [1, Theorem 3.24], [33, p. 270], [17].

Lemma 1 (Zakai equation). *Let $(t,x) \mapsto f_t(x)$ be a bounded Borel function on $\mathbb{R}_+^2$. Then,*

$$\rho_t f_t = \pi_0(f_0) + \int_0^t \rho_s(Af)\, ds + \int_0^t \rho_s(f_s h_s)\, dY_s + \frac{1}{\mu}\int_0^t \rho_{s-}(f_s(\Lambda_{s-}) - \mu)\,(dN_s - \mu\, ds), \tag{11}$$

where A is the generator of X given in (2).

Lemma 1 expresses the evolution of $\rho_t f_t$ in terms of stochastic integrals with respect to the driving processes Y and N. We recall that both Y and $(N_t - \mu t)$ are $\mathbb{P}_0$-martingales, so that the drift of $\rho_t f_t$ is given by the first term in (11).

For completeness we briefly recall the normalized filter π described in the Introduction. Let $\mathscr{P}(\mathbb{R}_+) \subset \mathscr{M}(\mathbb{R}_+)$ be the set of all probability measures on $\mathbb{R}_+$. The normalized filter $\pi_t \in \mathscr{P}(\mathbb{R}_+)$ satisfies $\pi_t f_t = \mathbb{E}[f_t(X_t)|\mathscr{Y}_t]$ for any smooth $f_t(x)$. By Bayes formula, we have the Kallianpur-Striebel relationship

$$\pi_t f_t = \frac{\mathbb{E}_0[L_t f_t(X_t)|\mathscr{Y}_t]}{\mathbb{E}_0[L_t|\mathscr{Y}_t]} = \frac{\rho_t f_t}{\rho_t 1}. \tag{12}$$

The analogue of (11) is the Kushner Stratonovich equation:

$$\begin{aligned}\pi_t f_t = \pi_0 f_0 + \int_0^t \pi_s(Af)\, ds + \int_0^t (\pi_s(f_s h_s) - \pi_s(f_s)\pi_s(h_s))\,(dY_s - \pi_s(h_s)\, ds) \\ + \frac{1}{\mu}\int_0^t (\pi_{s-}(f_s\Lambda_s) - \pi_{s-}(f_s)\pi_{s-}(\Lambda_s))\,(dN_s - \pi_s(\Lambda_s)\, ds).\end{aligned} \tag{13}$$

Above, the driving processes are $(\mathbb{P}, \mathscr{Y})$-martingales known as the innovation processes $Y - \int \pi_s(h_s)\, ds$ and $N - \int \pi_s(\Lambda_s)\, ds$.

3.1 Conditional Moments

The filters $\pi_t \in \mathscr{P}(\mathbb{R}_+)$ and $\rho_t \in \mathscr{M}(\mathbb{R}_+)$ in general do not admit any finite-dimensional Markovian sufficient statistics. Indeed, the equations (11) and (13) are not autonomous and to e.g. compute $\pi_t f$ we must also compute $\pi_t h$, $\pi_t \chi$, $\pi_t(Af), \pi_t(f\chi)$ and $\pi_t(fh)$. Rather than directly working with the measure-valued π (or ρ), the original method of Kalman-Bucy consisted of deriving the dynamics of the conditional moments of X. For comparison we recall the following

Lemma 2. *Define the unnormalized conditional moments $\phi_0(t) = \mathbb{E}_0[L_t 1_{\{\theta \le t\}} | \mathscr{Y}_t]$ and for $k = 1, 2, \ldots$, $\phi_k(t) = \mathbb{E}_0[L_t X_t^k | \mathscr{Y}_t]$. Let us assume that $h(x) = \alpha x$. Then ϕ_k satisfy*

$$\begin{cases} d\phi_0(t) = \alpha\phi_1(t)\, dY_t + (\phi_1(t) - \mu)(dN_t - \mu\, dt) + \lambda_t(1 - \phi_0(t))\, dt \qquad \textit{and} \\ d\phi_k(t) = \alpha\phi_{k+1}(t)\, dY_t + (\phi_{k+1}(t) - \mu\phi_k(t))(dN_t - \mu\, dt) + \lambda_t(1 - \phi_0(t))(\mathbb{E}[\chi^k] - \mu^k)\, dt. \end{cases} \tag{14}$$

Proof. This result easily follows from the Zakai equation (11) after noting that for $f(x,t) = 1 - 1_{\{x=\mu\}}$, $Af(X_t,t) = A1_{\{\theta \le t\}} = \lambda_t 1_{\{X_t = \mu\}}$, so that $\rho_t(Af) = \lambda_t(1 - \phi_0(t))$. Similarly for $f(x) = x^k$, we have $\rho_t(Af) = \rho_t(\lambda_t(X^k - \mu^k)1_{\{X_t=\mu\}}) = \lambda_t \mathbb{E}[\chi^k - \mu^k](1 - \phi_0(t))$ since conditional on $X_t = \mu$, $\chi \perp\!\!\!\perp \mathscr{Y}_t$. A similar result for the normalized moments and Wiener-only observations was given in [31, Theorem 4.2]. □

Lemma 2 demonstrates that the (unnormalized) conditional moments do not constitute a closed system of equations. Indeed, the evolution of the k-th conditional moment ϕ_k depends on ϕ_{k+1} so that the entire infinite sequence $(\phi_k(t))_{k=0}^{\infty}$ is necessary to solve (14). The system in (14) can be closed under some special circumstances, e.g. a finitely-supported distribution $F(\cdot)$ of χ that implies that X_t has a finite state space. In that case, [3, Corollary 3.3] explicitly shows the closure equations satisfied by the conditional moments. Alternatively, artificial closure equations can be introduced. For instance, a Gaussian-type filter (see [31, Section 6]) can be obtained by setting the conditional centered third moment to be zero, $\mathbb{E}[(X_t - \mathbb{E}[X_t|\mathscr{Y}_t])^3 | \mathscr{Y}_t] = 0 \forall t$ which leads to an expression relating ϕ_3 to ϕ_0, ϕ_1 and ϕ_2 and a 3-dimensional sufficient statistic. However, quantification of the corresponding error is difficult. For this reason, in our approach we will be working directly with (11) rather than (14).

The Zakai equation (11) can also be interpreted in its strong form, namely as a stochastic partial differential equation for the density of the measure ρ_t (or π_t). We do not detail here the additional technical assumptions needed for such a representation, and only recall that in the case where X_t has finite-state space E, the measure ρ_t reduces to a vector on $\mathbb{R}^{|E|}$ and we can directly derive its evolution, see [33, sec. 7.3] or [1, remark 3.26]. More precisely, identify $E = (e_1, e_2, \ldots)$ and denote by $\tilde{H} = \mathrm{diag}(h)$ where $h_i = h(e_i)$ and $\tilde{\Lambda} = \mathrm{diag}(\lambda)$ with $\lambda_i = \Lambda(e_i)$. Then the column vector $\tilde{\rho} = (\tilde{\rho}_1, \ldots, \tilde{\rho}_{|E|})$ defined by $\tilde{\rho}_i \triangleq \mathbb{E}_0[L_t 1_{\{X_t = e_i\}} | \mathscr{Y}_t]$ solves the linear stochastic differential equation

$$\tilde{\rho}_{\mathbf{t}} = \pi_0 + \int_0^t A^T \tilde{\rho}_{\mathbf{s}}\, ds + \int_0^t \tilde{H}\tilde{\rho}_{\mathbf{s}}\, dY_s + \int_0^t \frac{(\tilde{\Lambda} - \mu)}{\mu} \tilde{\rho}_{\mathbf{s}-}\, (dN_s - \mu\, ds), \tag{15}$$

where A^T is the transpose of the transition matrix of X. Equation (15) gives a complete description of the conditional distribution of X_t via the (unnormalized) posterior probabilities of each state e_i and is known as the Wonham filter.

3.2 Particle Filters

As we have seen in Lemma 2, for continuous post-disorder distribution F, the filter ρ does not admit finite-dimensional sufficient statistics. To obtain a representation of ρ that is amenable to computational solution, we approximate ρ through a purely atomic measure. Namely, we replace the diffuse ρ with a particle cloud $\rho^{(n)}$ of n particles. The dynamics of $\rho^{(n)}$ are described through a *sequential Monte Carlo* procedure that is summarized via the two main steps of mutation and selection. We refer to [1, 5, 15] for general references on SMC methods in nonlinear filtering of Markov processes.

Our description below is based on the basic continuous-time filter in [1, Ch. 9]. Fix $n > 0$; the particle system $\rho^{(n)} = (\rho_t^{(n)})_{0\le t\le T}$ consists of a collection of n weights $a^j(t)$ and corresponding locations $x^j(t)$, $j = 1,\dots,n$. Before giving the full description of $\rho^{(n)}$ in (20), we first describe the evolution of the particles.

The particles are initialized as $a^j(0) = 1$ and $x^j(0) \sim \pi_0$, i.i.d. Let $\tau_1, \tau_2, \dots$, describe the selection or resampling times. These could be deterministic and lie on a grid $\tau_m = m\delta$ for some grid size δ or be stochastic, e.g. $\tau_k = \sigma_k$ the k-th arrival time of the counting process N, or be otherwise adaptive. Between selection times the particles undergo independent mutation according to the dynamics of X in (1). Thus, if $x^j(\tau_m) \neq \mu$ then $x^j(t) = x^j(\tau_m)$ for all $t \ge \tau_m$ and if $x^j(\tau_m) = \mu$, then for any $t \in (\tau_m, \tau_{m+1})$ we have

$$x^j(t) = \begin{cases} \mu, & \text{with prob.} \quad \frac{1-G(t)}{1-G(\tau_m)}, \\ \chi^j_{m+1}, & \text{with prob.} \quad \frac{G(t)-G(\tau_m)}{1-G(\tau_m)}, \end{cases} \tag{16}$$

where χ^j_{m+1} are i.i.d. with distribution $F(\cdot)$ and each x^j is independent of other particles. In the simulation procedure, given τ_m and $x^j(\tau_m) = \mu$, we first generate the post-disorder location χ^j_{m+1}, and then the particle's disorder time θ^j which has distribution $G|\theta > \tau_m$. Given this pair, the path of $x^j(t)$ on $[\tau_m, \tau_{m+1})$ is $x^j(t) = \mu 1_{t<\theta^j} + \chi^j_{m+1} 1_{t\le\theta^j}$, which clearly matches (1). Note that a new θ^j is simulated for each interval $[\tau_m, \tau_{m+1})$ as long as $x^j(\tau_m) = \mu$.

The corresponding weight a^j is assigned as

$$\begin{aligned} a^j_{m+1} \triangleq a^j(\tau_{m+1}-) &= \exp\left\{ \int_{\tau_m}^{\tau_{m+1}} \log\left(\frac{x^j(s-)}{\mu}\right) dN_s - \int_{\tau_m}^{\tau_{m+1}} \left(x^j(s)-\mu\right) ds \right\} \\ &\qquad \times \exp\left\{ \int_{\tau_m}^{\tau_{m+1}} h(x^j(s))\, dY_s - \frac{1}{2}\int_{\tau_m}^{\tau_{m+1}} h(x^j(s))^2 ds \right\} \\ &= \left(\prod_{\tau_m < t \le \tau_{m+1}: \Delta N_t = 1} \frac{x^j(t-)}{\mu} \right) \\ &\qquad \times \exp\left\{ \int_{\tau_m}^{\tau_{m+1}} h(x^j(s))\, dY_s - \int_{\tau_m}^{\tau_{m+1}} \left(\frac{1}{2} h(x^j(s))^2 + x(s)^j - \mu \right) ds \right\}. \end{aligned} \tag{17}$$

Since $x^j(\cdot)$ is piecewise constant, both terms can be computed and simulated exactly. The first term involving the Stieltjes integral with respect to N can be represented as a discrete product in terms of the arrival times of N. The second term with a stochastic integral with respect to Y can be written directly in terms of the Gaussian increments of Y. Thus, from a simulation point of view to obtain the

collection (a^j_{m+1}) it suffices to simulate the arrival times of the homogenous Poisson process $(N,\mathbb{P}_0)$ and the increments $Y_{\theta^j}-Y_{\tau_m}$ and $Y_{\tau_{m+1}}-Y_{\theta^j}$, $j=1,\dots,n$ of the Wiener process $(Y,\mathbb{P}_0)$.

Taking the normalized weights

$$\overline{a}^j_{m+1} \equiv \overline{a}_j(\tau_{m+1}-) \triangleq \frac{a^j_{m+1}}{\sum_{k=1}^n a^k_{m+1}}, \qquad j=1,\dots,n, \tag{18}$$

we apply a branching or resampling procedure at τ_{m+1}, such that the particle at x^j_{m+1} produces o^j_{m+1} offspring. Each offspring inherits the parent's location $x^j(\tau_{m+1}-)$ and, denoting by $\{a\}=a-\lfloor a\rfloor$ the fractional part of $a\in\mathbb{R}_+$, the integers o^j_{m+1} satisfy $\sum_{j=1}^n o^j_{m+1}=n$ with

$$o^j_{m+1} = \begin{cases} \lfloor n\overline{a}^j_{m+1}\rfloor, \text{ with prob.} & 1-\left\{n\overline{a}^j_{m+1}\right\}, \\ 1+\lfloor n\overline{a}^j_{m+1}\rfloor, \text{ with prob.} & \left\{n\overline{a}^j_{m+1}\right\}. \end{cases} \tag{19}$$

Note that $(o^j_{m+1})_{j=1}^n$ are therefore not independent. The branching procedure in (19) is detailed in [1, pages 226-230] and assures that the number of offspring for each particle has minimal variance while keeping $\mathbb{E}[o^j_{m+1}]=n\overline{a}^j_{m+1}$. After the branching/re-sampling step, all weights are then reset to $a^j(\tau_{m+1})=1$ and the mutation-selection loop is restarted.

Denote by

$$\xi_m \triangleq \prod_{\ell=1}^m \left(\frac{1}{n}\sum_{j=1}^n a^j(\tau_\ell-)\right).$$

The right-hand-side terms above correspond to the total sum of weights just before each branching, which measures the fitness score of the overall particle cloud with respect to the observations on $[\tau_\ell,\tau_{\ell+1})$. Then the unnormalized empirical measure $\rho^{(n)}$ is given by

$$\rho^{(n)}_t \triangleq \frac{\xi_m}{n}\sum_{j=1}^n a^j(t)\delta_{x^j(t)}(\cdot), \qquad \text{for} \quad t\in[\tau_m,\tau_{m+1}), \tag{20}$$

so that the filter of $f(X_t)$ becomes

$$\mathbb{E}_0[L_t f(X_t)|\mathscr{Y}_t] \simeq \rho^{(n)}_t f = \frac{\xi_m}{n}\sum_{j=1}^n a^j(t) f(x^j(t)). \tag{21}$$

Note that $\xi_m \simeq \rho_{\tau_m}\mathbf{1}$ essentially corresponds to posterior likelihood. Accordingly, the normalized particle measure is simply

$$\pi^{(n)}_t \triangleq \sum_{j=1}^n \overline{a}^j(t)\delta_{x^j(t)}(\cdot) \qquad \text{and} \qquad \mathbb{E}[f(X_t)|\mathscr{Y}_t] \simeq \sum_{j=1}^n \overline{a}^j(t) f(x^j(t)).$$

Rigorous error analysis of a continuous-time particle filter for a diffusion signal X is given in [1, Sec. 9.4]. A key step is to establish uniform bounds on the particle weights in order to control the empirical filter $\rho^{(n)}_t f$, see eq. (9.40) in [1]. We straightforwardly may obtain the same bound in our additive jump-diffusion setup:

Lemma 3. *Suppose that h and Λ are bounded and that resampling takes place every δ time units. Then for $p=1,2,3,4$ there exist constants $C_1(p)$ that depend only on $\|h\|_\infty, \|\Lambda\|_\infty$ and T, such that for all*

$s \le T$

$$\sup_{n\ge 0}\max_{j\le n}\sup_{s\in[0,T]}\mathbb{E}_0\left[(\xi^{(n)}_{\lfloor s/\delta\rfloor}a^{(n),j}(s))^p\right]\le C_1(p),\qquad p=1,2,3,4. \tag{22}$$

Proof. Working with the definition of $a^{(n),j}(t)$ in (17) we have for $t<\delta$

$$\begin{aligned}\mathbb{E}_0[(a^{(n),j}(t))^p] &= \mathbb{E}_0\left[\mathscr{E}_1\cdot\exp(\mathscr{E}_2)\right] \qquad \text{with}\\ \mathscr{E}_1 &= \left(\prod_{s\le t:\Delta N_s=1}\left(\frac{\Lambda(X_s)}{\mu}\right)^p\right)\cdot e^{-p\int_0^t(\Lambda(X_s)-\mu)\,ds};\\ \mathscr{E}_2 &= p\int_0^t h(X_s)\,dW_s-\frac{p}{2}\int_0^t h^2(X_s)\,ds.\end{aligned}$$

Using the boundedness of Λ and h we find

$$\begin{aligned}\mathscr{E}_1 &\le \left(\frac{\|\Lambda\|_\infty}{\mu}\right)^{pN_t}e^{\mu pt},\\ \mathscr{E}_2 &\le p\int_0^t h(X_s)\,dW_s-\frac{p^2}{2}\int_0^t h^2(X_s)\,ds+\frac{p^2-p}{2}t\|h\|^2_\infty.\end{aligned} \tag{23}$$

Using the independence of N and Y and recognizing the first term on the second line above as the Wiener exponential martingale we obtain

$$\mathbb{E}_0[(a^{(n),j}(t))^p]\le \mathbb{E}_0\left[\left(\frac{\|\Lambda\|_\infty}{\mu}\right)^{pN_t}\right]e^{t(\mu p+\frac{p^2-p}{2}\|h\|^2_\infty)} \tag{24}$$

$$\le \exp\left(\delta\{\frac{\|\Lambda\|^p_\infty}{\mu^{p-1}}-\mu+\mu p+\frac{p^2-p}{2}\|h\|^2_\infty\}\right)\le\exp(C(p)\delta). \tag{25}$$

Thanks to the independent mutations of each particle and the resetting property of the resampling, it follows that for any m

$$\mathbb{E}_0\left[(\xi_m)^p\right]\le\prod_{\ell=1}^m\mathbb{E}_0\left[\sum_j\frac{1}{n}a^{(n),j}(\delta-)^p\right]\le\prod_{\ell=1}^m e^{C(p)\delta}=e^{C(p)m\delta}\le e^{C(p)T}.$$

Combining with (25) and using the Cauchy-Schwartz inequality the bound (22) follows. We observe that the assumed boundedness of Λ and h can be slighted relaxed, as long as sufficiently high exponential moments exist. □

Lemma 3 established uniform upper bounds on the particle weights on the full horizon $[0,T]$. Using it and straightforwardly adapting the proof of [1, Prop. 9.14] we obtain the $\mathscr{O}(1/\sqrt{n})$ convergence rate of the particle filtering algorithm.

Proposition 2. *There exists a constant $C_2(T)$ such that for any bounded function $f\in C_b(\mathbb{R})$*

$$\mathbb{E}_0\left[\{(\rho^{(n)}_t(f)-\rho_t(f)\}^2\right]\le\frac{C_2(T)}{n}\|f\|^2_\infty,\qquad\forall t\le T. \tag{26}$$

3.3 Particle Degeneracy

The post-disorder rate χ is an important component of the filtering/detection problem. Being a constant, its estimation is in many ways similar to estimating a parameter in a filtering model. This is in contrast to the filtering of the disorder indicator $1_{\{\theta\le t\}}$ that is inherently dynamic. It is well known that use of SMC methods for parameter estimation may be unstable and prone to error accumulation. Indeed, since each particle generates its copy $\chi^j = x^j(\theta^j)$ just once, the full particle cloud does not explore the state space of χ beyond this first mutation. Moreover, if re-sampling is performed regularly, with positive probability some particles will be discarded at each re-sampling step, so that as $t\to\infty$, the empirical measure $\rho^{(n)}$ will collapse to a single atom at some location $\check{x}\in\{x^1(\theta^1),x^2(\theta^2),\dots,x^n(\theta^n)\}$, rather than converging to the true location χ.

An early solution for such particle degeneracy proposed by [20] was to allow for *artificial evolution* of particles. Namely, introduce "roughening penalties", so that $x^j(t)$ remains dynamic for all time points, even for $t\ge\theta^j$. More precisely, take

$$x^j(\tau_m+) = x^j(\tau_m)+\varepsilon^j, \qquad \varepsilon^j\sim\mathcal{N}(0,v_m),$$

for j such that $\theta^j > t$, i.e. $x^j(\tau_m)\neq\mu$. These artificial disturbances ε^j can be justified through Gaussian kernel smoothing of the posterior measure $\rho^{(n)}$ that replace the point-masses in (21) with

$$\widetilde{\rho_t^{(n)}} f \triangleq \frac{\xi_m}{n}\sum_{j=1}^{n} a^j(t)\int_{\mathbb{R}} f(y)K(x^j(t);dy),$$

where $K(x,dy)$ is a smooth kernel satisfying $\int yK(x,dy)=x$, $\int y^2K(x,dy)<\infty$. A common choice is a Gaussian kernel $K(x,dy)=\frac{1}{\sqrt{2\pi v_m}}\exp(-\frac{(x-y)^2}{2v_m})\,dy$ with a small variance parameter v_m. Gordon and Salmond [20] suggested $v_m = Var(1_{\{X_{\tau_m}\neq\mu\}}X_{\tau_m}|\mathscr{Y}_{\tau_m})$, proportional to the true variance of the signal. With this additional kernel approximation, the branching procedure is re-interpreted as sampling n times with replacement from the smoothed empirical measure $\widetilde{\rho_{\tau_m-}^{(n)}}$ to obtain the new locations $x^j(\tau_m)$, $j=1,\dots,n$.

Such kernel convolution clearly introduces additional error by increasing the empirical variance of the particle cloud, since now the variance in $\widetilde{\rho}_n$ is bigger than the true variance of $X_{\tau_m-}|\mathscr{Y}_{\tau_m}$. This can be overcome by the method of [22] who suggested kernel *shrinkage*, namely setting

$$x^j(\tau_m+) = a\cdot x^j(\tau_m)+\varepsilon^j, \qquad \varepsilon^j\sim\mathcal{N}\left((1-a)\overline{x}(\tau_m),(1-a^2)v_m\right), \tag{27}$$

for $x^j(\tau_m)\neq\mu$, where $\overline{x}(\tau_m)=\frac{\sum_j x^j(\tau_m)1_{x\neq\mu}(x^j(\tau_m))}{\sum_j 1_{x\neq\mu}(x^j(\tau_m))}$ is the normalized average post-disorder rate among all particles, v_m is the variance associated with $\overline{x}(\tau_m)$ and $a\in[0.95,0.999]$ is the shrinkage parameter. Thus, particles experience random shocks that direct them slightly towards the empirical mean $\overline{x}(\tau_m)$ of the other post-disorder particles. The Liu and West method [22] ensures that no additional variance is added to the particle cloud, so that both the empirical mean and empirical variance are unaffected by the artificial mutation. We find that this procedure provides a good means to induce further exploration of the state space of χ by the particle cloud in the latter stages $t>\mathbb{E}_0[\theta]$ when most particles become static.

A second important issue we confront is the resampling *frequency*. On the one hand, resampling combats weight degeneracy which in particular allows pre-disorder particles to explore the θ-

distribution. On the other hand, resampling adds extra noise through the sampling variance. We partly addressed the latter issue by using the minimal-variance branching procedure. From (17) we see that if $h(X_t)$ is small, then the main source of weight variance is the Poisson process N, with particle weights increasing by a factor proportional to χ at each σ_k. This suggests adaptive resampling based on the arrival times of N, e.g. $\tau_k = \sigma_k$. A heuristic resampling rule can be derived by looking at statistics of the weight distribution becoming too low. We refer to [14] for details on possible $a(t)$-statistics, such as effective sample size $ESS = \{\sum_{j=1}^n (\bar{a}^j(t))^2\}^{-1}$, coefficient of variance $CoV = \sqrt{\frac{1}{n}\sum_{j=1}^n \{a^j(t)-1\}^2}$ and entropy $Ent = -\sum_{j=1}^n \bar{a}^j(t)\log \bar{a}^j(t)$. Finally, it is possible to not resample at all. Such scheme with independent particle mutation and adaptive weight assignment is known as the classical sequential importance sampling (SIS) filtering algorithm, see e.g. [5]. SIS filtering is effective in the cases where the post-disorder distribution is discrete, whereby with enough particles it is essentially equivalent to exact filtering according to (15).

A further possibility to reduce particle degeneracy is marginalization of the disorder time using the ideas of [6]. Namely, after each resampling, we split each pre-disorder particle into two copies, with copy 1 conditioned to have no disorder on the next interval $[\tau_m, \tau_{m+1})$, and copy 2 conditioned to have a disorder on $[\tau_m, \tau_{m+1})$. The weights of copy i are given by the corresponding conditional probabilities. In other words, we marginalize the random noise in the simulation of disorder times, similar to classical Rao-Blackwellization procedures. Note that we now would propagate more than n particles and therefore the minimum variance branching procedure no longer applies. Instead, we switch to standard sampling with replacement from $\rho^{(n)}_{\tau_m-}$, using e.g. stratified or systematic resampling procedures from [5, ch. 7.4].

Remark 3. Note that while classical filtering paradigms work on a fixed time horizon $[0,T]$, in our filtering-control problem the effective horizon is $[0,\tau^*]$. If the penalty for detection delay is significant, we can expect that $|\tau^* - \theta|$ is small, alleviating the problem of post-disorder static particle degeneracy.

4 Solving the Optimal Stopping Problem

Let us define the Snell envelope corresponding to (6),

$$U_t \triangleq \operatorname*{essinf}_{\tau \geq t} \mathbb{E}_0 \left[\int_t^\tau \rho_s H_1 \, ds + \rho_\tau H_2 | \mathscr{Y}_t \right],$$

The original problem consists in finding U_0.

Restricting the stopping decision to be made at discrete time instances $\tau \in \{0, \Delta t, 2\Delta t, \ldots, T = M\Delta t\}$, the discrete-time value function, which we continue to denote as U_t, satisfies the recursive equation

$$U_t = \min\left(\rho_t H_2, \rho_t H_1 \Delta t + \mathbb{E}_0[U_{t+\Delta t}|\mathscr{Y}_t]\right), \tag{28}$$

with $U_T = \rho_T H_2$. Therefore, one may compute $U_{m\Delta t}$ for $m\Delta t < T$ by iteratively solving (28) backwards in time. The key step in this approach becomes (approximately) computing the conditional expectation on the right-hand-side of (28). For later use we note that in discrete time the optimal stopping time $\tau^*(t)$ satisfies

$$U_t = \mathbb{E}_0\left[\sum_{s=t}^{\tau^*(t)-1} \rho_s H_1 \Delta t + \rho_{\tau^*(t)} H_2 \big| \mathscr{Y}_t\right];$$

$$\tau^*(t) = t 1_{S_t} + \tau^*(t+\Delta t) 1_{S_t^c}, \qquad S_t \triangleq \{\rho_t H_2 < \rho_t H_1 \Delta t + \mathbb{E}_0[U_{t+\Delta t}|\mathscr{Y}_t]\} \in \mathscr{Y}_t, \tag{29}$$

where S_t^c is the complement of the event S_t.

In our Markov setting and in view of (11), U_t is a function of ρ_t, i.e., a functional on the nonlocally compact space $\mathscr{M}(\mathbb{R}_+)$. Note that since Y and $N_t - \mu t$ are time-homogenous $\mathbb{P}_0$-martingales they only affect U_t through their stationary increments. As explained in Section 3, ρ is generally infinite-dimensional, presenting a computational roadblock. Indeed, we have that the conditional expectation $\mathbb{E}_0[U_{t+\Delta t}|\rho_t] = E(t;\rho_t)$ is a function of the measure-valued ρ_t for some functional $E : \mathscr{M}(\mathbb{R}_+) \to \mathbb{R}$. Evaluating this map E analytically appears hopeless. Here we propose to overcome this challenge by employing a Monte Carlo approach to (28) that additionally combines well with the SMC filtering method for approximating ρ with $\rho^{(n)}$.

4.1 Integrated Algorithm

Our method is based on the Longstaff-Schwartz [23] algorithm for classical optimal stopping problems. Its key idea is to focus on the decision rule, in other words to approximate (29), rather that U_t directly. This is achieved by replacing $U_{t+\Delta t} = \mathbb{E}_0[\sum_{s=t+\Delta t}^{\tau^*(t+\Delta t)-1} \rho_s H_1 \, ds + \rho_{\tau^*(t+\Delta t)} H_2 | \mathscr{Y}_{t+\Delta t}]$ with an *empirical pathwise continuation value* $u_{t+\Delta t}$ and to compute the conditional expectation $\mathbb{E}_0[u_{t+\Delta t}|\mathscr{Y}_t]$ through *cross-sectional regression* of a Monte Carlo collection $u^k_{t+\Delta t}$. This cross-sectional regression approximates via

$$\mathbb{E}_0[U_{t+\Delta t}|\mathscr{Y}_t] \simeq \sum_{i=1}^{r} \alpha^i(t) B_i(\rho_t), \tag{30}$$

where $(B_i(\rho))_{i=1}^r$ are the basis functions and $\alpha^i(t)$ the corresponding regression coefficients. Thus, we replace the conditional expectation operator $\mathbb{E}_0[\cdot|\mathscr{Y}_t]$ (characterized as the L^2-minimizer) with an L^2-projection onto the $span(B_i(\rho_t) : i = 1,\dots,r)$. Comparing the regression prediction $\sum_i \alpha^i(t) B_i(\rho_t)$ and the immediate payoff $\rho_t H_2$ we then construct the approximate stopping rule τ for (29).

The computational algorithm begins by generating a large scenario database, simulating (N,Y) under the reference measure $\mathbb{P}_0$ and computing the approximate filter $\rho^{(n)}$ along each realized path, $k = 1,\dots,K$, using the particle filter of Section 3.2. We then approximate $B_i(\rho_t^k) \simeq B_i(\rho_t^{(n),k})$ and using backward recursion implement (29) by regressing the empirical $(u^k_{t+\Delta t})$ against the simulated $\{B_i(\rho^{(n),k})\}_{k=1}^K$. Namely,

$$S_t^{(K)} = \left\{\rho_t^{(n),k} H_2 < \rho_t^{(n),k} H_1 \cdot \Delta t + \textstyle\sum_{i=1}^r \alpha^{(K),i}(t) B_i(\rho_t^{(n),k})\right\}; \tag{31}$$

$$\tau_t^k = t \cdot 1_{S_t^{(K)}} + \tau^k(t+\Delta t) \cdot 1_{(S_t^{(K)})^c}, \tag{32}$$

$$u_t^k = \textstyle\sum_{s=t}^{\tau_t^k - 1} \rho_s^{(n),k} H_1 \cdot \Delta t + \rho_{\tau_t^k}^{(n),k} H_2, \tag{33}$$

where $\alpha^{(K),\cdot}(t)$ are the empirical regression coefficients from the simulation of size K. Note that the whole algorithm operates under $\mathbb{P}_0$ which allows for easy simulation of (N,Y) and requires only the unnormalized filter $\rho^{(n)}$.

Iterating with $t=(M-1)\Delta t$ down to $t=0$ we eventually obtain an *in-sample estimate*

$$u_0 = \frac{1}{K}\sum_{k=1}^{K} u_0^k \simeq U_0.$$

Since it is Monte Carlo based, u_0 is itself a random variable and is typically slightly biased. To avoid this, we may straightforwardly compute an out-of-sample estimator $\hat{u}_0$. Indeed, it suffices to simulate a new independent set of $(N_t^k, Y_t^k, \rho_t^k)_{k=1}^{K'}$ and moving forward for $t=\Delta t, 2\Delta t, \ldots$, evaluate $\rho_t^k H_1, \rho_t^k H_2$ and the continuation value $\hat{q}_t^k \triangleq \sum_{i=1}^{r} \alpha^{(K),i}(t) B_i(\rho_t^k)$ using the previously computed regression coefficients $\alpha^{(K)}(t)$. Then comparing $\hat{q}^k$ and $\rho_t^{(n),k} H_2$ one can find the first time τ^k and set $\hat{u}_0^k = \sum_{s=0}^{\tau^k-1} \rho_s^k H_1 \cdot \Delta t + \rho_{\tau^k}^k H_2$ to be the corresponding scenario payoff. Finally, average to obtain $\hat{u}_0 = \frac{1}{K'}\sum_{k=1}^{K'} \hat{u}_0^k$, which by construction is a lower bound for U_0.

We summarize the full procedure in pseudo-code in Algorithm 1 in the Appendix. Let us stress the modularity of the code that only needs an implementation of a filtering method (with particle filters requiring only the ability to sample X_t) and a method for finding conditional expectations (with regression Monte Carlo relying on standard least-squares regression). For the basis functions $B_i(\cdot)$, the natural choices include conditional moments $B_i(\rho) = \rho x^\ell$ for integers ℓ, as well as nonlinear functions of quantities appearing in (10), such as $B_i(\rho) = (\rho 1_{\{x=\mu\}})^\ell$ or $(\rho x)^\ell$.

Remark 4. Subject to smoothness conditions, ρ_t has an L^2 density with respect to the Lebesgue measure. In that case, it can be completely characterized through its action against a basis $(\tilde{B}_\ell(\cdot))_{\ell=1}^{\infty}$ for $L^2(\mathbb{R}_+)$, namely as an infinite sequence $(a_\ell)_{\ell=1}^{\infty}$, $a_\ell \triangleq \rho_t(\tilde{B}_\ell)$. Accordingly, the functional E can be viewed as a function on the coordinates (a_ℓ) and the regression against (B_i) above is then a double projection, first truncating to a finite number of $(a_\ell)_{\ell_1}^{r'}$ and then projecting against a truncated basis of $L^2(\mathbb{R}^{r'})$.

4.2 Error analysis

Several steps of the overall algorithm above required approximation. In detail, these were: (i) filtering error between ρ and $\rho^{(n)}$; (ii) projection error of approximating conditional expectation with a projection onto the $span(B_i, i=1,\ldots,r)$; (iii) Monte Carlo error due to finite simulation database; (iv) Snell envelope error from making the wrong stopping decision in (31).

Observe that for the final value function, it is only the first and last errors that are intrinsically important: even if the projection and MC errors have been made, as long as the stopping strategy is correct they are not counted. This key observation is the motivation behind the algorithm's focus on (29) rather than (28).

The filtering error has been quantified in Lemma 2. Unfortunately, no sharp results are known for the Snell envelope error, and existing methods revert to its recursive decomposition into the projection and Monte Carlo errors. This produces very pessimistic theoretical error bounds and in particular implies that the number of simulation paths K should be exponential in number of basis functions r. Additionally, we note that the MC and filtering errors compound each other during the regression

step, due to the resulting error-in-variable that causes attenuated estimates of α_t. Consequently, it is important to minimize the filtering error as much as feasible.

For a given function $g : \mathscr{M}(\mathbb{R}_+) \to \mathbb{R}_+$ denote by $(\mathrm{pr}\circ g)(\rho) \triangleq \sum_i \hat{\alpha}^i B_i(\rho)$ the true projection operator onto $span(B_i, i = 1,\dots,r)$ and by $(\mathrm{pr}^{\mathrm{K}}\circ g)(\rho) \triangleq \sum_i \alpha^i B_i(\rho)$ the empirical regression based on a Monte Carlo sample of ρ_t of size K (thus formally, $\mathrm{pr}^{\mathrm{K}}\circ g$ is a random operator that is a function of the simulation set ω'). Let $\mathscr{R}_t(g) = \mathbb{E}_0[|\,\mathrm{pr}\circ g(\rho_t^{(n)}) - \mathrm{pr}^{\mathrm{K}}\circ g(\rho_t^{(n)})|^2]^{1/2}$ be the L^2 norm of this Monte Carlo error, where the expectation is over the possible simulation sets ω', the particle filter realizations $\rho_t^{(n)}$ and the paths of (N,Y). The following Lemma is a direct consequence of Theorem 1 in [24].

Lemma 4. *Denote by $\hat{U}_t \equiv \hat{U}_t(\rho_t)$ the approximate value function at date t computed by the algorithm of Section 4.1. Then there exist constants C_3, C_4 that depend only on model parameters, such that*

$$\mathbb{E}_0\left[|\hat{U}_0 - U_0|^2\right]^{1/2} \le 4^M \cdot C_1 \cdot \max_{0\le s\le T}\left\{\|\,\mathrm{pr}\circ\hat{U}_s - \hat{U}_s\|_2 + \mathscr{R}_s(\hat{U}_s)\right\} + \frac{C_2}{\Delta t\sqrt{n}}. \tag{34}$$

Typically, the Monte Carlo error is of order $\mathscr{R}_t = \mathscr{O}(K^{-1/2})$, so the algorithm accuracy is order-$(-1/2)$ in both n and K. The Lemma also suggests that the number of paths K should be exponential in Δt.

Remark 5. The classical disorder formulation of Shiryaev [30] is on an infinite horizon, so that there is no specific deadline by which an announcement must be made. For a discrete $F(\cdot)$ and memoryless θ, the results of [3, 25] imply that the finite horizon solution converges to the infinite horizon solution exponentially fast. In general, the convergence should depend on the right tail of the distribution of θ; in most practical applications this would imply (super)-exponential convergence.

From a global perspective, our algorithm employs approximation both at the filtering and optimization steps. The errors in computing Snell envelopes via the RMC method have been studied in [16, 18]. At the same time, the functional stability of Snell envelopes with respect to the underlying process (here we replace ρ with $\rho^{(n)}$) was considered in [27, 8]. Finally, the robustness of approximating (28) was addressed recently in [13]. However, combining these results together is difficult since no usable quantitative results are known for the *rate* of convergence of the Snell envelope.

5 Numerical Examples

5.1 Analysis of Particle Filter

We begin our numerical experiment with a study of the particle filter performance. To check the accuracy of our method we compare it against a case where an exact solution is possible, viz. a discrete $F(\cdot)$ and memoryless disorder time θ. In preparation for examples below, we assume that $\theta \sim Exp(0.5)$, $\mu = 10$ and χ is a mixture of two discretized "uniform" distributions on $[3,8]$ and $[15,25]$ respectively. More precisely,

$$F(\cdot) = \sum_{i=1}^{25} 0.02(\delta_{2.9+0.2i}(\cdot) + \delta_{14.8+0.4i}(\cdot)). \tag{35}$$

We consider just Poisson observations with $\Lambda(X_s) = X_s, h(X_s) \equiv 1$; in that case the exact solution is given by the Wonham filter (15) and can be readily computed with matrix exponentials, see [25]. Thus,

the posterior distribution is given by the 51-dimensional vector $\tilde{\pi}(t)$ with $\tilde{\pi}^j(t) \equiv \mathbb{P}(X_t = x_j|\mathscr{Y}_t)$ and x_j's given in the Dirac masses of (35).

We implement a particle filter with 500 particles and resampling frequency based on CoV criterion. In Figure 1 we show one sample path sequentially tracking the conditional expectation of the post-disorder rate $\chi|\mathscr{Y}_t$ given observations. As expected $\pi_t\chi$ decreases between arrivals and jumps up at arrival times. We observe that with 500 particles, the particle filter does an excellent job of matching the true χ-posterior. Also, despite the relatively high arrival rates, the observation process is still not very informative, so that the inference problem is rather difficult. Even after $T = 5$ time units, the 95% credibility interval has width of over 4; nevertheless the convergence to the true value of χ is evident.

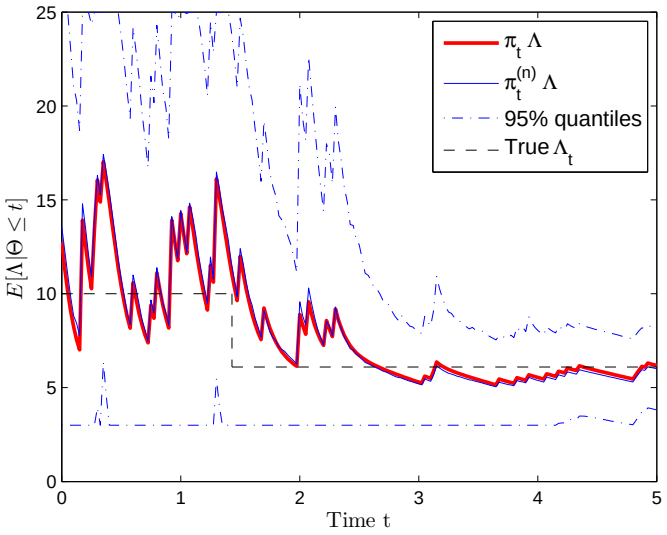

Fig. 1 Performance of the particle filter for the Poisson disorder problem with F given in (35). We show one sample path of the conditional expectation of $\chi|\mathscr{Y}_t$ using the exact and basic particle filter method with $n = 500$. We also indicate the true arrival rate X_t and the 95% quantile interval of the particle filter posterior.

Table 1 provides further details on the performance of the particle filter. The first two columns confirm the $\mathscr{O}(1/\sqrt{n})$ theoretical convergence rate of $\rho^{(n)}$. The last two columns show that in this example because of the discrete $F(\cdot)$ re-sampling increases the average error.

n	Std. Error	Resampling Freq	Std. Error
500	0.103	Every step: $\tau_k = 0.05k$	0.123
1000	0.074	Based on N: $\tau_k = \sigma_{2k}$	0.103
2000	0.053	None (SIS)	0.102

Table 1 Accuracy of the particle filter method. Standard error corresponds to $\mathbb{E}_0\left[\|\tilde{\pi}_T - \tilde{\pi}_T^{(n)}\|_2\right]$, where $\|\cdot\|_2$ is the Euclidean vector norm applied to the 51-dimensional vector of true posterior probability $\tilde{\pi}_T$ computed according to (15) and the approximate normalized vector $\tilde{\pi}_T^{(n)}$ returned by a particle filter with n particles. Standard error based on 4000 simulated paths.

5.2 Example 1

To be able to compare our Bayes risk computation to existing methods, we again consider sole Poisson observations together with a discrete post-disorder distribution $F(\cdot)$ and exponentially distributed θ. In that case, the filter π is finite-dimensional and the Bayes risk problem can be approached using the tools of [3, 25]. Thus, the drift is $h \equiv 1$ and for simplicity, we take a Bernoulli F, making π two-dimensional.

Our objective function is

$$U(\pi_0) = \inf_{\tau,d} \mathbb{E}^{\pi_0} \left[\int_0^\tau 0.2 \cdot 1_{X_t \neq \mu} \, dt + 1_{X_\tau = \mu} + 0.3 |1_{d>\mu} - 1_{\chi>\mu}| \right],$$

where $\mu = 3$ and $\chi \in \{2,4\}$. The rate of disorder is $\theta \sim Exp(0.5)$, $\pi_0 = (0.99, 0.005, 0.005)$ and the horizon is $T = 5$. As a benchmark we solve this problem using our regression Monte Carlo (RMC) method with $n = 4000$ particles and $K = K' = 50{,}000$ paths, with $\Delta t = 0.1$ and no resampling. With a Bernoulli $F(\cdot)$ there is no risk of particle degeneracy and we find that resampling noticeably lowers accuracy. We use the $r = 5$ basis functions, $\{1, \rho_t 1_{\{x=2\}}, \rho_t 1_{\{x=3\}}, \rho_t 1_{\{x=4\}}, \min(\rho_t 1_{\{x=2\}}, \rho_t 1_{\{x=4\}})\}$ which seems a minimal requirement looking at the payoff functions. After many numerical experiments, we find that adding further basis functions has a negligible effect, while using fewer than five leads to a much larger estimate of U_0. See also [4] in this volume for new ideas on how to automate this process. The standard error of the out-of-sample RMC estimate $\hat{u}_0$ is about $2.7 \cdot 10^{-4}$, producing three significant digits. The running time is about 2.5 hours on a 2.33GHz desktop, with the vast majority of time spent on simulating $\rho^{(n)}$.

To compare, we also re-solve this problem following the algorithm in [25] which relies on a recursion on the number of possible arrival times $k_t = \sup\{k : \sigma_k \leq T\}$ and discretizes the state space of the Wonham filter $\tilde{\pi} \in \{(\pi_1, \pi_2, \pi_3) : \pi_i \geq 0, \sum_i \pi_i = 1\}$. This approach is entirely separate from the RMC algorithm herein, utilizing different filtering and optimal stopping techniques and is only applicable for X a time-homogenous finite-state Markov chain. As shown in Table 2, the results are within 2% of each other, which is acceptable. The main error seems to be due to the time-discretization of the integral $\int_0^\tau \rho_s H_1 \, ds$.

Table 2 shows that in fact $n = 1000$ particles produce essentially the same answer, while decreasing the number of paths has mainly impact on the MC standard error of $\hat{u}_0$. Figure 2 shows the distribution of the optimal stopping time τ^*. We can see that τ^* varies quite a lot across different scenarios, with some paths stopping as early as $\tau^* = 1$, and others stopping as late as $\tau^* = 4$. Note that all paths stop before $T = 5$ indicating that the finite horizon constraint is irrelevant. The interpretation of the behavior of τ^* is difficult, since Figure 2 shows its distribution under the reference $\mathbb{P}_0$ (where all the simulations take place), while true observations will be under $\mathbb{P}$.

5.3 Example 2

Our second example is motivated by a financial application. It has joint Poisson and Wiener observations and more involved cost functions. We consider a manager that wishes to ramp up manufacturing of a product with demand $X = \{X_t\}$. To do so, she must set the production (supply) level d_t. Initially demand and supply are in equilibrium $d_0 = X_0 = \mu$, but at time θ demand is expected to ramp up to level χ. If the supplies are too high $d_t > X_t$, over-production costs are incurred; if $d_t < X_t$, opportunity

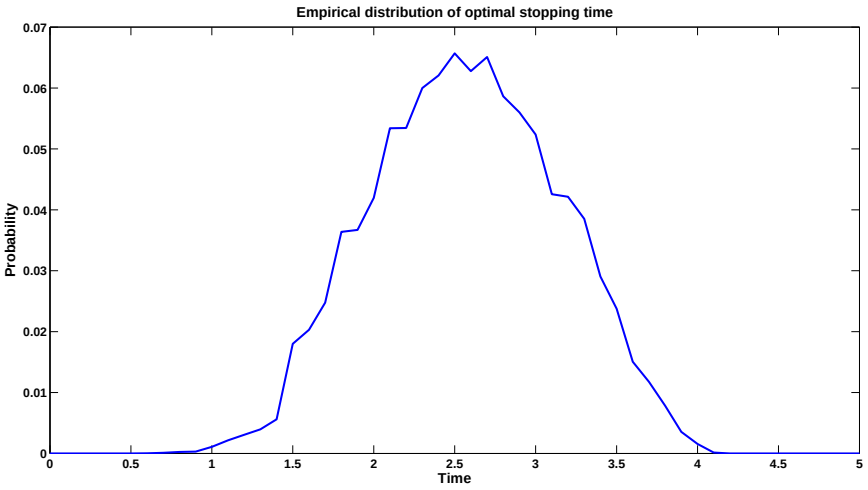

Fig. 2 Distribution of the empirical optimal stopping time τ^* under the reference measure $\mathbb{P}_0$. Out of sample simulation with $K' = 32,000$ and $\Delta t = 0.1$.

Case	Bayes Risk $U(\pi_0)$	MC Std. Error
LS08	0.532	–
RMC	0.542	$2.7 \cdot 10^{-4}$
$n = 500$	0.547	–
$K = 5000$	0.548	$19 \cdot 10^{-4}$
$K = 10000$	0.543	$9.1 \cdot 10^{-4}$
$K = 20000$	0.543	$7.7 \cdot 10^{-4}$

Table 2 RMC Algorithm solution for Example 1, compared to the HMM method of LS08 [25]. We also compare the sensitivity of our method to the number of particles n and number of paths K. The default case has $n = 4000$, $K = 50,000$ and $r = 5$. Monte Carlo standard error is based on 20 independent runs of the RMC algorithm.

costs of unmet demand result. We assume that $\chi > \mu$ a.s., so the disorder in terms of X is always upwards and $\theta < \tau$ leads to under-production. Based on her information so far, the manager must select a time τ to adjust her production levels to match the new demand χ so as to minimize these costs.

Demand is observed through two observation channels. Channel $\tilde{Y}$ consists of market demand for a related product, $d\tilde{Y}_t/\tilde{Y}_t = \alpha X_t\,dt + \sigma dW_t$, where $(W,\mathbb{P})$ is a Wiener process. Letting $Y_t = \frac{1}{\sigma}\log\tilde{Y}_t + \frac{1}{2}\sigma t$ we have $dY_t = \frac{\alpha}{\sigma}X_t\,dt + dW_t$ matching the setup in Proposition 1 with $h(X_t) = \frac{\alpha}{\sigma}X_t$. Channel N monitors economic events that are positively correlated with X and occur at frequency $\Lambda(X_t) = (X_t)^2$, with $N \perp\!\!\!\perp W$.

For the new demand level χ we take a truncated mixture of two normal distributions,

$$\begin{cases} \chi \sim \mathcal{N}(5, 0.5) \vee 3 & \text{with prob. } 0.6; \\ \chi \sim \mathcal{N}(4, 0.25) \vee 3 & \text{with prob. } 0.4. \end{cases}$$

and consider (non-symmetric) costs that depend on the posterior error in identifying χ:

$$\inf_{\tau\in\mathcal{S}, d\in\mathbb{R}_+} \mathbb{E}\left[\int_0^\tau c_1 1_{\{X_t\neq\mu\}}\,dt + 1_{\{X_\tau=\mu\}} + c_{21}1_{\{X_\tau>d\}}(X_\tau - d) + c_{22}1_{\{X_\tau<d\}}(d - X_\tau)\right].$$

Recall that in this case, the optimal new demand level d_τ^* is the $c_{22}/(c_{21}+c_{22})$ quantile of X_τ: $\mathbb{P}(X_\tau \leq d_\tau^* | \mathscr{Y}_\tau) =: \frac{c_{22}}{c_{21}+c_{22}}$. As far as we know, there are no nonlinear filtering models that admit sufficient statistics for posterior quantiles.

The respective costs are taken as $c_1 = 1, c_{21} = \frac{1}{2}, c_{22} = \frac{3}{2}$. The original rate of the Poisson process N is $\mu = 3^2$, and the horizon is $T = 3$. Finally, $\theta \sim Weibull(0.5, 1.5)$ so that the θ-hazard rate $\lambda(t) = 3\sqrt{2t}$ increases as the square-root of time passed.

Table 3 compares the resulting Bayes risk as we vary the observation schemes, namely the informativeness of the diffusion observation $\bar{\alpha} := \frac{\alpha}{\sigma}$ (see a related comparison for a 2-state X in [12, Figure 9]). Larger α makes the diffusion drift $h(X_s) = \bar{\alpha} X_s$ more sensitive to the signal level and therefore aids disorder detection. Therefore, the value function U is decreasing in $\bar{\alpha}$. We see that this effect is quite dramatic and nonlinear, with $\bar{\alpha} = 2$ reducing Bayes risk by nearly 70% compared to $\bar{\alpha} = 0$.

The presence of Wiener observations makes the variance of the likelihood weights ξ_t grow exponentially over time which affects all the unnormalized filtered quantities. At the same time, the median shrinks exponentially fast. For instance with $\alpha = 1$ we have that the median of $\rho_T^k 1$ is $\sim 10^{-8}$, while $\max_{k \leq K} \rho_T^k 1 > 10^4$. Such a large spread presents numerical round-off concerns for applying cross-sectional regression to find $\alpha^{(K)}(t)$. Accordingly, we apply regression on the normalized continuation values $u^k \cdot (\rho_t^k 1)^{-1}$ using the normalized basis functions $B_i(\pi_t^k)$. One may also apply local least squares or nonlinear regression to overcome the different scales across scenarios.

Case	Bayes risk $\hat{u}_0$
$\bar{\alpha} = 0$	0.78
$\bar{\alpha} = 0.25$	0.77
$\bar{\alpha} = 0.5$	0.76
$\bar{\alpha} = 1$	0.62
$\bar{\alpha} = 2$	0.26

Table 3 Bayes risk and average optimal stopping time for different jump-diffusion observation schemes. All the results are based on $K = K' = 32,000$ paths with $\Delta t = 0.04$ and particle filters implementing the Liu-West adjustment moves with $n = 1000$ particles, and $a = 0.99$.

6 Extensions

The algorithm in Section 4.1 is certainly computationally intensive and it would be desirable to find further speed/efficiency improvements. However, at this stage there are few alternatives to handle a problem like Example 2 above. Most existing methods, such as those in [25, 12] for finite HMMs observed via point processes, are designed to handle low-dimensional optimal stopping problems and would not apply. Certainly replacing the continuous distribution F in Example 2 with a two- or three-point discrete approximation is a poor choice. A quantization approximation of ρ was investigated by [29] in a simple diffusion setting, but it is arguably as computationally expensive as ours. An interesting alternative consisting in improving the stochastic mesh method for very high-dimensional problems, was very recently proposed by [13]. Alternatively, one can consider other curve-fitting statistical tools, such as smoothing splines, local regression or convex regression, to refine the approximation of the set

S_t in (31). In this vein, importance sampling techniques to generate more Monte Carlo samples $\rho_t^{(n),k}$ in the region of interest, namely close to the boundaries of S_t, also merit investigation.

A significant advantage of our simulation-based approach is its flexibility which could be exploited to handle many extensions of the basic model (1)-(3).

6.1 Compound Poisson Process Observations

In many practical applications, it is useful to consider a richer structure of counting process observations. Namely, we may replace the simple Poisson process N with a compound Poisson process, or more generally a marked point process. Hence, we now assume that N consists of a double sequence $(\sigma_1, Z_1; \sigma_2, Z_2, \dots)$ where $0 < \sigma_1 < \sigma_2 < \dots$ are the arrival times and $Z_k \in \mathscr{Z}$ are the corresponding marks. Similar to the intensity of N being driven by the signal X, we assume that the mark distribution $f_Z(dz;x)$ depends on X. As a simple example, consider Gaussian marks with mean X_t, $\mathbb{P}(Z_k \in A | \sigma_k = t, \mathscr{F}_t) = \int_A \frac{1}{\sqrt{2\pi}} \exp(-(z-X_t)^2/2) dz$, $A \in \mathscr{B}(\mathbb{R})$.

Let us assume that all the distributions $f_Z(\cdot;x)$ are absolutely continuous with respect to some reference measure $f_Z^0(\cdot)$ (e.g. the standard Gaussian distribution in the example above). Then we may repeat the construction of the reference probability measure $\mathbb{P}_0$ by starting with a canonical Poisson random measure $(N, \mathbb{P}_0)$ with compensator $\nu(dt;dz) = \mu dt f_Z^0(z) dz$, i.e. the process $N([0,t] \times \mathscr{Z}) - \int_{(0,t]} \int_{\mathscr{Z}} \nu(dt \times dz)$ is a $\mathbb{P}_0$-martingale.

The likelihood process L_t is then re-defined as

$$L_t = \exp\left\{ \int_0^t h_s dY_s - \frac{1}{2} \int_0^t h_s^2 ds \right\} \cdot \exp\left\{ \int_{(0,t]} \int_{\mathscr{Z}} \log\left(\frac{\Lambda_{s-} f_Z(z; X_{s-})}{\mu f_Z^0(z)} \right) N(ds \times dz) - \int_0^t (\Lambda_s - \mu) ds \right\},$$

so that $L_{\sigma_k} = L_{\sigma_k -} \cdot \frac{\Lambda_{s-}}{\mu} \frac{f_Z(Z_k; X_{s-})}{f_Z^0(Z_k)}$. The rest of the analysis proceeds as before, with straightforward adjustments to the particle filtering algorithm in (17). Clearly, by providing further information, presence of marks facilitates disorder detection and will reduce the Bayes risk faced by the controller.

6.2 Jump Markov Signal

In our original setup, the signal process X undergoes a single transition, representing a permanent disorder in the signal. Practically, more complex signal dynamics may be modeled. Hence, instead of letting X consist of a single jump time θ and level χ, we may imagine a *sequence* of disorders, corresponding to a marked point process (θ_k, χ_k). In the simplest jump-Markov setting, we assume that different pairs (θ_k, χ_k) are independent and have the same common distribution (G, F). Since only the last disorder time affects the probability of the next jump, the pair $(t - \theta^t, X_t)$, with $\theta^t = \sup\{\theta_k : \theta_k \le t\}$, remains Markov. The first component $t - \theta^t$ undergoes deterministic evolution between disorder times, while the second component X_t is constant between disorders. SMC methods of Section 3.2 can be adapted to filter the evolution of such pair $(t - \theta^t, X_t)$, see e.g. [7, 32]. Note that the multiple disorders in X make sure that particles are never static, alleviating particle degeneracy concerns. With multiple disorders the measure change (3) remains the same, except that Λ_s and h_s are piecewise constant.

The above framework which has signal components undergoing deterministic evolution between disorder times can be extended to allow generic piecewise-deterministic, or jump-Markov process for X. A practical example would be the shot-noise process that experiences jumps at disorder times and decays exponentially towards a long-run mean otherwise, $X_t = \bar{x} + \sum_{k:\theta_k<t} Z_k e^{-r_k(t-\theta_k)}$ where r_k is the decay rate associated with the k-th jump of size $Z_k \in \mathbb{R}$. Bayesian decision making in such more complex models will be explored in a separate forthcoming work.

Acknowledgement

Assistance with the computational examples was provided by Chunhsiung Lu. I am grateful to two anonymous referees for numerous helpful comments and to Jarad Niemi for many discussions on SMC methods.

References

1. Bain, A., Crisan, D.: Fundamentals of stochastic filtering, *Stochastic Modelling and Applied Probability*, vol. 60. Springer, New York (2009)
2. Baron, M., Tartakovsky, A.: Asymptotic optimality of change-point detection schemes in general continuous-time models. Sequential Analysis **25**(3), 257–296 (2006)
3. Bayraktar, E., Dayanik, S., Karatzas, I.: Adaptive Poisson disorder problem. Ann. Appl. Probab. **16**(3), 1190–1261 (2006)
4. Bouchard, B., Warin, X.: Monte-Carlo valorisation of American options: facts and new algorithms to improve existing methods. In: R. Carmona, P.D. Moral, P. Hu, N. Oudjane (eds.) Numerical Methods in Finance, Springer Proceedings in Mathematics. Springer (2011)
5. Cappé, O., Moulines, E., Rydén, T.: Inference in hidden Markov models. Springer Series in Statistics. Springer, New York (2005)
6. Chopin, N.: Dynamic detection of change points in long time series. Annals of the Institute of Statistical Mathematics **59**(2), 349–366 (2007)
7. Chopin, N., Varini, E.: Particle filtering for continuous-time hidden Markov models. In: ESAIM: Proceedings, vol. 19, pp. 12–17 (2007)
8. Coquet, F., Toldo, S.: Convergence of values in optimal stopping and convergence of optimal stopping times. Electr. J. Probab **12**, 207–228 (2007)
9. Davis, M.H.A.: Markov Models and Optimization. Chapman & Hall, London (1993)
10. Dayanik, S., Poor, V., Sezer, S.: Multisource Bayesian sequential change detection. Annals of Applied Probability **2**, 552–590 (2008)
11. Dayanik, S., Poor, V., Sezer, S.: Sequential multi-hypothetis testing for compound Poisson processes. Stochastics **80**(1), 19–50 (2008)
12. Dayanik, S., Sezer, S.: Multisource Bayesian sequential hypothesis testing (2010). Preprint
13. Del Moral, P., Hu, P., Oudjane, N., Rémillard, B.: On the robustness of the Snell envelope. SIAM J. Financial Mathematics p. to Appear (2011)
14. Douc, R., Cappé, O.: Comparison of resampling schemes for particle filtering. In: Image and Signal Processing and Analysis, 2005. ISPA 2005. Proceedings of the 4th International Symposium on, pp. 64–69. IEEE (2005)
15. Doucet, A., de Freitas, N., Gordon, N. (eds.): Sequential Monte Carlo methods in practice. Statistics for Engineering and Information Science. Springer-Verlag, New York (2001)
16. Egloff, D.: Monte Carlo algorithms for optimal stopping and statistical learning. Annals of Applied Probability **15**(2), 1396–1432 (2005)
17. Elliott, R.J., Wu, P.: Hidden Markov chain filtering for a jump diffusion model. Stoch. Anal. Appl. **23**(1), 153–163 (2005)
18. Gobet, E., Lemor, J., Warin, X.: Rate of convergence of an empirical regression method for solving generalized backward stochastic differential equations. Bernoulli **12(5)**, 889–916 (2006)

19. Godsill, S.: Particle filters for continuous-time jump models in tracking applications. In: Conference Oxford sur les méthodes de Monte Carlo séquentielles, *ESAIM Proc.*, vol. 19, pp. 39–52. EDP Sci., Les Ulis (2007)
20. Gordon, N., Salmond, D., Smith, A.: Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In: IEE proceedings, vol. 140, pp. 107–113 (1993)
21. Lai, T., Xing, H.: Sequential change-point detection when the pre- and post-change parameters are unknown. Sequential Anal. **29**, 162–175 (2010)
22. Liu, J., West, M.: Combined parameter and state estimation in simulation-based filtering. In: Sequential Monte Carlo methods in practice, Stat. Eng. Inf. Sci., pp. 197–223. Springer, New York (2001)
23. Longstaff, F., Schwartz, E.: Valuing American options by simulations: a simple least squares approach. The Review of Financial Studies **14**, 113–148 (2001)
24. Ludkovski, M.: A simulation approach to optimal stopping under partial observations. Stoch. Proc. Appl. **119**(12), 4061–4087 (2009)
25. Ludkovski, M., Sezer, S.: Finite horizon decision timing with partially observable Poisson processes. Stochastic Models p. to Appear (2011). Available at arxiv.org/abs/1105.1484
26. Mei, Y.: Is average run length to false alarm always an information criterion? Sequential Anal. **27**, 351–376 (2008)
27. Mulinacci, S., Pratelli, M.: Functional convergence of Snell envelopes: applications to American options approximations. Finance and Stochastics **2**(3), 311–327 (1998)
28. Peskir, G., Shiryaev, A.N.: Optimal Stopping and Free-boundary problems. Birkhauser-Verlag, Lectures in Mathematics, ETH Zurich (2006)
29. Pham, H., Runggaldier, W., Sellami, A.: Approximation by quantization of the filter process and applications to optimal stopping problems under partial observation. Monte Carlo Methods Appl. **11**(1), 57–81 (2005)
30. Shiryaev, A.N.: Optimal methods in quickest detection problems. Teor. Verojatnost. i Primenen. **8**, 26–51 (1963)
31. Vellekoop, M.H., Clark, J.M.C.: A nonlinear filtering approach to changepoint detection problems: direct and differential-geometric methods. SIAM J. Control Optim. **42**(2), 469–494 (electronic) (2003)
32. Whiteley, N., Johansen, A., Godsill, S.: Monte carlo filtering of piecewise deterministic processes. Journal of Computational and Graphical Statistics pp. 1–21 (2010)
33. Wong, E., Hajek, B.: Stochastic processes in engineering systems. Springer Texts in Electrical Engineering. Springer-Verlag, New York (1985)

Appendix A

Algorithm 1 Adaptive Disorder Detection using Regression Monte Carlo

Input: K (number of paths); n (number of particles per path); Δt (time step for Snell envelope); $B_i(\rho)$ (regression basis functions); r (number of basis functions).

for $k \in \{1,2,\ldots,K\}$ **do**

Simulate standard Poisson process $(N_t^k, \mathbb{P}_0)$ with intensity μ, and standard Wiener process $(Y_t^k, \mathbb{P}_0)$ on $[0,T]$.

Sample n particles forming $\rho_0^k \equiv \rho_0^{(n),k}$ from the prior π_0 of X_0.

Use the particle filter algorithm to compute ρ_t^k along the path (N_t^k, Y_t^k) for $t = 0, \Delta t, \ldots, T$.

Initialize $u^k(T) = \rho_T^k H_2$, $\tau^k(T) = T$.

end for

for $t = (M-1)\Delta t, \ldots, \Delta t, 0$ **do**

Evaluate the basis functions $B_i(\rho_t^k)$, for $i = 1,\ldots,r$ and $k = 1,\ldots,K$.

Regress

$$\alpha^{(K)}(t) \triangleq \underset{(\alpha^1,\ldots,\alpha^r)\in\mathbb{R}^r}{\operatorname{argmin}} \sum_{k=1}^{K} \left| u^k(t+\Delta t) - \sum_{i=1}^{r} \alpha^i B_i(\rho_t^k) \right|^2.$$

for $k = 1,\ldots,K$ **do**

Set $h^{1,k}(t) := \rho_t^k H_1$ and $h^{2,k}(t) := \rho_t^k H_2$.

Set $\hat{q}^k(t) = h_t^{1,k}\Delta t + \sum_{i=1}^{r} \alpha^{(K),i}(t) B_i(\rho_t^k)$. // predicted continuation value

Set $u^k(t) = \begin{cases} h^{2,k}(t) & \text{if } \hat{q}^k(t) > h^{2,k}(t); \\ u^k(t+\Delta t) + h^{1,k}(t)\Delta t & \text{otherwise.} \end{cases}$

Update $\tau^k(t) = \begin{cases} t & \text{if } \hat{q}^k(t) > h^{2,k}(t); \\ \tau^k(t+\Delta t) & \text{otherwise.} \end{cases}$

end for

end for

return $u(\pi_0) \simeq \frac{1}{K}\sum_{k=1}^{K} u^k(0)$.