Adaptive robust hedging under model uncertainty

Areski Cousin

ISFA, Université Lyon 1

Based on a work in progress with
T. Bielecki, T. Chen, I. Cialenco, M. Jeanblanc

MURF, 2016

Milan, November 11, 2016

General setting and motivation

- Robust control may be overly conservative when applied to the true unknown system
- We develop an adaptive robust methodology for solving a discrete-time Markovian control problem subject to Knightian uncertainty
- We focus on a financial hedging problem, but the methodology can be applied to any kind of Markov decision process under model uncertainty
- As in the classical robust case, the uncertainty comes from the fact that the true law of the driving process is only known to belong to a certain family of probability laws

General setting and motivation

- T : terminal date of our finite horizon control problem
- $\mathcal{T} = \{0, 1, 2, \ldots, T\}$: time grid
- $\mathcal{T}' = \{0, 1, 2, \ldots, T-1\}$: time grid without last date
- $S = \{S_t, \; t \in \mathcal{T}\}$: stochastic process that drives the random system

We assume that :

- S is observable and we denote by $\mathbb{F}^S = (\mathscr{F}_t^S, t \in \mathcal{T})$ its natural filtration.
- The law of S is not known but it belongs to a family of parametrized distributions $\mathbf{P}(\mathbf{\Theta}) := \{\mathbb{P}_\theta, \theta \in \mathbf{\Theta}\}$, $\mathbf{\Theta} \subset \mathbb{R}^d$
- The unknown (true) law of S is denoted by $\mathbb{P}_{\theta^*}$ and is such that $\theta^* \in \mathbf{\Theta}$

Model uncertainty occurs if $\mathbf{\Theta} \neq \{\theta^*\}$

We consider the following stochastic control problem

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

where

- $\mathcal{A}$ is a set of admissible control processes : $\mathbb{F}^S$–adapted processes $\varphi = \{\varphi_t,\ t \in \mathcal{T}'\}$
- L is a measurable functional (loss or error to minimize in our case)

Obviously, the problem cannot be dealt with directly since we do not know the value of θ^*

Robust control problem : Başar and Bernhard (1995), Hansen et al. (2006), Hansen and Sargent (2008)

$$\inf_{\varphi \in \mathcal{A}} \sup_{\theta \in \boldsymbol{\Theta}} \mathbb{E}_{\theta} \left(L(S, \varphi) \right). \tag{1}$$

- Best strategy over the worst possible model parameter in $\boldsymbol{\Theta}$
- If the true model is close to the best one, the solution to this problem could perform very badly

General setting and motivation

Strong robust control problem : Sirbu (2014), Bayraktar, Cosso and Pham (2014)

$$\inf_{\varphi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathcal{Q}^{\varphi, \boldsymbol{\Psi}_K}} \mathbb{E}_{\mathbb{Q}} \left(L(S, \varphi) \right), \qquad (2)$$

- $\boldsymbol{\Psi}_K$ is the set of strategies chosen by a Knightian adversary (the nature) that may keep changing the system distribution over time
- $\mathcal{Q}^{\varphi, \boldsymbol{\Psi}_K}$ represents all possible model dynamics resulting from φ and when nature plays strategies in $\boldsymbol{\Psi}_K$
- Solution is even more conservative than in the classical robust case
- No learning mechanism to reduce model uncertainty

General setting and motivation

Bayesian adaptive control problem : Kumar and Varaiya (1986), Runggaldier et al. (2002), Corsi et al. (2007)

$$\inf_{\varphi \in \mathcal{A}} \int_{\Theta} \mathbb{E}_\theta \left(L(S, \varphi) \right) \nu_0(d\theta). \tag{3}$$

- The unknown parameter θ is treated as an unobserved state variable with a prior distribution ν_0
- Control problem with partial information solved by transforming the original problem into a full-information separated problem (adding the posterior distribution as a new state variable)
- No reduction of uncertainty is really involved

Bayesian adaptive control vs robust control

Proposition

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

Thus, for any given prior distribution ν_0 we have :

$$\inf_{\varphi \in \mathcal{A}} \sup_{\theta \in \boldsymbol{\Theta}} \mathbb{E}_\theta \left(L(S, \varphi) \right) \geq \inf_{\varphi \in \mathcal{A}} \int_{\boldsymbol{\Theta}} \mathbb{E}_\theta \left(L(S, \varphi) \right) \nu_0(d\theta).$$

$\Rightarrow$ The Bayesian adaptive problem is less conservative than the classical robust one.

General setting and motivation

Adaptive control problem : Kumar and Varaiya (1986), Chen and Guo (1991)

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

$$\inf_{\varphi \in \mathcal{A}} \mathbb{E}_\theta \left(L(S, \varphi) \right). \tag{4}$$

- Let φ^θ be a corresponding optimal control
- At each time t, compute a point estimate $\widehat{\theta}_t$ of θ^*, using a chosen, $\mathscr{F}_t^S$ measurable estimator and apply control value $\varphi_t^{\widehat{\theta}_t}$.
- Known to have poor performance for finite horizon problems

Problem : Hedging a short position on an European-type option with maturity T, payoff function Φ and underlying asset S with price dynamics

$$S_0 = s_0 \in (0, \infty),$$
$$S_{t+1} = Z_{t+1} S_t, \qquad t \in \mathcal{T}'$$

where

- $Z = \{Z_t,\ t = 1, \ldots, T\}$ is a non-negative random process
- Under each measure $\mathbb{P}_\theta$, Z_{t+1} is independent from $\mathscr{F}_t^S$ for each $t \in \mathcal{T}$
- The true law $\mathbb{P}_{\theta*}$ of Z is not known.

Hedging under model uncertainty

Hedging is made using a self-financing portfolio composed of the underlying risky asset S and of a risk-free asset (with constant value equal to 1).

The hedging portfolio has the following dynamics

$$\begin{aligned} V_0 &= v_0, \\ V_{t+1} &= V_t + \varphi_t(S_{t+1} - S_t), \quad t = 0, \ldots, T-1 \end{aligned}$$

Exact replication is out of reach in our setting (v_0 may be too small), so that the nominal control problem (without uncertainty) is

$$\inf_{\varphi \in \mathcal{A}} \mathbb{E}_{\theta^*} \left(\ell[(\Phi(S_T) - V_T(\varphi))^+] \right),$$

where l is a loss function, i.e., an increasing function such that $\ell(0) = 0$ (shortfall risk minimization approach)

The methodology relies on recursive construction of confidence regions. We assume that :

1) A point estimator $\widehat{\theta}_t$ of θ^* can be constructed recursively

$$\widehat{\theta}_0 = \theta_0,$$
$$\widehat{\theta}_{t+1} = R(t, \widehat{\theta}_t, Z_{t+1}), \quad t = 0, \ldots, T-1$$

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

2) An approximate α-confidence region $\mathbf{\Theta}_t$ of θ^* can be constructed from $\widehat{\theta}_t$ by a deterministic rule :

$$\mathbf{\Theta}_t = \tau(t, \widehat{\theta}_t)$$

where $\tau(t, \cdot) : \mathbb{R}^d \to 2^{\mathbf{\Theta}}$ is a deterministic set valued function. The region $\mathbf{\Theta}_t$ should be such that $\mathbb{P}_{\theta^*}(\theta^* \in \mathbf{\Theta}_t) \approx 1 - \alpha$ and $\lim_{t\to\infty} \mathbf{\Theta}_t = \{\theta^*\}$ where the convergence is understood $\mathbb{P}^{\theta^*}$ almost surely, and the limit is in the Hausdorff metric.

Adaptive robust control methodology

We consider the following (augmented) state process

$$X_t = (S_t, V_t, \widehat{\theta}_t), \quad t \in \mathcal{T}$$

with state space $E_X := \mathbb{R}_+ \times \mathbb{R} \times \mathbb{R}^d$.

In our hedging problem, $X = (S, V, \widehat{\theta})$ is a Markov process with dynamics :

$$\begin{aligned} S_{t+1} &= Z_{t+1} S_t, \\ V_{t+1} &= V_t + \varphi_t S_t (Z_{t+1} - 1), \\ \widehat{\theta}_{t+1} &= R(t, \widehat{\theta}_t, Z_{t+1}) \end{aligned}$$

We denote by

$$Q(B \mid t, x, a, \theta) := \mathbb{P}_\theta(X_{t+1} \in B \mid X_t = x, \varphi_t = a)$$

the time-t Markov transition kernel under probability $\mathbb{P}_\theta$ when strategy a is applied

Adaptive robust control methodology

Let us denote by

$$H_t := ((S_0, V_0, \widehat{\theta}_0), (S_1, V_1, \widehat{\theta}_1), \ldots, (S_t, V_t, \widehat{\theta}_t)), \ t \in \mathcal{T},$$

the history of the state process up to time t.

Note that, for any admissible trading strategy φ, H_t is $\mathscr{F}_t^S$ measurable and

$$H_t \in \mathbf{H}_t := \underbrace{E_X \times E_X \times \ldots \times E_X}_{t+1 \text{ times}}.$$

We denote by

$$h_t = (x_0, x_1, \ldots, x_t) = (s_0, v_0, c_0, s_1, v_1, c_1, \ldots, s_t, v_t, c_t)$$

a realization of H_t.

Adaptive robust control methodology

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

The controller plays history-dependent strategies φ that belong to

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

where φ_t is a measurable mapping.

Strong robust case : nature plays history-dependent strategies ψ that belong to

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

Adaptive robust case : nature plays history-dependent strategies ψ that belong

to

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

where $\mathbf{\Theta}_t = \tau(t, \widehat{\theta}_t)$ is the α-confidence region of θ^* at time t

Given that the controller plays φ and nature plays ψ, using Ionescu-Tulcea theorem, we define the canonical law of the state process X on E_X^T as

$$\mathbb{Q}_{h_0}^{\varphi,\psi}(B_1,\dots,B_T) = \int_{B_1}\cdots\int_{B_T} Q(dx_T \mid T-1, x_{T-1}, \varphi_{T-1}(h_{T-1}), \psi_{T-1}(h_{T-1})) \cdots Q(dx_2 \mid 1, x_1, \varphi_1(h_1), \psi_1(h_1))\, Q(dx_1 \mid 0, x_0, \varphi_0(h_0), \psi_0(h_0)).$$

For a given strategy φ, we define

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

and

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

Adaptive robust control methodology

The strong robust hedging problem :

$$\inf_{\varphi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathcal{Q}_{h_0}^{\varphi, \Psi_K}} \mathbb{E}_{\mathbb{Q}} \left(\ell[(\Phi(S_T) - V_T)^+] \right)$$

The adaptive robust hedging problem :

$$\inf_{\varphi \in \mathcal{A}} \sup_{\mathbb{Q} \in \mathcal{Q}_{h_0}^{\varphi, \Psi_A}} \mathbb{E}_{\mathbb{Q}} \left(\ell[(\Phi(S_T) - V_T)^+] \right)$$

Without uncertainty

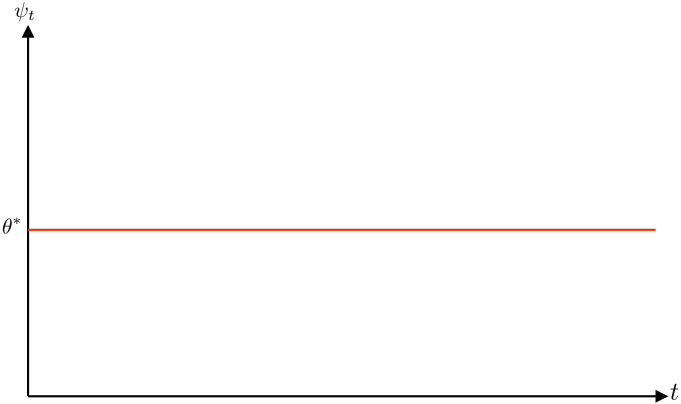

Robust

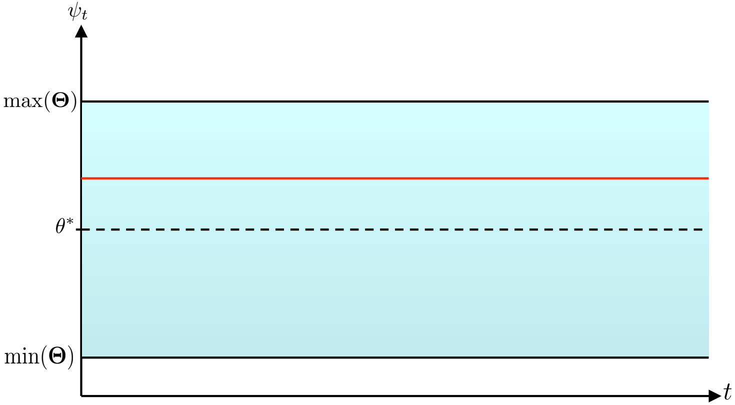

Adaptive robust control methodology

Strong robust

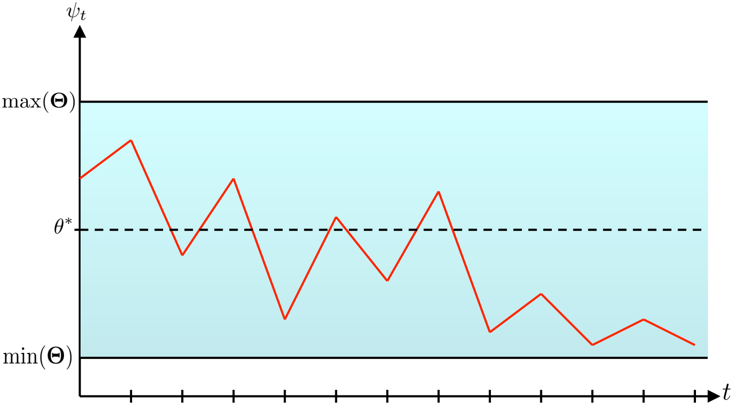

Adaptive robust

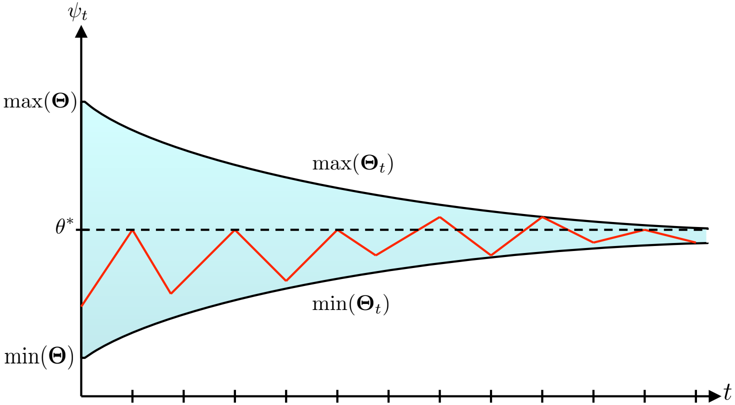

Proposition

The following inequalities hold

$$\inf_{\varphi\in\mathcal{A}} \mathbb{E}_{\theta^*}\left(\ell[(\Phi(S_T)-V_T)^+]\right) \le \inf_{\varphi\in\mathcal{A}} \sup_{\mathbb{Q}\in\mathcal{Q}_{h_0}^{\varphi,\Psi_A}} \mathbb{E}_{\mathbb{Q}}\left(\ell[(\Phi(S_T)-V_T)^+]\right)$$

$$\le \inf_{\varphi\in\mathcal{A}} \sup_{\mathbb{Q}\in\mathcal{Q}_{h_0}^{\varphi,\Psi_K}} \mathbb{E}_{\mathbb{Q}}\left(\ell[(\Phi(S_T)-V_T)^+]\right).$$

and

$$\inf_{\varphi\in\mathcal{A}} \sup_{\theta\in\Theta} \mathbb{E}_{\theta}\left(\ell[(\Phi(S_T)-V_T)^+]\right) \le \inf_{\varphi\in\mathcal{A}} \sup_{\mathbb{Q}\in\mathcal{Q}_{h_0}^{\varphi,\Psi_K}} \mathbb{E}_{\mathbb{Q}}\left(\ell[(\Phi(S_T)-V_T)^+]\right).$$

We conjecture that

$$\inf_{\varphi\in\mathcal{A}} \sup_{\mathbb{Q}\in\mathcal{Q}_{h_0}^{\varphi,\Psi_A}} \mathbb{E}_{\mathbb{Q}}\left(\ell[(\Phi(S_T)-V_T)^+]\right) \le \inf_{\varphi\in\mathcal{A}} \sup_{\theta\in\Theta} \mathbb{E}_{\theta}\left(\ell[(\Phi(S_T)-V_T)^+]\right)$$

Dynamic programming principle

Proposition

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

$$\inf_{\varphi\in\mathcal{A}} \sup_{\mathbb{Q}\in\mathcal{Q}^{\varphi,\Psi_A}_{h_0}} \mathbb{E}_{\mathbb{Q}}\left(\ell[(\Phi(S_T)-V_T)^+]\right)$$

coincides with the solution of the following robust Bellman equation :

$$W_T(x) = \ell\left[(\Phi(s)-v)^+\right], \quad x=(s,v,\widehat{\theta})\in E_X,$$
$$W_t(x) = \inf_{a\in A} \sup_{\theta\in\tau(t,\widehat{\theta})} \int_{E_X} W_{t+1}(y)Q(dy \mid t,x,a,\theta),$$

for any $x=(s,v,\widehat{\theta})\in E_X$ and $t=T-1,\dots,0$.

Note that the optimal strategy at time t is such that $\varphi^*_t(h_t) = \varphi^*_t(x_t)$.

Example : uncertain log-normal model

We consider that the stock price is driven by an uncertain log-normal model

$$S_{t+1} = Z_{t+1} S_t$$

where Z_t is an iid sequence such that $\ln Z_t \overset{\mathbb{P}_{\theta^*}}{\sim} N(\mu^*, (\sigma^*)^2)$.

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

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

with $\widehat{\mu}_1 = \ln Z_1 = \ln \frac{S_1}{S_0}$ and $\widehat{\sigma}^2_1 = 0$.

Example : uncertain log-normal model

Due to asymptotic normality of the MLE $\widehat{\theta}_t = (\widehat{\mu}_t, \widehat{\sigma}_t^2)$, we have

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

So that, if κ_α is the $(1-\alpha)$–quantile of the χ_2^2 distribution,

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

is an approximate α-confidence region of θ^*, i.e., $\mathbf{\Theta}_t$ is such that

$$\mathbb{P}_{\theta^*}(\theta^* \in \mathbf{\Theta}_t) \approx 1 - \alpha$$

[See Bielecki et al. (2016) for more details]

The adaptive robust control problem can be solved using the following dynamic programming principle :

$$W_T(x) = \ell\left[(\Phi(s) - v)^+\right], \quad x = (s, v, \widehat{\mu}, \widehat{\sigma}^2) \in E_X,$$

$$W_t(x) = \inf_{a \in A} \sup_{(\mu,\sigma^2) \in \tau(t,\widehat{\mu},\widehat{\sigma}^2)} \int_{E_X} W_{t+1}(y) Q(dy \mid t, x, a; \mu, \sigma^2)$$

where $x = (s, v, \widehat{\mu}, \widehat{\sigma}^2) \in E_X = \mathbb{R}_+ \times \mathbb{R} \times \mathbb{R} \times \mathbb{R}_+$, $t = T - 1, \ldots, 0$

Example : uncertain log-normal model

The integral in the previous slide can be written as

$$\int_{\mathbb{R}} W_{t+1}\left(se^{\mu+\sigma z}, v + as(e^{\mu+\sigma z} - 1), R(t, \widehat{\mu}, \widehat{\sigma}^2, \mu + \sigma z)\right)\phi(z)dz$$

where ϕ is the density of the standard normal distribution and R is such that

$$R\left(t, \widehat{\mu}, \widehat{\sigma}^2, y\right) = \left(\frac{t}{t+1}\widehat{\mu} + \frac{1}{t+1}y, \frac{t}{t+1}\widehat{\sigma}^2 + \frac{t}{(t+1)^2}(\widehat{\mu} - y)^2\right)$$

Perspectives

- Numerically solve Bellman equation for the considered hedging problem : challenging issue due to the curse of dimensionality (optimal quantization, approximate dynamic programming could be used)
- Compare hedging performance with other approaches : control without uncertainty, standard robust, adaptive robust, Bayesian adaptive robust

Thanks for your attention.

References I

Başar, T. and Bernhard, P. (1995).
H^∞-optimal control and related minimax design problems.
Systems & Control : Foundations & Applications. Birkhäuser Boston, Inc., Boston, MA, second edition.
A dynamic game approach.

Bayraktar, E., Cosso, A., and Pham, H. (2014).
Robust feedback switching control : dynamic programming and viscosity solutions.
http ://arxiv.org/pdf/1409.6233v1.pdf.

Bielecki, T., Chen, T., and Cialenco, I. (2016).
Recursive Construction of Confdence Regions.
Preprint.

Chen, H. F. and Guo, L. (1991).
Identification and stochastic adaptive control.
Systems & Control : Foundations & Applications. Birkhäuser Boston, Inc., Boston, MA.

Corsi, M., Pham, H., and Runggaldier, W. (2007).
Numerical approximation by quantization of control problems in finance under partial observations.
preprint.

Hansen, L. P., Sargent, T. J., Turmuhambetova, G., and Williams, N. (2006).
Robust control and model misspecification.
J. Econom. Theory, 128(1) :45–90.

References II

Hansen, P. L. and Sargent, T. J. (2008).
Robustness.
Princeton University Press.

Iyengar, G. N. (2005).
Robust dynamic programming.
Mathematics of Operations Research, 30(2) :pp. 257–280.

Kumar, P. R. and Varaiya, P. (1986).
Stochastic systems : estimation, identification and adaptive control.
Prentice-Hall, Inc.

Runggaldier, W., Trivellato, B., and Vargiolu, T. (2002).
A Bayesian adaptive control approach to risk management in a binomial model.
In *Seminar on Stochastic Analysis, Random Fields and Applications, III (Ascona, 1999)*, volume 52 of *Progr. Probab.*, pages 243–258.

Sirbu, M. (2014).
A note on the strong formulation of stochastic control problems with model uncertainty.
Electronic Communications in Probability, 19.