Adaptive Robust Control Under Model Uncertainty

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

Joint work with
T.R. Bielecki, I. Cialenco, A. Cousin, M. Jeanblanc

Mathematical Finance Colloquia
University Of Southern California
Oct. 23 2017

Motivations

- To control the risk due to model uncertainty
  (error in model estimation or model misspecification)

- To solve control problems when the true law of the underlying stochastic phenomenon is unknown and belongs to a parameterized family of probability laws

- Existing robust methodologies may be overly conservative when applied to unknown system

- Available robust frameworks do not consider reduction of uncertainty

Main Goals

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

- The methodology can be applied to any Markov decision process under model uncertainty, and in particular, to some financial problems

- To build a theory of recursive construction of confidence regions

Preliminaries

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

General 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 $\varphi = \{\varphi_t,\ t \in \mathcal{T}'\}$); L is a measurable functional (loss function, utility function, etc).

Example. Find a self-financing portfolio that maximizes the expected utility of the terminal wealth

$$\sup_{\varphi \in \mathcal{A}} \mathbb{E}_{\theta^*}(u(V_T^{\varphi})),$$

where V_T^{φ} is the wealth at time T using strategy φ, u is an utility function, $\mathcal{A}$ is the set of s.f. strategies.

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

Classical Approaches

Robust control problem:

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

- Select the best strategy φ over the worst possible model.
- Başar and Bernhard (1995), Hansen et al. (2006), Hansen and Sargent (2008).
- If the '$\sup$ model' is far from the true model, this approach could produce really bad results.

Strong robust control problem:

$$\inf_{\varphi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathcal{Q}^{\varphi, \boldsymbol{\Psi}_K}} \mathbb{E}_{\mathbb{Q}}(L(X, \varphi)).$$

- $\boldsymbol{\Psi}_K$ is the set of strategies chosen by Knightian adversary or the nature.
- $\mathcal{Q}^{\varphi, \boldsymbol{\Psi}_K}$ is a set of probability measures on canonical space that represents all possible models resulting from φ and $\boldsymbol{\Psi}_K$.
- Bayraktar, Cosso and Pham (2014), Sîrbu (2014).

Bayesian adaptive control problem:

$$\inf_{\varphi \in \mathcal{A}} \int_{\mathbf{\Theta}} \mathbb{E}_\theta(L(X,\varphi))\nu_0(d\theta).$$

- The parameter θ is modeled as a random variable taking values in $\mathbf{\Theta}$ with a prior distribution ν_0.
- Kumar, Varaiya (1986), Runggaldier et al. (2002), Corsi et al. (2007).
- Note that

$$\inf_{\varphi \in \mathcal{A}} \sup_{\theta \in \mathbf{\Theta}} \mathbb{E}_\theta\left(L(X,\varphi)\right) = \inf_{\varphi \in \mathcal{A}} \sup_{\nu_0 \in \mathcal{P}(\mathbf{\Theta})} \int_{\mathbf{\Theta}} \mathbb{E}_\theta\left(L(X,\varphi)\right)\nu_0(d\theta).$$

No reduction of uncertainty about θ.

Adaptive control problem:

For each $\theta \in \mathbf{\Theta}$ solve

$$\inf_{\varphi \in \mathcal{A}} \mathbb{E}_\theta \left(L(X, \varphi) \right).$$

- Let φ^θ be a corresponding optimal control.
- At each time t, compute an $\mathscr{F}_t$-measurable point estimate $\widehat{\theta}_t$ of θ^*.
- Apply the control value $\varphi_t^{\widehat{\theta}_t}$.
- Known to have poor performance for finite time horizon problems
- Kumar and Varaiya (1986), Chen and Guo (1991)

- 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.**"

We propose **Adaptive Robust Control Methodology**

$$\inf_{\varphi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathcal{Q}_{h_0}^{\varphi, \boldsymbol{\Psi}}} \mathbb{E}_{\mathbb{Q}} \left(L(X, \varphi) \right),$$

where $\mathcal{Q}_{h_0}^{\varphi, \boldsymbol{\Psi}}$ is a family of probability measures constructed in a way that allows for *dynamic reduction of uncertainty* about θ^*.

We chose the family $\mathcal{Q}_{h_0}^{\varphi, \boldsymbol{\Psi}}$ in terms of *confidence regions* for the parameter θ^*.

Without uncertainty

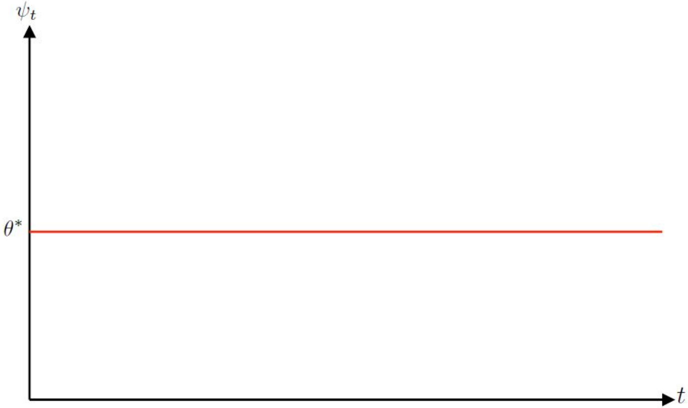

Robust

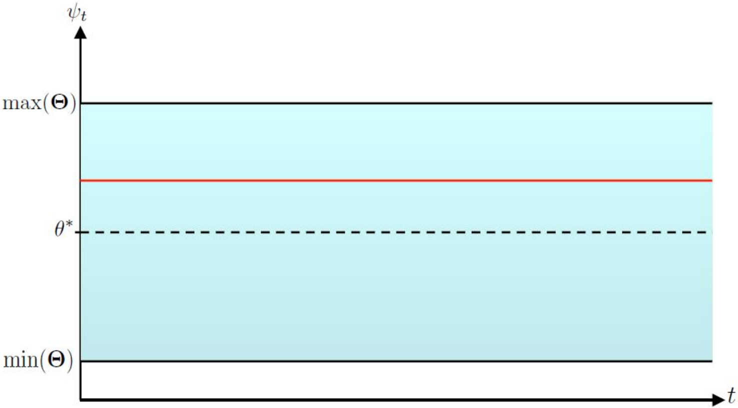

Strong robust

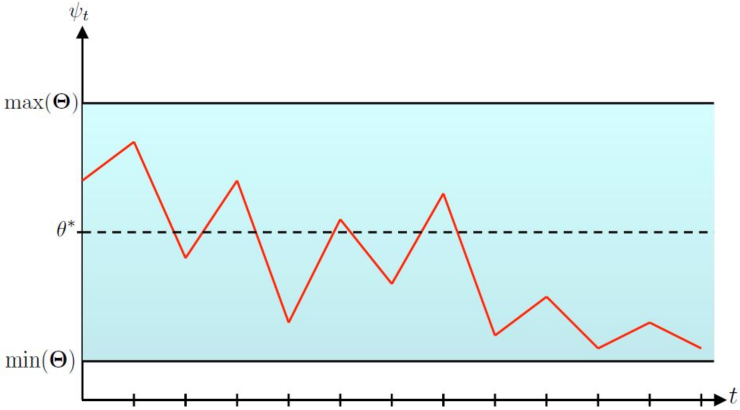

Adaptive Robust

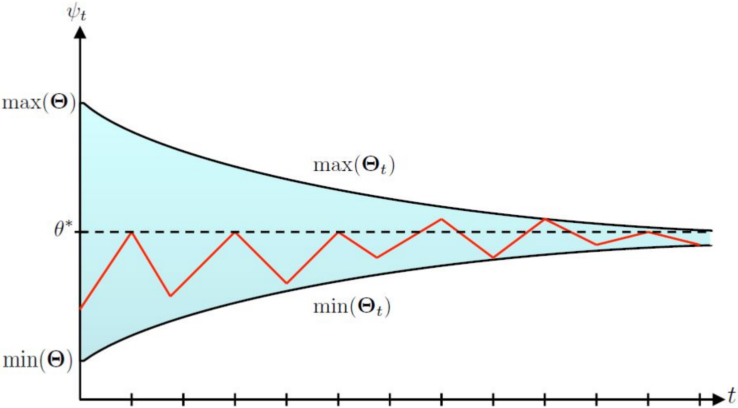

We assume that the observed state process X follows the dynamics

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

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

- $\mathbb{F}$–adapted,
- observed,
- i.i.d. under $\mathbb{P}_\theta,\ \theta \in \mathbf{\Theta} \subset \mathbb{R}^d$,
- the law of Z_1 is unknown but it belongs to a family of parameterized distributions $\{\mathbb{P}_\theta,\ \theta \in \mathbf{\Theta}\}$.

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

The adaptive robust control methodology relies *essentially* on recursive construction of confidence regions. In Bielecki, Cialenco and Chen (2016), for time-homogenous 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 $\mathbf{\Theta}_t$ of θ^* can be constructed by a deterministic rule:

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

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

We consider the augmented state process

$$Y_t := (X_t, \hat{\theta}_t), \quad t \in \mathcal{T},$$

with state space $E_Y := \mathbb{R}^n \times \mathbb{R}^d$, and dynamics

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

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

We define the histories

$$H_t := (Y_0, \ldots, Y_t) \in \mathbf{H}_t := \underbrace{E_Y \times E_Y \times \ldots \times E_Y}_{t+1 \text{ times}}.$$

The sequence of confidence regions $\mathbf{\Theta}_t$, $t \in \mathcal{T}$, will represent the learning about θ^* based on the observed history H_t, $t \in \mathcal{T}$.

A robust control problem can be viewed as a game between a controller and the nature.

The controller plays history-dependent strategies φ that belong to

$$\mathcal{A} = \{(\varphi_t)_{t\in\mathcal{T}'} \mid \varphi_t : \boldsymbol{H}_t \to A, t \in \mathcal{T}'\}.$$

Adaptive robust control: the nature plays history-dependent strategies ψ that belong to

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

Remark. The *strong robust case* corresponds to

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

Define the Markov transition probability kernel on $\mathcal{E}_Y$

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

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

Using Ionescu-Tulcea theorem, 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

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

where

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

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 $\varphi^* = (\varphi_t^*(h_t))_{t\in\mathcal{T}'}$ *of*

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

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

$$W_T(y) = \ell(x),$$
$$W_t(y) = \inf_{a\in A} \sup_{\theta\in\tau_\alpha(t,\tilde{\theta})} \int_{E_Y} W_{t+1}(y)Q(dy \mid t, y, a, \theta),$$

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

Note that the optimal strategy at time t is such that $\varphi_t^*(h_t) = \varphi_t^*(y_t)$.

Proof:

We define the functions U_t and U_t^* as follows: for $\varphi^t \in \mathcal{A}^t$ and $h_t \in H_t$

$$
\begin{aligned}
U_t(\varphi^t, h_t) &= \sup_{\mathbb{Q} \in \mathcal{Q}_{h_t}^{\varphi^t, \boldsymbol{\Psi}^t}} \mathbb{E}_{\mathbb{Q}} \ell(X_T), \ t \in \mathcal{T}', \\
U_t^*(h_t) &= \inf_{\varphi^t \in \mathcal{A}^t} U_t(\varphi^t, h_t), \ t \in \mathcal{T}', \\
U_T^*(h_T) &= \ell(x_T).
\end{aligned}
$$

We proceed via backward induction in $t = T, T-1, \ldots, 1, 0$.

Take $t = T$. Clearly, $U_T^*(h_T) = W_T(y_T)$. For $t = T-1$ we have

$$
\begin{aligned}
U_{T-1}^*(h_{T-1}) &= \inf_{\varphi^{T-1} = \varphi_{T-1} \in \mathcal{A}^{T-1}} \ \sup_{\mathbb{Q} \in \mathcal{Q}_{h_{T-1}}^{\varphi^{T-1}, \boldsymbol{\Psi}^{T-1}}} \mathbb{E}^{\mathbb{Q}} \ell(X_T) \\
&= \inf_{\varphi^{T-1} = \varphi_{T-1} \in \mathcal{A}^{T-1}} \ \sup_{\theta \in \tau(T-1, c_{T-1})} \int_{E_Y} U_T^*(h_{T-1}, y) Q(dy \mid T-1, y_{T-1}, \varphi_{T-1}(h_{T-1}), \theta \\
&= \inf_{\varphi^{T-1} = \varphi_{T-1} \in \mathcal{A}^{T-1}} \ \sup_{\theta \in \tau(T-1, c_{T-1})} \int_{E_Y} W_T(y)\, Q(dy \mid T-1, y_{T-1}, \varphi_{T-1}(h_{T-1}), \theta) \\
&= \inf_{a \in A} \sup_{\theta \in \tau(T-1, c_{T-1})} \int_{E_Y} W_T(y)\, Q(dy \mid T-1, y_{T-1}, a, \theta) = W_{T-1}(y_{T-1}).
\end{aligned}
$$

For $t = T-1, \ldots, 1, 0$ we have by induction

$$\begin{aligned}
U_t^*(h_t) &= \inf_{\varphi^t \in \mathcal{A}^t} \sup_{\mathbb{Q} \in \mathcal{Q}_{h_t}^{\varphi^t, \Psi^t}} \mathbb{E}^{\mathbb{Q}} \ell(X_T) \\
&= \inf_{\varphi^t = (\varphi_t, \varphi^{t+1}) \in \mathcal{A}^t} \sup_{\theta \in \tau(t, c_t)} \int_{E_Y} \sup_{\widehat{\mathbb{Q}} \in \mathcal{Q}_{h_t, y}^{\varphi^{t+1}, \Psi^{t+1}}} \mathbb{E}^{\widehat{\mathbb{Q}}} \ell(X_T) Q(dy \mid t, y_t, \varphi_t(h_t), \theta) \\
&\geq \inf_{\varphi^t = (\varphi_t, \varphi^{t+1}) \in \mathcal{A}^t} \sup_{\theta \in \tau(t, c_t)} \int_{E_Y} U_{t+1}^*(h_t, y)\, Q(dy \mid t, y_t, \varphi_t(h_t), \theta) \\
&= \inf_{a \in A} \sup_{\theta \in \tau(t, c_t)} \int_{E_Y} U_{t+1}^*(h_t, y)\, Q(dy \mid t, y_t, a, \theta) \\
&= \inf_{a \in A} \sup_{\theta \in \tau(t, c_t)} \int_{E_Y} W_{t+1}(y)\, Q(dy \mid t, y_t, a, \theta) = W_t(y_t).
\end{aligned}$$

Now, fix $\epsilon > 0$, and let $\varphi^{t+1,\epsilon}$ denote an ϵ-optimal control process starting at time $t+1$, so that

$$U_{t+1}(\varphi^{t+1,\epsilon}, h_{t+1}) \leq U^*_{t+1}(h_{t+1}) + \epsilon.$$

Then we have

$$\begin{aligned}
U^*_t(h_t) &= \inf_{\varphi^t \in \mathcal{A}^t} \sup_{\mathbb{Q} \in \mathcal{Q}^{\varphi^t, \Psi^t}_{h_t}} \mathbb{E}^{\mathbb{Q}} \ell(X_T) \\
&= \inf_{\varphi^t = (\varphi_t, \varphi^{t+1}) \in \mathcal{A}^t} \sup_{\theta \in \tau(t, c_t)} \int_{E_Y} \sup_{\widehat{\mathbb{Q}} \in \mathcal{Q}^{\varphi^{t+1}, \Psi^{t+1}}_{h_t, y}} \mathbb{E}^{\widehat{\mathbb{Q}}} \ell(X_T) Q(dy \mid t, y_t, \varphi_t(h_t), \theta) \\
&\leq \inf_{\varphi^t = (\varphi_t, \varphi^{t+1}) \in \mathcal{A}^t} \sup_{\theta \in \tau(t, c_t)} \int_{E_Y} \sup_{\widehat{\mathbb{Q}} \in \mathcal{Q}^{\varphi^{t+1,\epsilon}, \Psi^{t+1}}_{h_t, y}} \mathbb{E}^{\widehat{\mathbb{Q}}} \ell(X_T) Q(dy \mid t, y_t, \varphi_t(h_t), \theta) \\
&\leq \inf_{a \in A} \sup_{\theta \in \tau(t, c_t)} \int_{E_Y} U^*_{t+1}(h_t, y)\, Q(dy \mid t, y_t, a; \theta) + \epsilon \\
&= \inf_{a \in A} \sup_{\theta \in \tau(t, c_t)} \int_{E_Y} W_{t+1}(y)\, Q(dy \mid t, y_t, a; \theta) + \epsilon = W_t(y_t) + \epsilon.
\end{aligned}$$

Since ϵ was arbitrary, the proof is done.

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 $u(V_T)$ of the terminal wealth.

- r^f - the constant risk free rate
- e^{Z_t} - the excess 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

$$V_{t+1} = V_t(1 + r^f + \varphi_t(e^{Z_{t+1}} - 1 - r)), \quad t \in \mathcal{T}',$$

with $V_0 = v_0$, and where φ_t is the proportion of the wealth invested in the risky asset from t to $t+1$.

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:

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

Due to asymptotic normality of the MLEs, we have

$$\frac{t}{\hat{\sigma}_t^2}(\hat{\mu}_t - \mu^*)^2 + \frac{t}{2\hat{\sigma}_t^4}(\hat{\sigma}_t^2 - (\sigma^*)^2)^2 \xrightarrow[t\to\infty]{d} \chi_2^2.$$

Consequently, the recursive $1-\alpha$ confidence regions take the form

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

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

The Markov decision process $Y_t = (V_t^{\varphi}, \hat{\mu}_t, \hat{\sigma}_t^2)$ has dynamics

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

where

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

The corresponding robust Bellman equation becomes

$$\begin{aligned} W_T(v, \widehat{\mu}, \widehat{\sigma}^2) &= u(v), \\ W_t(v, \widehat{\mu}, \widehat{\sigma}^2) &= \sup_{a \in A} \inf_{(\mu, \sigma^2) \in \tau_\alpha(t, \widehat{\mu}, \widehat{\sigma}^2)} \mathbb{E}_{\mu,\sigma}\left[W_{t+1}\left(\mathbf{T}(t, v, \widehat{\mu}, \widehat{\sigma}^2, a, Z_{t+1}) \right) \right] \end{aligned}$$

We consider the following setting:

- CRRA utility function $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$, with $\gamma = 40$,
- Return of the risk-free asset : $r^f = 0$,
- Excess return true mean and standard deviation : $\mu^* = 1\%$, $\sigma^* = 4\%$
- Initial guess on model parameters : $\widehat{\mu}_0 = 2\%$ and $\widehat{\sigma}_0 = 1\%$
- Initial endowment : $V_0 = 100$
- $A = [0, 1]$, i.e. optimal strategies at time t to be between 0 and 1.

It can be showed that for $t \in \mathcal{T}$ and any $c \in [a,b]$ the ratio $W_t(v,c)/v^{1-\gamma}$ does not depend on v, and that the functions $\widetilde{W}_t$ defined as $\widetilde{W}_t(c) := W_t(v,c)/v^{1-\gamma}$ satisfy the following backward recursion

$$\begin{cases} \widetilde{W}_T(c) = \frac{1}{1-\gamma}, \\ \widetilde{W}_t(c) = \inf_{a\in A} \sup_{\mu\in\tau_\alpha(t,c)} \mathbb{E}\Big[(1+r+a(e^{\mu+\sigma\varepsilon_{t+1}}-1-r))^{1-\gamma}\widetilde{W}_{t+1}(\frac{t}{t+1}c+\frac{1}{t+1}(\mu+\sigma\varepsilon_{t+1}) \\ t \in \mathcal{T}'. \end{cases} \tag{2.1}$$

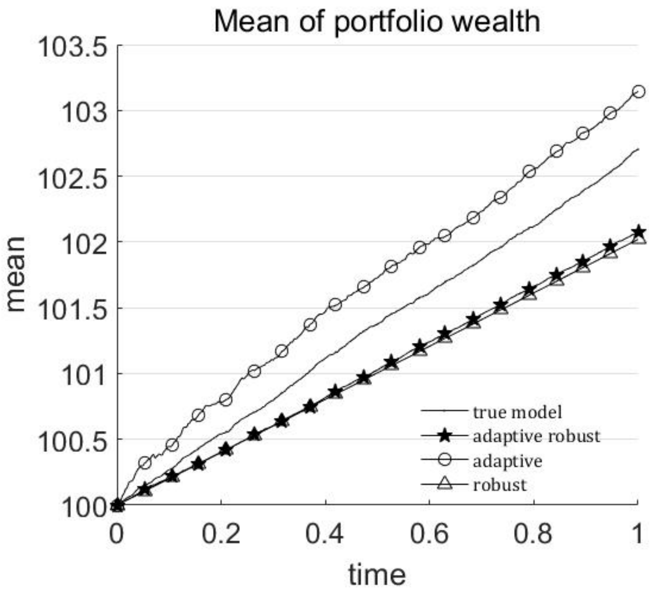
Mean of portfolio wealth
103.5
103
102.5
102
101.5
101
100.5
100
mean
0
0.2
0.4
0.6
0.8
1
time
true model
adaptive robust
adaptive
robust

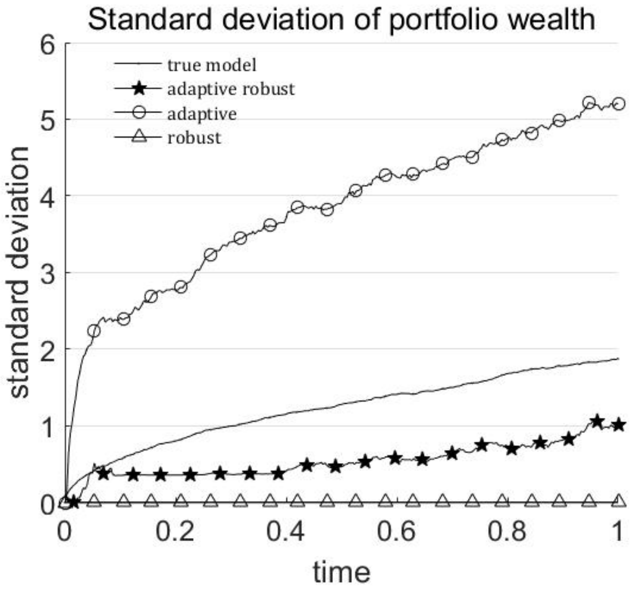
Standard deviation of portfolio wealth
standard deviation
time
true model
adaptive robust
adaptive
robust
0
1
2
3
4
5
6
0
0.2
0.4
0.6
0.8
1

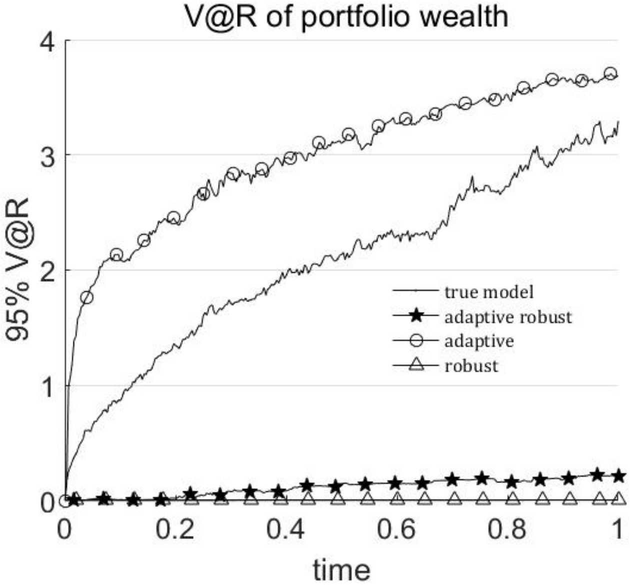

$$V@R(V_T) = \inf\{v \in \mathbb{R} : \mathbb{P}_{\theta^*}(V_T + v < 0) \leq 95\%\}$$

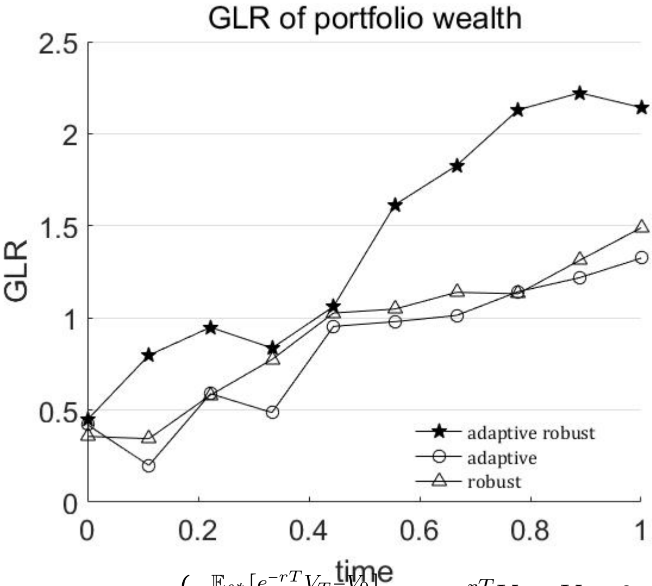

$$GLR(V) = \begin{cases} \frac{\mathbb{E}_{\theta^*}[e^{-rT}V_T - V_0]}{\mathbb{E}_{\theta^*}[(e^{-rT}V_T - V_0)^-]}, & e^{-rT}V_T - V_0 > 0, \\ 0, & \text{otherwise,} \end{cases}$$

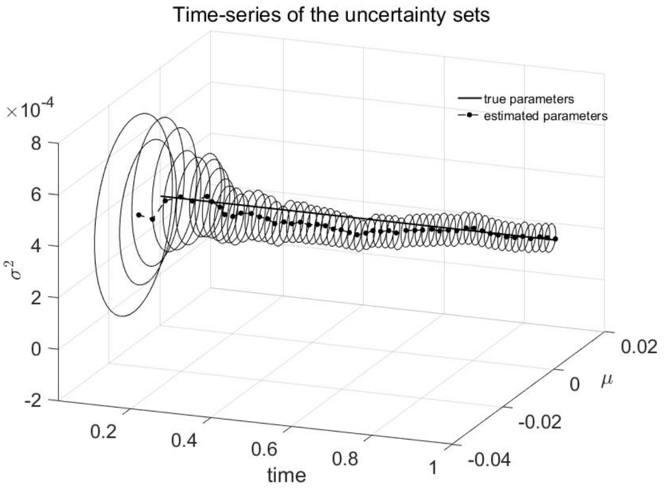

Time-series of the uncertainty sets
true parameters
estimated parameters
×10^-4
8
6
4
2
0
-2
σ^2
0.2
0.4
0.6
0.8
1
time
0.02
0
-0.02
-0.04
μ

Thank You !

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