Adaptive Robust Stochastic Control and Statistical Surrogates

Tao Chen
Department of Statistics and Applied Probability
University of California, Santa Barbara
tchen@pstat.ucsb.edu

Joint work with
M. Ludkovski

Robust Techniques in Quantitative Finance
Oxford, UK
September 3 2018

Motivations

Adaptive Robust Parametric Markovian Control

- To control the risk due to model uncertainty
  (error in model estimation or model misspecification)
- A robust framework adaptively reduces uncertainty through learning

Numerical Implementation

- Solve discrete time robust Bellman equation
- Three challenges: continuous state space, integration, optimization
- Main hurdle: curse of dimensionality

Main Goals

- To propose and study an **adaptive robust control approach** for solving a discrete time Markovian control problem subject to Knightian uncertainty

- To develop a new numerical methodology in response to the challenges arise when scaling the approach to higher dimensions.

T. R. Bielecki, T. Chen, I. Cialenco, A. Cousin, and M. Jeanblanc *Adaptive Robust Control Under Model Uncertainty.* Submitted for Publication, 2017.

T. Chen and M. Ludkovski *Robust Stochastic Control and Statistical Surrogates.* In preparation, 2018.

Example: Dynamic Optimal Portfolio Selection

An investor is deciding on investing in a risky asset and a risk-free banking account by maximizing the expected utility $U(W_T)$ of the terminal wealth.

- r^f - the constant risk free rate
- e^{Z_t} - the return on the risky asset
- Assume that $Z_t = \mu^* + \sigma^* \varepsilon_t$, where ε_t are i.i.d. $\mathcal{N}(0,1)$
- The dynamics of the wealth process produced by a s.f. strategy

$$W_{t+1} = W_t(1 + r^f + \varphi_t(e^{Z_{t+1}} - 1 - r^f)), \quad t \in \mathcal{T}', W_0 = w_0.$$

- Stochastic Control Problem if μ^* and σ^* are known:

$$\sup_{\varphi \in \mathcal{A}} \mathbb{E}_{\mu*,\sigma^*}[U(W_T^{\varphi})].$$

Notations

- $(\Omega, \mathcal{F})$ - measurable space
- $T \in \mathbb{N}$ - fixed time horizon
- $\mathcal{T} = \{0, 1, \dots, T\}$ and $\mathcal{T}' = \{0, 1, \dots, T-1\}$
- $\mathbf{\Theta} \subset \mathbb{R}^d$ - known parameter space
- $X = \{X_t, t \in \mathcal{T}\}$ - observed process taking values in $\mathbb{R}^k$
- $\mathbb{F} = \{\mathscr{F}_t, t \in \mathcal{T}\}$ - the natural filtration of X
- $\{\mathbb{P}_\theta, \theta \in \mathbf{\Theta} \subset \mathbb{R}^d\}$ - set of plausible laws of X
- $\mathbb{P}_{\theta^*}$ - (unknown) true law of X.

General Stochastic Control Problem

Consider the following general form of stochastic control problem

$$\inf_{\varphi \in \mathcal{A}} \mathbb{E}_{\theta^*}[L(X, \varphi)],$$

where $\mathcal{A}$ is the set of admissible control processes (some $\mathbb{F}$–adapted processes $\varphi = \{\varphi_t,\ t \in \mathcal{T}'\}$ taking values in A); L is a measurable functional (loss function, utility function, etc).

Since the true parameter $\theta^* \in \boldsymbol{\Theta}$ is unknown, the question is how to handle the stochastic control problem subject to this type of *model uncertainty*.

Classical Robust Approach

$$\inf_{\varphi \in \mathcal{A}} \sup_{\theta \in \boldsymbol{\Theta}} \mathbb{E}_\theta[L(X, \varphi)]$$

(in some cases) solved by Bellman equation of the form

$$V(t, x) = \inf_{a \in A} \sup_{\theta \in \boldsymbol{\Theta}} \mathbb{E}_\theta[V(t+1, X_{t+1}^{a,\theta}(x))]. \tag{1.1}$$

- Select the best strategy a^* over the worst possible model $\underline{\theta}$.
- "static robustness" and no reduction of uncertainty: fixed adversarial choice of $\underline{\theta}$ and $\boldsymbol{\Theta}$.
- If $\underline{\theta}$ is far from the true model θ^*, this approach is overly conservative.

- The classical robust control problem does not involve any reduction of uncertainty about θ^*; the parameter space is not "updated" with incoming information about the signal process X.
- Incorporating "learning" into the robust control paradigm appears like a good idea.
- Anderson, Hansen, Sargent (2003) state:
  "We see three important extensions to our current investigation. Like builders of rational expectations models, we have side-stepped the issue of how decision-makers select an approximating model. ... Just as we have not formally modelled how agents learned the approximating model, neither have we formally justified why they do not bother to learn about potentially complicated misspecifications of that model. **Incorporating forms of learning would be an important extension of our work.**"

Adaptive Robust Approach

To achieve reduction of uncertainty, we want to have the following dynamic programming equation different from (1.1):

$$V(t,x) = \inf_{a \in A} \sup_{\theta \in \mathbf{\Theta}_t} \mathbb{E}_\theta [V(t+1, X_{t+1}^{a,\theta}(x))], \tag{1.2}$$

where $\mathbf{\Theta}_t \in \mathscr{F}_t$ with a shrinking size, and $\mathbf{\Theta}_t \to \{\theta^*\}$.

We incorporate the idea of parameter learning to formulate set $\mathbf{\Theta}_t$

- point estimator $\hat{\theta}_t$ can be updated as observations come in
- choose $\mathbf{\Theta}_t$ as the confidence region centered at $\hat{\theta}_t$
- adaptive robust control

Without uncertainty

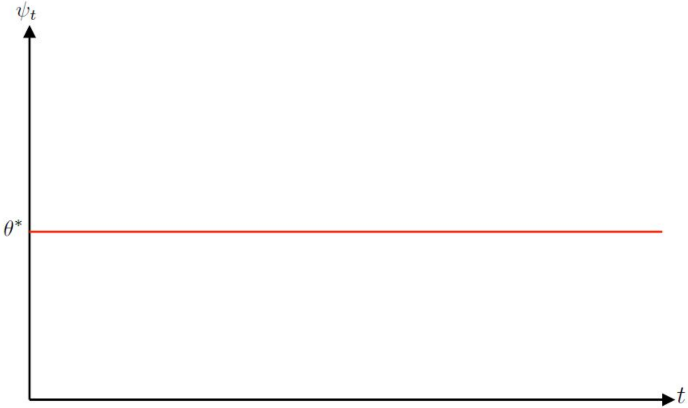

Robust

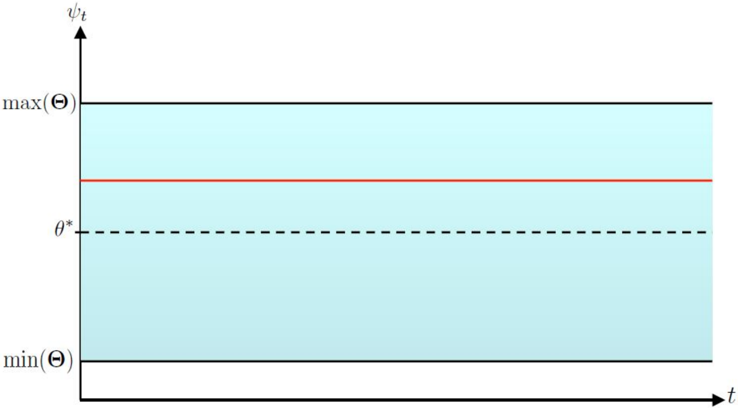

Strong robust

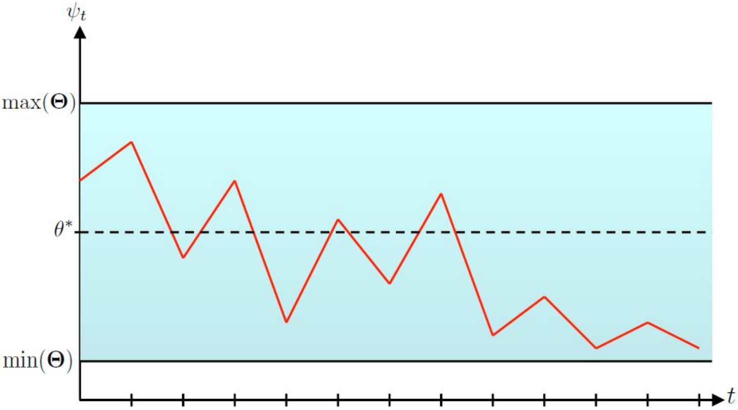

Adaptive Robust

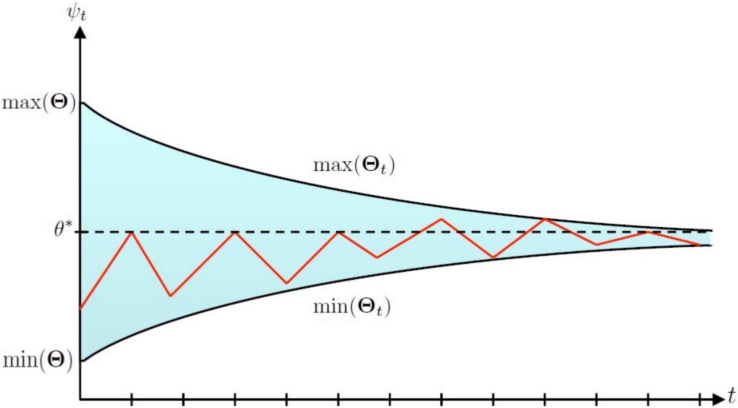

Problem Set-up

We assume that the observed state process X follows the dynamics

$$X_0 = x_0,$$
$$X_{t+1} = f(X_t, \varphi_t, Z_{t+1}), \quad t \in \mathcal{T},$$

where $Z = \{Z_t\}_{t \in \mathcal{T}'}$ is an $\mathbb{R}^n$-valued random sequence that is

- $\mathbb{F}$–adapted,
- observed,
- i.i.d. or ergodic Markovian under $\mathbb{P}_\theta$, $\theta \in \mathbf{\Theta} \subset \mathbb{R}^d$,

Example: Z_t is the excess return on risk assets in the optimal portfolio selection problem.

Dynamics of set Θ_t

Recall the adaptive Bellman equation that we want to have

$$V(t,x) = \inf_{a\in A} \sup_{\theta\in\mathbf{\Theta}_t} \mathbb{E}_\theta[V(t+1, f(x,a,Z^\theta_{t+1})].$$

Since $\hat{\theta}_t$ and $\mathbf{\Theta}_t$ now affect the inner optimization, $\hat{\theta}_t$ is augmented to the system state $Y_t = (X_t, \hat{\theta}_t)$.

$$V(t,x,\hat{\theta}) = \inf_{a\in A} \sup_{\theta\in\mathbf{\Theta}_t(\hat{\theta})} \mathbb{E}_\theta[V(t+1, \mathbf{T}(x,\hat{\theta},a,Z^\theta_{t+1})].$$

- To formulate the control problem, dynamics of $\hat{\theta}_t$ and $\mathbf{\Theta}_t$ are needed
- Recursive construction of confidence regions
- This brings in additional difficulty for numerics because the dimension of the state space is increased by d

Recursive Construction of Confidence Regions

In Bielecki, Chen and Cialenco (2017), for ergodic Markov chains Z, we showed that:

- A point estimator $\hat{\theta}_t$ of θ^* can be computed recursively

$$\hat{\theta}_0 = \theta_0,$$
$$\hat{\theta}_{t+1} = R(t, \hat{\theta}_t, Z_{t+1}),$$

where $R(t, c, z)$ is a deterministic measurable function.

- An approximate $1 - \alpha$-confidence region $\boldsymbol{\Theta}_t$ of θ^* can be constructed by a deterministic rule:

$$\boldsymbol{\Theta}_t = \tau_\alpha(t, \hat{\theta}_t)$$

where $\tau_\alpha(t, \cdot) : \mathbb{R}^d \to 2^{\boldsymbol{\Theta}}$ is a is a *deterministic* set valued function, $\mathbb{P}_{\theta^*}(\theta^* \in \boldsymbol{\Theta}_t) \approx 1 - \alpha$, and $\lim_{t\to\infty} \boldsymbol{\Theta}_t = \{\theta^*\}$.

How to formulate the stochastic control problem that is solved by the adaptive Bellman equation?

$$V(t,x,\hat{\theta}) = \inf_{a \in A} \sup_{\theta \in \mathbf{\Theta}_t(\hat{\theta})} \mathbb{E}_\theta[V(t+1, \mathbf{T}(x,\hat{\theta},a,Z^{\theta}_{t+1})].$$

- The expectation $\mathbb{E}_\theta[V(t+1, \mathbf{T}(x,\hat{\theta},a,Z^{\theta}_{t+1})]$ is computed according to transition probability of X with $\theta = \hat{\theta}^*_t$
- $\hat{\theta}^*_t$ depends on $\hat{\theta}_t$ which is updated according to the observations
- The transition probability function is now path dependent
- That leads to consideration of canonical construction of the augmented state process

Adversary Selector

Denote by $E_Y := \mathbb{R}^n \times \mathbb{R}^d$ the state space of augmented state process $Y_t = (X_t, \hat{\theta}_t)$ with dynamics

$$Y_{t+1} = \mathbf{T}(t, Y_t, \varphi_t, Z_{t+1}),$$

where $\mathbf{T}(t, y, a, z) := (f(x, a, z), R(t, c, z))$ and $y = (x, c)$.

We define the set of (adversary) selectors

$$\mathbf{\Psi} = \{(\psi_t)_{t\in\mathcal{T}'} \mid \psi_t : E_Y^{t+1} \to \mathbf{\Theta}_t, t \in \mathcal{T}'\}.$$

Remark. From a game point of view, process ψ is strategy played by controller's opponent. In classical robust approach, $\mathbf{\Theta}_t = \mathbf{\Theta}$, $\psi_t = \underline{\theta}$ fixed through time.

Canonical Law of the Augmented State Process

Define the Markov transition probability kernel on $\mathcal{E}_Y$ (Borel σ-algebra of E_Y)

$$Q(B \mid t, y, a, \theta) := \mathbb{P}_\theta(Y_{t+1} \in B \mid Y_t = y, \varphi_t = a),$$

for each $(t, y, a, \theta) \in \mathcal{T}' \times E_Y \times A \times \mathbf{\Theta}$,

Define the canonical law of the state process Y on E_Y^{T+1} as

$$\begin{aligned}
&\mathbb{Q}_{h_0}^{\varphi,\psi}(B_0, B_1, \ldots, B_T) \\
&\quad = \int_{B_0} \cdots \int_{B_T} Q(dx_T | T-1, x_{T-1}, \varphi_{T-1}(h_{T-1}), \psi_{T-1}(h_{T-1})) \\
&\quad \cdots Q(dx_1 \mid 0, x_0, \varphi_0(h_0), \psi_0(h_0))\delta_{h_0}(dx_0)
\end{aligned}$$

The adaptive robust control problem

$$\inf_{\varphi\in\mathcal{A}} \sup_{\mathbb{Q}\in\mathcal{Q}_{h_0}^{\varphi,\mathbf{\Psi}}} \mathbb{E}_{\mathbb{Q}}[L(X_T^{\varphi})],$$

where

$$\mathcal{Q}_{h_0}^{\varphi,\mathbf{\Psi}} := \{\mathbb{Q}_{h_0}^{\varphi,\psi}, \psi \in \mathbf{\Psi}\},$$

and

$$\mathbf{\Psi} = \{(\psi_t)_{t\in\mathcal{T}'} \mid \psi_t : \boldsymbol{H}_t \to \mathbf{\Theta}_t, t \in \mathcal{T}'\}.$$

Adaptive Robust DPP

Theorem

The solution φ^* *of*

$$\inf_{\varphi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathcal{Q}^{\varphi,\Psi}_{h_0}} \mathbb{E}_{\mathbb{Q}}[L(X^{\varphi}_T)]$$

can be obtained from the solution of the following Bellman equations:

$$V(T, x, \hat{\theta}) = L(x),$$

$$V(t, x, \hat{\theta}) = \inf_{a \in A} \sup_{\theta \in \tau_\alpha(t, \hat{\theta})} \int_{E_Y} V(t+1, y) Q(dy \mid t, x, \hat{\theta}, a, \theta),$$

for any $(x, \hat{\theta}) \in E_Y$ *and* $t = T-1, \ldots, 0$.

Outline for Numerical Implementation

We want to find a numerical solver for

$$V(t,x,\hat{\theta}) = \inf_{a\in A} \sup_{\theta\in\tau(t,\hat{\theta})} \mathbb{E}_\theta[V(t+1,\mathbf{T}(t,x,\hat{\theta},a,Z^{\theta}_{t+1})].$$

What needs to be done:

- Discretization of the continuous state space for Y, accompanied by interpolation in order to evaluate $V(t,x,\hat{\theta})$
- Approximation of the integral since integrand is not analytically available
- Approximation of the optimizers $a^*(x,\hat{\theta})$ and $\hat{\theta}^*(x,\hat{\theta})$
- The key idea is to recursively construct a functional approximation $\hat{V}(t,\cdot)$ that is used for interpolation and prediction

Curse of Dimensionality

As far as we know, no existing schemes are available when the state dimension is higher than 2. Specific difficulties include:

- Traditionally one constructs a grid of $(x, \hat{\theta})$-values, which is extremely inefficient for $k + d > 2$ and essentially impossible for $k + d > 4$
- Parametric representation of $\hat{V}$ (eg. in terms of polynomials in x and $\hat{\theta}$) is difficult for $k + d > 2$ and brings the concern for overfitting/underfitting
- The control a affects the evolution of the state x and prevents direct simulation of X as is done in the popular regression Monte Carlo paradigm

Existing Methods for Discretization and Interpolation

- Grid is immediately out of picture since $k + d$ will most likely be greater than 2
- Monte-Carlo based paradigm: simulate N trajectories according to a fixed measure $\mathbb{Q}^0$ and solve Bellman equation pathwise
- Such paradigm uses pre-specified $\mathbb{Q}^0$ that leads to non-adaptive experimental design
- Accompanied linear interpolation works very badly for high dimensional problem
- Estimated value functions are not smooth
- No way to deal with out-of-sample path

Spatial Modeling and Statistical Surrogates

- After discretization, $V(t+1,\cdot)$ is only evaluated at sampled sites
- Popular interpolation methods will have to go through all sites all the time, which is very expensive for high dimensions
- If two state points y^1 and y^2 are close, then $V(t+1,y^1)$ and $V(t+1,y^2)$ should also be close
- Leverage already obtained solutions of similar optimization problems
- Build a spatial statistical model $\hat{V}(t+1,\cdot)$ for $V(t+1,\cdot)$ over the domain by learning the correlation structure of $V(t+1,\cdot)$ at sample sites

Our Contribution

- We develop a machine learning framework tailored to generic stochastic min-max optimization problem.
- We recast the task of solving the Bellman equation as a statistical learning problem of fitting a surrogate (i.e. a statistical model) for $(x, \hat{\theta}) \mapsto \hat{V}(t, x, \hat{\theta})$.
- Gaussian Process surrogate for $\hat{V}$, coupled with an adaptive Experimental Design.
- Gaussian Process surrogate for $\hat{a}$ allows us to predict the optimal control (for out-of-sample) paths without directly optimizing.

Basic Loop

$$V(t,x,\hat{\theta}) = \inf_{a\in A} \sup_{\theta\in\mathcal{T}(t,\hat{\theta})} \mathbb{E}_\theta[V(t+1,f(x,a,Z^\theta_{t+1}))] =: F(V(t+1,\cdot),x,\hat{\theta})$$

We have a fit – predict – optimize – fit loop:

1. (Assume that the surrogate $\hat{V}(t+1,\cdot)$ has been fitted)
2. Select an experimental Design $\mathcal{D}_t$ of N_t sites y^n, $n = 1,\dots,N_t$;
3. Solve the optimization problem at each y^n, using $\mathrm{predict}(\hat{V}(t+1,y'))$ for the expectation. This yields the outputs $e^n = F(\hat{V}(t+1,\cdot),y^n)$ and optimal control a^n at y^n;
4. Fit $\hat{V}(t,\cdot)$ based on data $(y^{1:N_t},e^{1:N_t})$ and $\hat{a}(t,\cdot)$ based on $(y^{1:N_t},a^{1:N_t})$;
5. Goto 1: start the next recursion for $t-1$

Gaussian Process Surrogates

- Non-parametric regression, similar to splines or kernel regression
- Multivariate Gaussian structure to describe the shape of $\hat{V}$ and $\hat{a}$: covariance matrix $\mathbf{K}_{i,j} := K(y^i, y^j)$
- Train the model corresponds to applying the Gaussian conditional equations, and posterior is still Gaussian
- Statistical model for $\hat{V}$ and $\hat{a}$ are described by the corresponding posterior means

$$m_*(y_*) = k(y_*)[\mathbf{K} + \sigma^2 \mathbf{I}]^{-1} \vec{e},$$
$$s_*(y_*, y'_*) = K(y_*, y'_*) - k(y_*)[\mathbf{K} + \sigma^2 \mathbf{I}]^{-1} k(y'_*).$$

Fitting the emulator $=$ learning the hyperparameters in $\mathbf{K}$.

Experimental Design

- Accuracy of $\hat{V}(t, y_*)$ at given input y_* is directly related to the density of the design $\mathcal{D}_t$ around y_*
- Extrapolate for inputs far away from the simulated sites and depend on the prior mean
- Design $\mathcal{D}_t$ should be statistically sound by avoiding points that can't be observed in practice
- $\mathcal{D}_t$ should increase in time to avoid extrapolation when solving the Bellman equation backwards in time
- For a design, the choosing sites must fill in the space well
- Adaptively choose the design based on result of the previous step of numerical recursion

Dynamic Optimal Portfolio Selection

Recall the dynamic optimal portfolio selection problem where an investor wants to maximize the expected utility $U(W_T)$ of the terminal wealth.

- log return of the risky asset $Z_t = \mu + \sigma\varepsilon_t$, where ε_t are i.i.d. $\mathcal{N}(0,1)$
- Dynamics of the wealth process

$$\begin{aligned} W_{t+1} &= W_t(1 + r^f + \varphi_t(e^{Z_{t+1}} - 1 - r^f)) \\ &= W_t(1 + r^f + \varphi_t(e^{\mu+\sigma\varepsilon_t} - 1 - r^f)), \quad t \in \mathcal{T}', \ W_0 = w_0. \end{aligned}$$

- Adaptive Robust Stochastic Control Problem:

$$\sup_{\varphi\in\mathcal{A}} \inf_{\mathbb{Q}\in\mathcal{Q}^{\varphi,\Psi}} \mathbb{E}_{\mathbb{Q}}[U(W_T^{\varphi})].$$

Confidence Region

The MLE $\hat{\theta}_t = (\hat{\mu}_t, \hat{\sigma}_t^2)$ of the unknown parameter $\theta^* = (\mu^*, (\sigma^*)^2)$ can be expressed in the following recursive way:

$$
\begin{aligned}
\hat{\mu}_{t+1} &= \frac{t}{t+1}\hat{\mu}_t + \frac{1}{t+1} Z_{t+1}, \\
\hat{\sigma}_{t+1}^2 &= \frac{t}{t+1}\hat{\sigma}_t^2 + \frac{t}{(t+1)^2}(\hat{\mu}_t - Z_{t+1})^2
\end{aligned}
$$

Due to asymptotic normality of the MLEs, we have the recursive $1-\alpha$ confidence regions take the form

$$
\boldsymbol{\Theta}_t = \tau_\alpha(t, \hat{\mu}, \hat{\sigma}) := \left\{ (\mu, \sigma^2) \in \mathbb{R}^2 : \frac{t}{\hat{\sigma}^2}(\hat{\mu} - \mu)^2 + \frac{t}{2\hat{\sigma}^4}(\hat{\sigma}^2 - \sigma^2)^2 \le \kappa_\alpha \right\}
$$

with κ_α being the $(1-\alpha)$–quantile of the χ_2^2 distribution.

Bellman Equation

The Markov decision process $Y_t = (W_t^{\varphi}, \hat{\mu}_t, \hat{\sigma}_t)$ has dynamics

$$Y_{t+1} = \mathbf{T}(t, Y_t, \varphi_t, Z_{t+1})$$

where

$$\mathbf{T}(t, w, \widehat{\mu}, \widehat{\sigma}, a, z) = \left(w(1 + r^f + az), \frac{t}{t+1}\widehat{\mu} + \frac{1}{t+1}z, \sqrt{\frac{t}{t+1}\widehat{\sigma}^2 + \frac{t}{(t+1)^2}(\widehat{\mu} - z)^2} \right)$$

The corresponding adaptive robust Bellman equation becomes

$$V(T, w, \widehat{\mu}, \widehat{\sigma}) = u(w),$$
$$V(t, w, \widehat{\mu}, \widehat{\sigma}) = \sup_{a \in A} \inf_{(\mu,\sigma) \in \tau_\alpha(t, \widehat{\mu}, \widehat{\sigma})} \mathbb{E}_{\mu,\sigma} \left[W_{t+1} \left(\mathbf{T}(t, w, \widehat{\mu}, \widehat{\sigma}, a, Z_{t+1}) \right) \right]$$

Dimension-reduced Bellman Equation

By choosing the CRRA utility function $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$, we can be show that the ratio $\overline{V}$ defined as $\overline{V}(t, \hat{\mu}, \hat{\sigma}) := V(t, w, \hat{\mu}, \hat{\sigma})/w^{1-\gamma}$ satisfy the following backward recursion

$$\overline{V}(T, \hat{\mu}, \hat{\sigma}) = \frac{1}{1-\gamma},$$

$$\overline{V}(t, \hat{\mu}, \hat{\sigma}) = \inf_{a \in A} \sup_{(\mu,\sigma) \in \tau_\alpha(t,\hat{\mu},\hat{\sigma})} \mathbb{E}\Big[\big(1 + r + a\big(e^{\mu+\sigma\varepsilon_{t+1}} - 1 - r\big)\big)^{1-\gamma}$$

$$\overline{V}\Big(t+1, \frac{t}{t+1}\hat{\mu} + \frac{1}{t+1}(\mu + \sigma\varepsilon_{t+1}), \frac{t}{t+1}\hat{\sigma}^2 + \frac{t}{(t+1)^2}(\hat{\mu} - \mu - \sigma\varepsilon_{t+1})^2\Big)\Big].$$

Algorithm Part I

Backward recursion over the time steps for $t = T - 1, \dots, 0$

1. Create a design $\mathcal{D}_t = (\mu_t^{1:N_t}, \theta_t^{1:N_t})$ that will be used to estimate $\hat{V}(t, \cdot)$. The design is based on Monte Carlo paradigm, Sobol space filling, and adaptively adding more points.
2. For $n = 1, 2, \dots, N_t$, let $f_2(a, \mu_t^n, \sigma_t^n) = \inf_{(\mu,\sigma)\in\tau(t,\mu_t^n,\sigma_t^n)} \hat{E}\Big[\hat{V}(t+1, \mathbf{T}(t, \mu_t^n, \sigma_t^n, a, \mu + \sigma\varepsilon))\Big]$. $\hat{E}$ is an approximate operator to estimate the expectation using quadrature rule or Monte Carlo.
3. Let $e_t^n := \sup_{a\in A} f_2(a, \mu_t^n, \sigma_t^n)$. Record the estimated optimal control a_t^n.
4. Build a GP model $\hat{V}(t, \cdot)$ for the link between (μ_t^n, σ_t^n) and (e_t^n). Build a GP model $\hat{a}(t, \cdot)$ for the link between (μ_t^n, σ_t^n) and (a_t^n).

Algorithm Part II

Forward simulation on fresh paths, over the time steps for $t \in \mathcal{T}'$, to evaluate the performance of the strategy assuming the true probability model.

1. For path n, draw iid sequence of Z_t^n via ε_t^n.
2. Using the GP model to predict the control $a_{t-1}^n = \hat{a}(t-1, \mu_{t-1}^n, \sigma_{t-1}^n)$.
3. Update the states according to $(w_{t+1}^n, \mu_{t+1}^n, \sigma_{t+1}^n) = \mathbf{T}(t+1, w_t^n, \mu_t^n, \sigma_t^n, a_{t-1}^n, Z_{t+1}^n)$.

The final answer is the average $\underline{V}(0, w_0, \mu_0, \sigma_0) = \frac{1}{N}\sum_{n=1}^{N} u(w_T^n)$.

Simulation Design

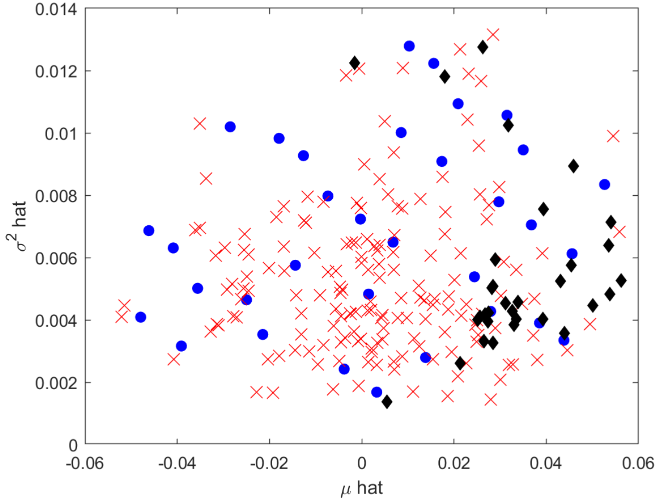

GP Fitting and Extrapolation

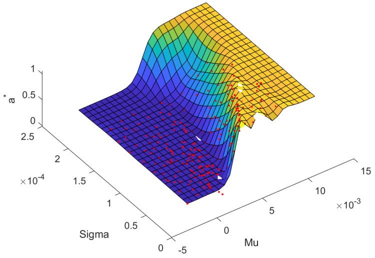

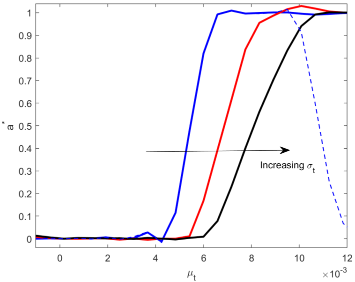

- In the area where points are sampled, GP surrogate works very well
- For GP, outputs corresponding to inputs far away from simulated sites are decided by mean function
- We shift (a_t^n) by a sigmoid function $M(\mu) = (1 + e^{-(A\mu - B)})^{-1}$ and set the mean function as 0
- It improves the extrapolation but doesn't affect predictions inside the training domain

Stability of Hyperparameters

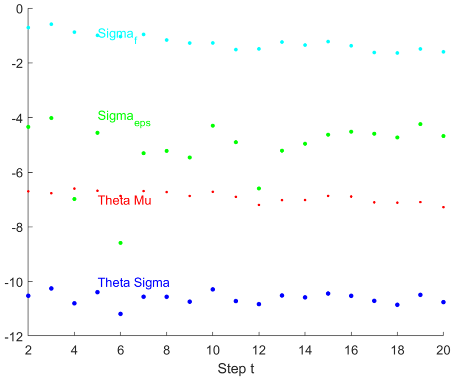

A major worry is that the GP model is mis-estimated, which can be diagnosed by an outlier in the hyperparameters.

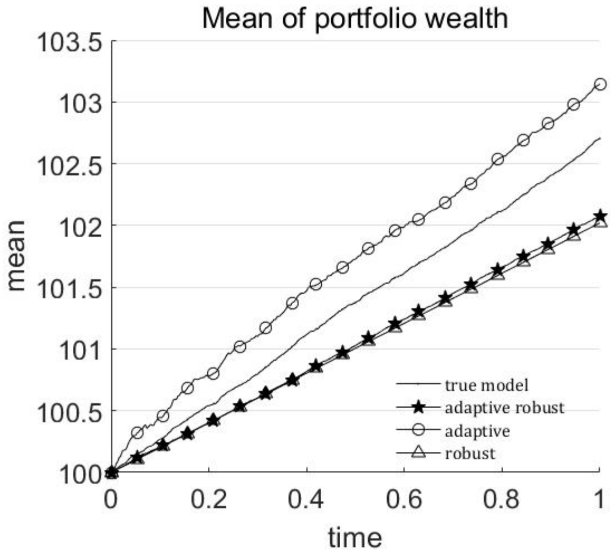
Mean of portfolio wealth
103.5
103
102.5
102
101.5
101
100.5
100
mean
0
0.2
0.4
0.6
0.8
1
time
true model
adaptive robust
adaptive
robust

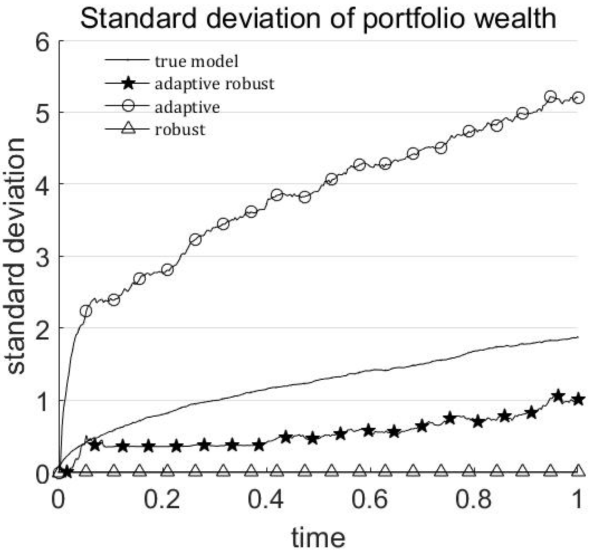
Standard deviation of portfolio wealth
true model
adaptive robust
adaptive
robust
standard deviation
time
0
1
2
3
4
5
6
0
0.2
0.4
0.6
0.8
1

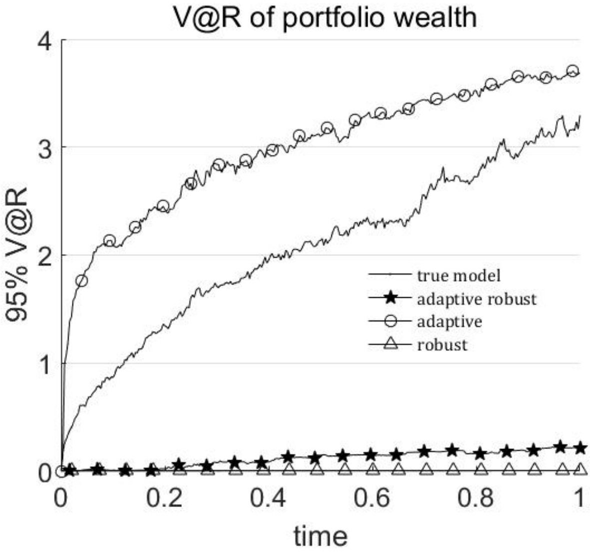

$$V@R(V_T) = \inf\{v \in \mathbb{R} : \mathbb{P}_{\theta^*}(V_T + v < 0) \leq 95\%\}$$

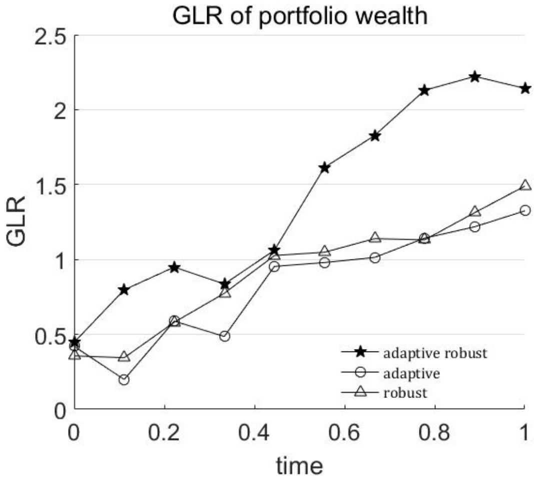

$$GLR(V) = \begin{cases} \frac{\mathbb{E}_{\theta^*}[e^{-rT}V_T - V_0]}{\mathbb{E}_{\theta^*}[(e^{-rT}V_T - V_0)^-]}, & e^{-rT}V_T - V_0 > 0, \\ 0, & \text{otherwise.} \end{cases}$$

Adaptive Robust Hedging

In a similar setup, the investor wants to minimize the expected superhedging risk $\ell[(\Phi(S_T) - W_T)^+]$.

- $\ell : \mathbb{R}_+ \to \mathbb{R}_+$ is an increasing function such that $\ell(0) = 0$
- State process: $(S_t, W_t, \hat{\mu}_t, \hat{\sigma}_t)$
- Need to modify the Monte Carlo paradigm as state W_t depends on control such that direct simulation is impossible
- $\mathcal{D}_t = (S_t^{1:N_t}, \mu_t^{1:N_t}, \sigma_t^{1:N_t}, \mathsf{BS}(S_t)^{1:N_t}) \bigcup (S_t^{1:N_t}, \mu_t^{1:N_t}, \sigma_t^{1:N_t}, 0.6 * \mathsf{BS}(S_t)^{1:N_t}) \bigcup (S_t^{1:N_t}, \mu_t^{1:N_t}, \sigma_t^{1:N_t}, 1.4\mathsf{BS}(S_t)^{1:N_t})$
- Randomization of the starting simulation sites are important due to strong correlation between S_t and $\hat{\mu}_t$

Thank You !

The end of the talk ...
but not of the story