Parameter Uncertainty in the Kalman–Bucy Filter

Andrew Allan

Joint work with Samuel Cohen

Oxford University

September 2018

Suppose we have a hidden Markov process X, that affects another process Y which we can see.

Suppose we have a hidden Markov process X, that affects another process Y which we can see.

Stochastic filtering is the problem of estimating the current value of X, from our observations of Y.

Suppose we have a hidden Markov process X, that affects another process Y which we can see.

Stochastic filtering is the problem of estimating the current value of X, from our observations of Y.

That is, we are interested in the distribution of $X_t|\mathcal{Y}_t$, where $\mathcal{Y}_t = \sigma(Y_s : 0 \leq s \leq t)$ is all the information we have about Y up to time t.

Suppose we have a hidden Markov process X, that affects another process Y which we can see.

Stochastic filtering is the problem of estimating the current value of X, from our observations of Y.

That is, we are interested in the distribution of $X_t|\mathcal{Y}_t$, where $\mathcal{Y}_t = \sigma(Y_s : 0 \leq s \leq t)$ is all the information we have about Y up to time t.

As in any Bayesian framework, we recursively update this distribution as we make new observations.

In this talk we focus on a continuous time model where the underlying dynamics of X and Y are linear.

In this talk we focus on a continuous time model where the underlying dynamics of X and Y are linear.

That is, we assume the signal X and observation Y satisfy

$$\mathrm{d}X_t = \alpha_t X_t \,\mathrm{d}t + \sqrt{\beta_t}\,\mathrm{d}B_t,$$
$$\mathrm{d}Y_t = c_t X_t \,\mathrm{d}t + \mathrm{d}W_t,$$

with the initial conditions $Y_0 = 0$ and $X_0 \sim N(\mu_0, \sigma_0^2)$.

In this talk we focus on a continuous time model where the underlying dynamics of X and Y are linear.

That is, we assume the signal X and observation Y satisfy

$$\mathrm{d}X_t = \alpha_t X_t \,\mathrm{d}t + \sqrt{\beta_t}\,\mathrm{d}B_t,$$
$$\mathrm{d}Y_t = c_t X_t \,\mathrm{d}t + \mathrm{d}W_t,$$

with the initial conditions $Y_0 = 0$ and $X_0 \sim N(\mu_0, \sigma_0^2)$.

In this case the posterior distriution of X_t is Gaussian, i.e. $X_t|\mathcal{Y}_t \sim N(q_t, R_t)$.

In this talk we focus on a continuous time model where the underlying dynamics of X and Y are linear.

That is, we assume the signal X and observation Y satisfy

$$\mathrm{d}X_t = \alpha_t X_t \,\mathrm{d}t + \sqrt{\beta_t}\,\mathrm{d}B_t,$$
$$\mathrm{d}Y_t = c_t X_t \,\mathrm{d}t + \mathrm{d}W_t,$$

with the initial conditions $Y_0 = 0$ and $X_0 \sim N(\mu_0, \sigma_0^2)$.

In this case the posterior distriution of X_t is Gaussian, i.e. $X_t|\mathcal{Y}_t \sim N(q_t, R_t)$. Moreover, the conditional mean q_t and variance R_t satisfy the equations

$$\mathrm{d}q_t = \alpha_t q_t \,\mathrm{d}t + c_t R_t(\mathrm{d}Y_t - c_t q_t \,\mathrm{d}t),$$
$$\frac{\mathrm{d}R_t}{\mathrm{d}t} = \beta_t + 2\alpha_t R_t - c_t^2 R_t^2.$$

This is the standard Kalman–Bucy filter.

This is all well and good,

This is all well and good, except this assumes that we have exact knowledge of the system dynamics, i.e. of the parameters α, β, c, μ_0 and σ_0.

This is all well and good, except this assumes that we have exact knowledge of the system dynamics, i.e. of the parameters α, β, c, μ_0 and σ_0.

In practice these parameters must be estimated from data, which introduces statistical uncertainty.

This is all well and good, except this assumes that we have exact knowledge of the system dynamics, i.e. of the parameters α, β, c, μ_0 and σ_0.

In practice these parameters must be estimated from data, which introduces statistical uncertainty.

This is problematic, as stochastic filters are generally sensitive to perturbations in the parameters.

This is all well and good, except this assumes that we have exact knowledge of the system dynamics, i.e. of the parameters α, β, c, μ_0 and σ_0.

In practice these parameters must be estimated from data, which introduces statistical uncertainty.

This is problematic, as stochastic filters are generally sensitive to perturbations in the parameters.

Thus, our goal is to construct filters which are robust to parameter uncertainty.

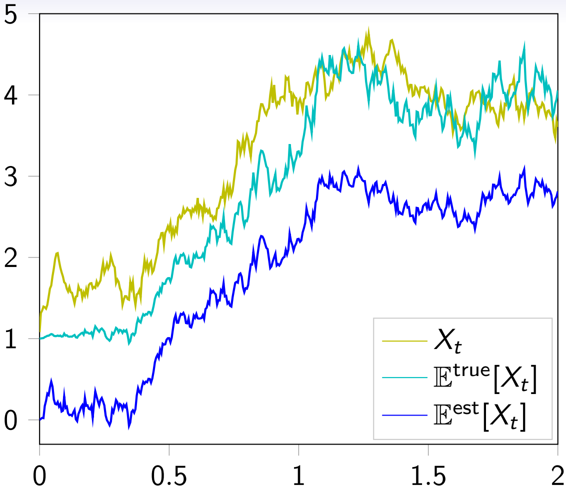

	α	β	μ_0	σ_0
true parameters	0.5	1.5	1	0.2
estimated parameters	0	1	0	1

We can calculate estimates $\mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t) \,|\, \mathcal{Y}_t]$ for any choice of parameters, but we don't know which choice is the correct one.

We can calculate estimates $\mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t)\,|\,\mathcal{Y}_t]$ for any choice of parameters, but we don't know which choice is the correct one.

We could consider the (essential) supremum over all choices of parameters:

$$\operatorname*{ess\,sup}_{\alpha,\beta,\mu_0,\sigma_0} \mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t)\,|\,\mathcal{Y}_t],$$

which corresponds to taking a **sublinear expectation** of $\varphi(X_t)$.

We can calculate estimates $\mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t) \mid \mathcal{Y}_t]$ for any choice of parameters, but we don't know which choice is the correct one.

We could consider the (essential) supremum over all choices of parameters:

$$\operatorname*{ess\,sup}_{\alpha,\beta,\mu_0,\sigma_0} \mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t) \mid \mathcal{Y}_t],$$

which corresponds to taking a **sublinear expectation** of $\varphi(X_t)$.

This is rather pessimistic, as we consider all possible models equally, including very implausible ones.

We can calculate estimates $\mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t)\,|\,\mathcal{Y}_t]$ for any choice of parameters, but we don't know which choice is the correct one.

We could consider the (essential) supremum over all choices of parameters:

$$\operatorname*{ess\,sup}_{\alpha,\beta,\mu_0,\sigma_0} \mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t)\,|\,\mathcal{Y}_t],$$

which corresponds to taking a **sublinear expectation** of $\varphi(X_t)$.

This is rather pessimistic, as we consider all possible models equally, including very implausible ones.

We propose that one can do better by using a **convex expectation**, which amounts to introducing a penalty term:

$$\mathcal{E}(\varphi(X_t)\,|\,\mathcal{Y}_t) = \operatorname*{ess\,sup}_{\alpha,\beta,\mu_0,\sigma_0} \Big\{\mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t)\,|\,\mathcal{Y}_t] - \text{`penalty'}_{\alpha,\beta,\mu_0,\sigma_0}\Big\}.$$

More precisely, we define

$$\mathcal{E}(\varphi(X_t)\,|\,\mathcal{Y}_t) = \operatorname*{ess\,sup}_{\alpha,\beta,\mu_0,\sigma_0}\Big\{\mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t)\,|\,\mathcal{Y}_t] - \Big(\frac{1}{k_1}\Big(\int_0^t \gamma(s,\alpha_s,\beta_s)\,\mathrm{d}s + \kappa_0(\mu_0,\sigma_0^2)\Big)\Big)^{k_2}\Big\},$$

More precisely, we define

$$\mathcal{E}(\varphi(X_t)\,|\,\mathcal{Y}_t) = \operatorname*{ess\,sup}_{\alpha,\beta,\mu_0,\sigma_0}\Big\{\mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t)\,|\,\mathcal{Y}_t] - \Big(\frac{1}{k_1}\Big(\int_0^t \gamma(s,\alpha_s,\beta_s)\,\mathrm{d}s + \kappa_0(\mu_0,\sigma_0^2)\Big)\Big)^{k_2}\Big\},$$

where γ and κ_0 are penalty functions, and the parameters k_1, k_2 can be chosen depending on how averse we are to uncertainty.

More precisely, we define

$$\mathcal{E}(\varphi(X_t)\,|\,\mathcal{Y}_t) = \operatorname*{ess\,sup}_{\alpha,\beta,\mu_0,\sigma_0}\Big\{\mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t)\,|\,\mathcal{Y}_t] - \Big(\frac{1}{k_1}\Big(\int_0^t \gamma(s,\alpha_s,\beta_s)\,\mathrm{d}s + \kappa_0(\mu_0,\sigma_0^2)\Big)\Big)^{k_2}\Big\},$$

where γ and κ_0 are penalty functions, and the parameters k_1, k_2 can be chosen depending on how averse we are to uncertainty.

The functions γ, κ_0 are nonnegative, and equal to zero for our a priori estimate of the parameters.

More precisely, we define

$$\mathcal{E}(\varphi(X_t)\,|\,\mathcal{Y}_t) = \operatorname*{ess\,sup}_{\alpha,\beta,\mu_0,\sigma_0}\Big\{\mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t)\,|\,\mathcal{Y}_t] - \Big(\frac{1}{k_1}\Big(\int_0^t \gamma(s,\alpha_s,\beta_s)\,\mathrm{d}s + \kappa_0(\mu_0,\sigma_0^2)\Big)\Big)^{k_2}\Big\},$$

where γ and κ_0 are penalty functions, and the parameters k_1, k_2 can be chosen depending on how averse we are to uncertainty.

The functions γ, κ_0 are nonnegative, and equal to zero for our a priori estimate of the parameters.

In particular, the standard Kalman–Bucy filter is recovered in the limit as $k_1 \to 0$.

In filtering we perform estimation based on observations of Y.
Therefore, we actually want to work with a fixed realisation y of Y.

In filtering we perform estimation based on observations of Y. Therefore, we actually want to work with a fixed realisation y of Y.

Rearranging things a bit, for any such realisation y, we have that

$$\mathcal{E}(\varphi(X_t)\,|\,y) = \sup_{(\mu,\sigma)\in\mathbb{R}\times(0,\infty)} \left\{ \mathbb{E}^{\mu,\sigma^2}[\varphi(X_t)\,|\,y] - \left(\frac{1}{k_1}\kappa_t(\mu,\sigma^2\,|\,y)\right)^{k_2} \right\},$$

where $X_t \sim N(\mu,\sigma^2)$ under $\mathbb{E}^{\mu,\sigma^2}$,

In filtering we perform estimation based on observations of Y. Therefore, we actually want to work with a fixed realisation y of Y.

Rearranging things a bit, for any such realisation y, we have that

$$\mathcal{E}(\varphi(X_t)\,|\,y) = \sup_{(\mu,\sigma)\in\mathbb{R}\times(0,\infty)} \left\{ \mathbb{E}^{\mu,\sigma^2}[\varphi(X_t)\,|\,y] - \left(\frac{1}{k_1}\kappa_t(\mu,\sigma^2\,|\,y)\right)^{k_2} \right\},$$

where $X_t \sim N(\mu,\sigma^2)$ under $\mathbb{E}^{\mu,\sigma^2}$, and

$$\kappa_t(\mu,\sigma^2\,|\,y) = \inf\left\{ \int_0^t \gamma(s,\alpha_s,\beta_s)\,\mathrm{d}s + \kappa_0(\mu_0,\sigma_0^2) \;\middle|\; \begin{array}{c} \alpha,\beta,\mu_0,\sigma_0 \quad \text{such that} \\ (q(y)_t, R_t) = (\mu,\sigma^2) \end{array} \right\}.$$

In filtering we perform estimation based on observations of Y. Therefore, we actually want to work with a fixed realisation y of Y.

Rearranging things a bit, for any such realisation y, we have that

$$\begin{aligned}&\mathcal{E}(\varphi(X_t)\,|\,y)\\&\quad=\sup_{(\mu,\sigma)\in\mathbb{R}\times(0,\infty)}\left\{\mathbb{E}^{\mu,\sigma^2}[\varphi(X_t)\,|\,y]-\left(\frac{1}{k_1}\kappa_t(\mu,\sigma^2\,|\,y)\right)^{k_2}\right\},\end{aligned}$$

where $X_t \sim N(\mu,\sigma^2)$ under $\mathbb{E}^{\mu,\sigma^2}$, and

$$\begin{aligned}&\kappa_t(\mu,\sigma^2\,|\,y)\\&\quad=\inf\left\{\int_0^t\gamma(s,\alpha_s,\beta_s)\,\mathrm{d}s+\kappa_0(\mu_0,\sigma_0^2)\;\middle|\;\begin{array}{c}\alpha,\beta,\mu_0,\sigma_0\quad\text{such that}\\(q(y)_t,R_t)=(\mu,\sigma^2)\end{array}\right\}.\end{aligned}$$

Our goal is therefore to calculate κ_t, which represents how reasonable we believe different posterior distributions of the signal to be.

$$\kappa_t(\mu, \sigma^2 \,|\, y) = \inf\left\{ \int_0^t \gamma(s, \alpha_s, \beta_s)\,\mathrm{d}s + \kappa_0(\mu_0, \sigma_0^2) \;\middle|\; \begin{array}{c} \alpha, \beta, \mu_0, \sigma_0 \quad \text{such that} \\ (q(y)_t, R_t) = (\mu, \sigma^2) \end{array} \right\},$$

$$\kappa_t(\mu, \sigma^2 \,|\, y) = \inf\left\{ \int_0^t \gamma(s, \alpha_s, \beta_s)\, \mathrm{d}s + \kappa_0(\mu_0, \sigma_0^2) \;\middle|\; \begin{array}{c} \alpha, \beta, \mu_0, \sigma_0 \quad \text{such that} \\ (q(y)_t, R_t) = (\mu, \sigma^2) \end{array} \right\},$$

where q and R satisfy

$$\mathrm{d}q_t = \alpha_t q_t \,\mathrm{d}t + c_t R_t(\mathrm{d}Y_t - c_t q_t \,\mathrm{d}t),$$

$$\frac{\mathrm{d}R_t}{\mathrm{d}t} = \beta_t + 2\alpha_t R_t - c_t^2 R_t^2.$$

$$\kappa_t(\mu, \sigma^2 \,|\, y) = \inf\left\{ \int_0^t \gamma(s, \alpha_s, \beta_s)\,\mathrm{d}s + \kappa_0(\mu_0, \sigma_0^2) \;\middle|\; \begin{array}{c} \alpha, \beta, \mu_0, \sigma_0 \quad \text{such that} \\ (q(y)_t, R_t) = (\mu, \sigma^2) \end{array} \right\},$$

where q and R satisfy

$$\mathrm{d}q_t = \alpha_t q_t\,\mathrm{d}t + c_t R_t(\mathrm{d}Y_t - c_t q_t\,\mathrm{d}t),$$
$$\frac{\mathrm{d}R_t}{\mathrm{d}t} = \beta_t + 2\alpha_t R_t - c_t^2 R_t^2.$$

We recognise this as an **optimal control problem**, where κ is the value function, q, R are the state trajectories, and α, β are the controls.

$$\kappa_t(\mu, \sigma^2 \,|\, y) = \inf\left\{ \int_0^t \gamma(s, \alpha_s, \beta_s)\,\mathrm{d}s + \kappa_0(\mu_0, \sigma_0^2) \;\middle|\; \begin{array}{c} \alpha, \beta, \mu_0, \sigma_0 \quad \text{such that} \\ (q(y)_t, R_t) = (\mu, \sigma^2) \end{array} \right\},$$

where q and R satisfy

$$\mathrm{d}q_t = \alpha_t q_t\,\mathrm{d}t + c_t R_t(\mathrm{d}Y_t - c_t q_t\,\mathrm{d}t),$$

$$\frac{\mathrm{d}R_t}{\mathrm{d}t} = \beta_t + 2\alpha_t R_t - c_t^2 R_t^2.$$

We recognise this as an **optimal control problem**, where κ is the value function, q, R are the state trajectories, and α, β are the controls.

Note that this control problem is posed ‘backwards in time’.

$$\kappa_t(\mu, \sigma^2 \,|\, y) = \inf \left\{ \int_0^t \gamma(s, \alpha_s, \beta_s)\, \mathrm{d}s + \kappa_0(\mu_0, \sigma_0^2) \;\middle|\; \begin{array}{l} \alpha, \beta, \mu_0, \sigma_0 \quad \text{such that} \\ (q(y)_t, R_t) = (\mu, \sigma^2) \end{array} \right\},$$

where q and R satisfy

$$\mathrm{d}q_t = \alpha_t q_t \,\mathrm{d}t + c_t R_t(\mathrm{d}Y_t - c_t q_t \,\mathrm{d}t),$$

$$\frac{\mathrm{d}R_t}{\mathrm{d}t} = \beta_t + 2\alpha_t R_t - c_t^2 R_t^2.$$

We recognise this as an **optimal control problem**, where κ is the value function, q, R are the state trajectories, and α, β are the controls.

Note that this control problem is posed ‘backwards in time’.

Moreover, since y is a realisation of the stochastic process Y, this is actually a **pathwise stochastic** control problem.

$$\kappa_t(\mu,\sigma^2\,|\,y) = \inf\left\{ \int_0^t \gamma(s,\alpha_s,\beta_s)\,\mathrm{d}s + \kappa_0(\mu_0,\sigma_0^2) \;\middle|\; \begin{array}{c} \alpha,\beta,\mu_0,\sigma_0 \quad \text{such that} \\ (q(y)_t, R_t) = (\mu,\sigma^2) \end{array} \right\},$$

where q and R satisfy

$$\mathrm{d}q_t = \alpha_t q_t\,\mathrm{d}t + c_t R_t(\mathrm{d}Y_t - c_t q_t\,\mathrm{d}t),$$
$$\frac{\mathrm{d}R_t}{\mathrm{d}t} = \beta_t + 2\alpha_t R_t - c_t^2 R_t^2.$$

We recognise this as an **optimal control problem**, where κ is the value function, q, R are the state trajectories, and α,β are the controls.

Note that this control problem is posed ‘backwards in time’.

Moreover, since y is a realisation of the stochastic process Y, this is actually a **pathwise stochastic** control problem. In particular, κ has ‘Brownian-like’ regularity in time.

We transformation the state trajectories via

$$w(t) = \left(\frac{q_t}{R_t} - \eta_t,\ \frac{1}{R_t} \right),$$

We transformation the state trajectories via

$$w(t) = \left(\frac{q_t}{R_t} - \eta_t,\ \frac{1}{R_t}\right),$$

where, for a fixed $\omega \in \Omega$, we fix the value of the stochastic integral as

$$\eta_t = \left(\int_0^t c_s \, \mathrm{d}Y_s\right)(\omega).$$

We transformation the state trajectories via

$$w(t) = \left(\frac{q_t}{R_t} - \eta_t,\ \frac{1}{R_t}\right),$$

where, for a fixed $\omega \in \Omega$, we fix the value of the stochastic integral as

$$\eta_t = \left(\int_0^t c_s \,\mathrm{d}Y_s\right)(\omega).$$

Then w satisfies

$$\frac{\mathrm{d}w}{\mathrm{d}s}(s) = f(w(s), s, \alpha_s, \beta_s),$$

We transformation the state trajectories via

$$w(t) = \left(\frac{q_t}{R_t} - \eta_t,\ \frac{1}{R_t} \right),$$

where, for a fixed $\omega \in \Omega$, we fix the value of the stochastic integral as

$$\eta_t = \left(\int_0^t c_s \, \mathrm{d} Y_s \right)(\omega).$$

Then w satisfies

$$\frac{\mathrm{d} w}{\mathrm{d} s}(s) = f(w(s), s, \alpha_s, \beta_s),$$

where the function f is given by

$$f(x, t, a, b) = \begin{pmatrix} -(x_1 + \eta_t)(a + b x_2) \\ -b x_2^2 - 2 a x_2 + c_t^2 \end{pmatrix}.$$

We also transform the value function via

$$v(x, t) = \kappa_t(\mu, \sigma^2 \,|\, y),$$

where

$$x = \left(\frac{\mu}{\sigma^2} - \eta_t,\ \frac{1}{\sigma^2}\right).$$

We thus obtain a deterministic control problem:

$$v(x,t) = \inf_{\alpha,\beta}\left\{ \int_0^t \gamma(s,\alpha_s,\beta_s)\,\mathrm{d}s + v_0(w(0))\right\},$$

where

$$\frac{\mathrm{d}w}{\mathrm{d}s}(s) = f(w(s),s,\alpha_s,\beta_s).$$

We thus obtain a deterministic control problem:

$$v(x,t) = \inf_{\alpha,\beta} \left\{ \int_0^t \gamma(s,\alpha_s,\beta_s)\,\mathrm{d}s + v_0(w(0)) \right\},$$

where

$$\frac{\mathrm{d}w}{\mathrm{d}s}(s) = f(w(s),s,\alpha_s,\beta_s).$$

The problem is still somewhat degenerate, as the ODE satisfied by w is not Lipschitz. We must therefore work to avoid 'blow ups' in the state trajectories.

We thus obtain a deterministic control problem:

$$v(x,t) = \inf_{\alpha,\beta} \left\{ \int_0^t \gamma(s,\alpha_s,\beta_s)\,\mathrm{d}s + v_0(w(0)) \right\},$$

where

$$\frac{\mathrm{d}w}{\mathrm{d}s}(s) = f(w(s), s, \alpha_s, \beta_s).$$

The problem is still somewhat degenerate, as the ODE satisfied by w is not Lipschitz. We must therefore work to avoid 'blow ups' in the state trajectories.

Moreover, the value function itself blows up at the boundary of its domain, as extreme values of the posterior mean and variance are considered to be very implausible.

Nonetheless, we establish a dynamic programming principle, and can then prove:

Nonetheless, we establish a dynamic programming principle, and can then prove:

Theorem

The value function v is the unique viscosity solution of the Hamilton–Jacobi–Bellman (HJB) equation

$$\frac{\partial v}{\partial t}(x,t) + \sup_{a,b}\left\{ f(x,t,a,b)\cdot\nabla v(x,t) - \gamma(t,a,b)\right\} = 0$$

which blows up at the boundary.

Consider again the parameters

true parameters	$\alpha = 0.5$	$\beta = 1.5$	$\mu_0 = 1$	$\sigma_0 = 0.2$
estimated parameters	$\alpha^* = 0$	$\beta^* = 1$	$\mu_0^* = 0$	$\sigma_0^* = 1$

Consider again the parameters

true parameters	$\alpha = 0.5$	$\beta = 1.5$	$\mu_0 = 1$	$\sigma_0 = 0.2$
estimated parameters	$\alpha^* = 0$	$\beta^* = 1$	$\mu_0^* = 0$	$\sigma_0^* = 1$

We simulate the signal X using the true parameters, and simulate the observation Y using $c = 1$.

Consider again the parameters

true parameters	$\alpha = 0.5$	$\beta = 1.5$	$\mu_0 = 1$	$\sigma_0 = 0.2$
estimated parameters	$\alpha^* = 0$	$\beta^* = 1$	$\mu_0^* = 0$	$\sigma_0^* = 1$

We simulate the signal X using the true parameters, and simulate the observation Y using $c = 1$.

We use the penalty functions given by

$$\gamma(t, a, b) = 5(a - \alpha^*)^2 + 10(b - \beta^*)^2,$$
$$v_0(x_1, x_2) = 15(x_1 - x_1^*)^2 + 15(x_2 - x_2^*)^2,$$

where $x_1^* = \frac{\mu_0^*}{(\sigma_0^*)^2}$ and $x_2^* = \frac{1}{(\sigma_0^*)^2}$.

Consider again the parameters

true parameters	$\alpha = 0.5$	$\beta = 1.5$	$\mu_0 = 1$	$\sigma_0 = 0.2$
estimated parameters	$\alpha^* = 0$	$\beta^* = 1$	$\mu_0^* = 0$	$\sigma_0^* = 1$

We simulate the signal X using the true parameters, and simulate the observation Y using $c = 1$.

We use the penalty functions given by

$$\gamma(t, a, b) = 5(a - \alpha^*)^2 + 10(b - \beta^*)^2,$$
$$v_0(x_1, x_2) = 15(x_1 - x_1^*)^2 + 15(x_2 - x_2^*)^2,$$

where $x_1^* = \frac{\mu_0^*}{(\sigma_0^*)^2}$ and $x_2^* = \frac{1}{(\sigma_0^*)^2}$.

We then numerically solve the HJB equation to find the value function v.

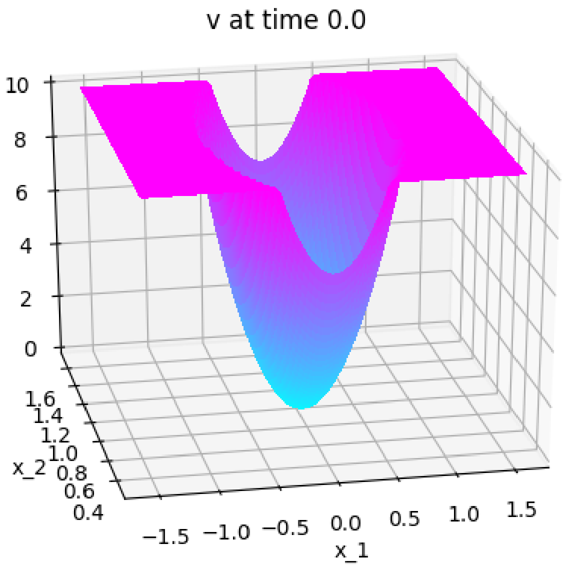

v at time 0.0
10
8
6
4
2
0
1.6
1.4
1.2
1.0
x_2
0.8
0.6
0.4
−1.5
−1.0
−0.5
0.0
0.5
1.0
1.5
x_1

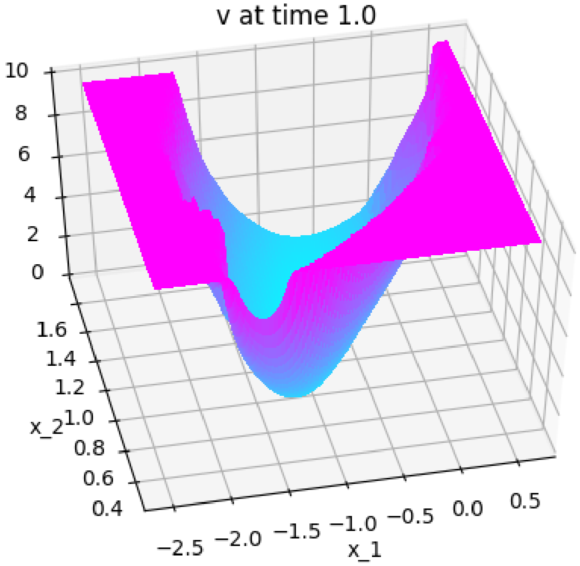
v at time 1.0
10
8
6
4
2
0
1.6
1.4
1.2
1.0
0.8
0.6
0.4
x_2
−2.5
−2.0
−1.5
−1.0
−0.5
0.0
0.5
x_1

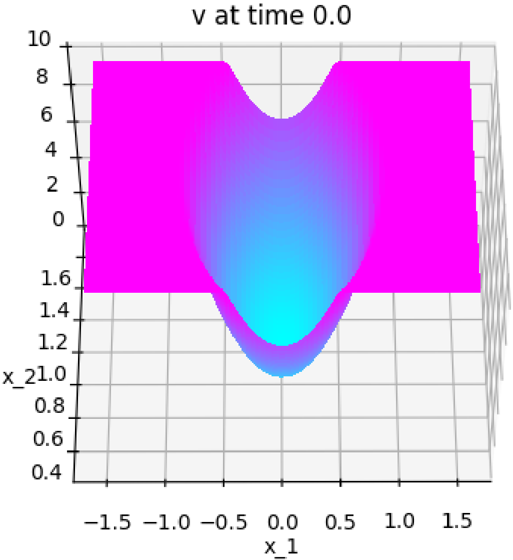
v at time 0.0
10
8
6
4
2
0
1.6
1.4
1.2
x_2
1.0
0.8
0.6
0.4
−1.5
−1.0
−0.5
0.0
0.5
1.0
1.5
x_1

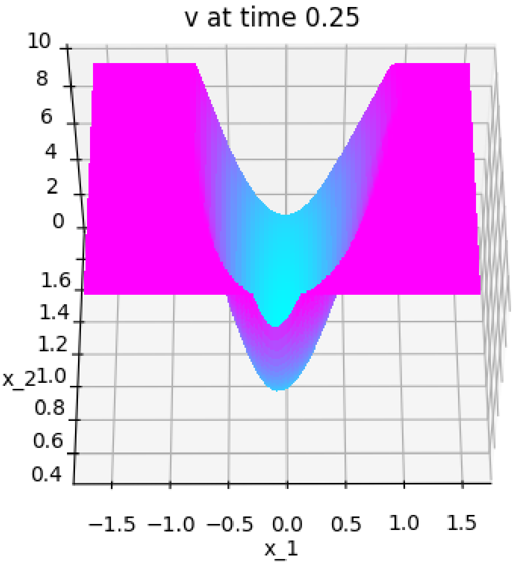
v at time 0.25
10
8
6
4
2
0
1.6
1.4
1.2
x_2
1.0
0.8
0.6
0.4
−1.5
−1.0
−0.5
0.0
0.5
1.0
1.5
x_1

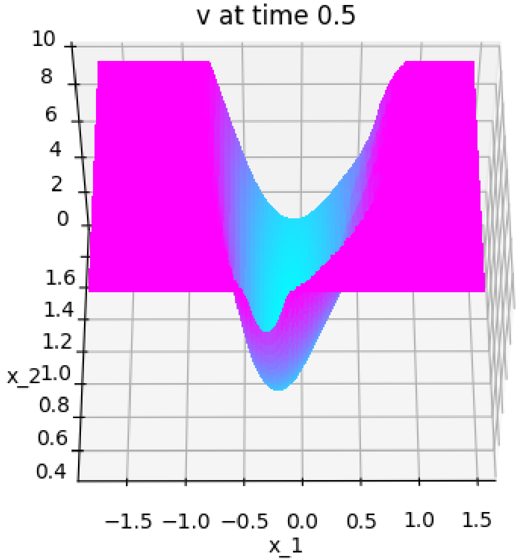
v at time 0.5
10
8
6
4
2
0
1.6
1.4
1.2
1.0
0.8
0.6
0.4
x_2
−1.5
−1.0
−0.5
0.0
0.5
1.0
1.5
x_1

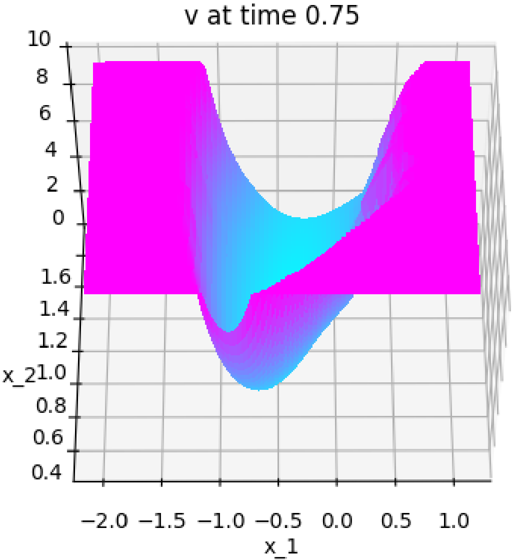

v at time 0.75
10
8
6
4
2
0
1.6
1.4
1.2
x_2
1.0
0.8
0.6
0.4
−2.0
−1.5
−1.0
−0.5
0.0
0.5
1.0
x_1

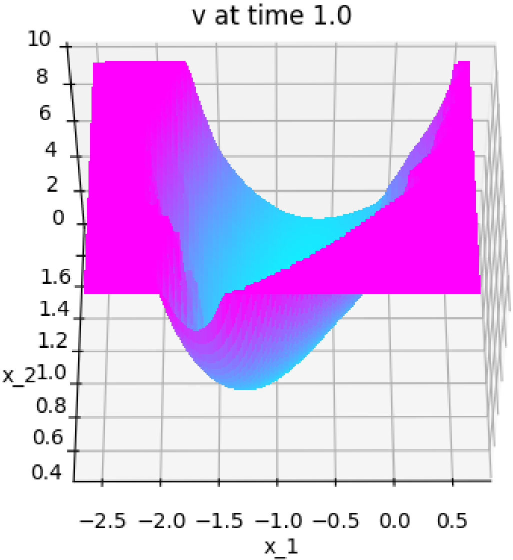
v at time 1.0
10
8
6
4
2
0
1.6
1.4
1.2
x_2
1.0
0.8
0.6
0.4
−2.5
−2.0
−1.5
−1.0
−0.5
0.0
0.5
x_1

Having solved for v, we can reverse the change of variables we did earlier to obtain κ.

Having solved for v, we can reverse the change of variables we did earlier to obtain κ.

Recalling the relation

$$\mathcal{E}(\varphi(X_t) \,|\, y) = \sup_{(\mu,\sigma)\in\mathbb{R}\times(0,\infty)} \left\{ \mathbb{E}^{\mu,\sigma^2}[\varphi(X_t) \,|\, y] - \left(\frac{1}{k_1} \kappa_t(\mu, \sigma^2 \,|\, y) \right)^{k_2} \right\},$$

we are now in a position to calculate robust estimates of the signal.

Having solved for v, we can reverse the change of variables we did earlier to obtain κ.

Recalling the relation

$$\mathcal{E}(\varphi(X_t) \mid y) = \sup_{(\mu,\sigma)\in\mathbb{R}\times(0,\infty)} \left\{ \mathbb{E}^{\mu,\sigma^2}[\varphi(X_t) \mid y] - \left(\frac{1}{k_1} \kappa_t(\mu, \sigma^2 \mid y) \right)^{k_2} \right\},$$

we are now in a position to calculate robust estimates of the signal.

In the following we take $k_1 = 10$ and $k_2 = 5$.

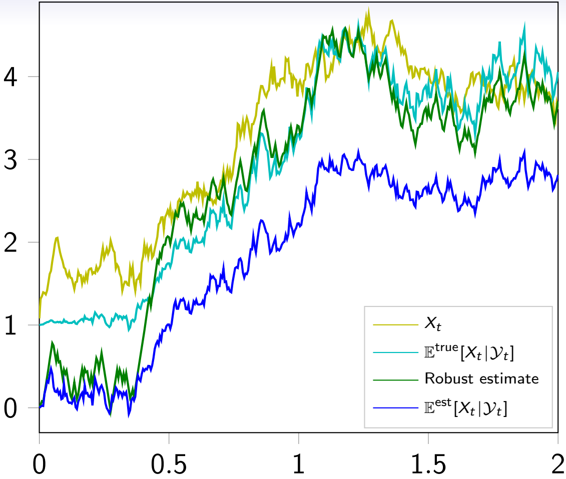

Here the robust estimate is calculated as

$$\arg\min_{\xi} \mathcal{E}\big((X_t - \xi)^2 \,\big|\, \mathcal{Y}_t\big).$$

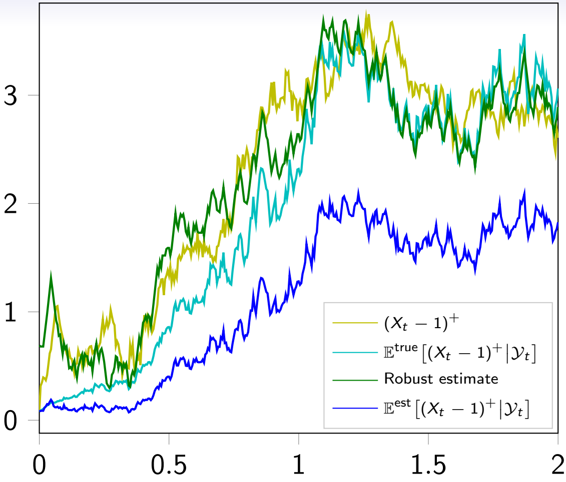

Here the robust estimate is calculated as

$$\arg\min_{\xi} \mathcal{E}\Big(\big((X_t - 1)^+ - \xi\big)^2 \,\Big|\, \mathcal{Y}_t\Big).$$

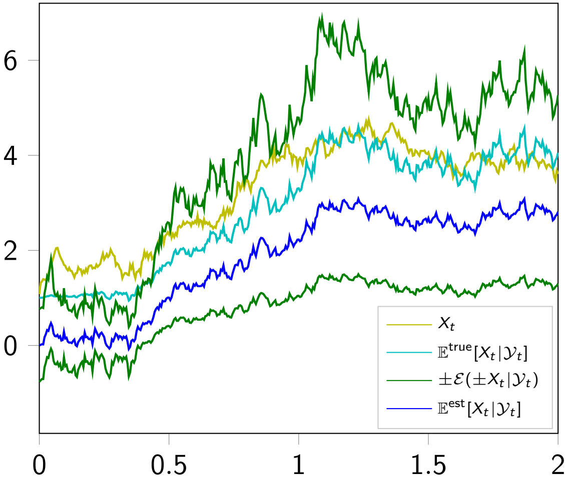

We also calculate the ‘upper’ and ‘lower’ expectations of the signal, i.e. $\mathcal{E}(X_t \,|\, \mathcal{Y}_t)$ and $-\mathcal{E}(-X_t \,|\, \mathcal{Y}_t)$ respectively.

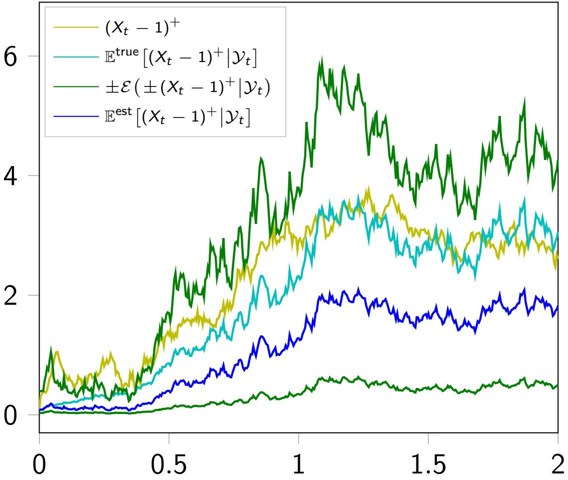

Finally, we calculate upper and lower expectations of $(X_t - 1)^+$.

Some extensions we are working on are:

Some extensions we are working on are:

- considering uncertainty in the parameter c,

Some extensions we are working on are:

- considering uncertainty in the parameter c,
- including cases with correlation between the signal noise and observation noise,

Some extensions we are working on are:

- considering uncertainty in the parameter c,
- including cases with correlation between the signal noise and observation noise,
- allowing the nonlinear expectation term to ‘learn’ from new observations, and update the penalty terms accordingly.

Recall the filter dynamics

$$\mathrm{d}q_t = \alpha_t q_t \,\mathrm{d}t + c_t R_t (\mathrm{d}Y_t - c_t q_t \,\mathrm{d}t),$$
$$\frac{\mathrm{d}R_t}{\mathrm{d}t} = \beta_t + 2\alpha_t R_t - c_t^2 R_t^2.$$

Recall the filter dynamics

$$\mathrm{d}q_t = \alpha_t q_t \,\mathrm{d}t + c_t R_t(\mathrm{d}Y_t - c_t q_t \,\mathrm{d}t),$$
$$\frac{\mathrm{d}R_t}{\mathrm{d}t} = \beta_t + 2\alpha_t R_t - c_t^2 R_t^2.$$

When c is uncertain, our pathwise stochastic control problem becomes more interesting, as the control appears in the coefficient of ‘Brownian-like’ process Y.

Recall the filter dynamics

$$\mathrm{d}q_t = \alpha_t q_t \,\mathrm{d}t + c_t R_t (\mathrm{d}Y_t - c_t q_t \,\mathrm{d}t),$$
$$\frac{\mathrm{d}R_t}{\mathrm{d}t} = \beta_t + 2\alpha_t R_t - c_t^2 R_t^2.$$

When c is uncertain, our pathwise stochastic control problem becomes more interesting, as the control appears in the coefficient of 'Brownian-like' process Y.

This problem actually turns out to be degenerate if we aren't careful to penalise controls with very fast fluctuations.

Recall the filter dynamics

$$\mathrm{d}q_t = \alpha_t q_t \,\mathrm{d}t + c_t R_t(\mathrm{d}Y_t - c_t q_t \,\mathrm{d}t),$$
$$\frac{\mathrm{d}R_t}{\mathrm{d}t} = \beta_t + 2\alpha_t R_t - c_t^2 R_t^2.$$

When c is uncertain, our pathwise stochastic control problem becomes more interesting, as the control appears in the coefficient of 'Brownian-like' process Y.

This problem actually turns out to be degenerate if we aren't careful to penalise controls with very fast fluctuations.

That is, parameters which 'look like Y' over small time scales should be considered to be very implausible, and this has to be accounted for in the penalty terms.

Recall the filter dynamics

$$\mathrm{d}q_t = \alpha_t q_t \,\mathrm{d}t + c_t R_t (\mathrm{d}Y_t - c_t q_t \,\mathrm{d}t),$$
$$\frac{\mathrm{d}R_t}{\mathrm{d}t} = \beta_t + 2\alpha_t R_t - c_t^2 R_t^2.$$

When c is uncertain, our pathwise stochastic control problem becomes more interesting, as the control appears in the coefficient of 'Brownian-like' process Y.

This problem actually turns out to be degenerate if we aren't careful to penalise controls with very fast fluctuations.

That is, parameters which 'look like Y' over small time scales should be considered to be very implausible, and this has to be accounted for in the penalty terms.

Allowing either correlated noise or 'learning' leads to the same phenomenon.

The end