EE365: Risk Averse Control

Risk averse optimization

Exponential risk aversion

Risk averse control

Outline

Risk averse optimization

Exponential risk aversion

Risk averse control

Risk measures

- suppose f is a random variable we'd like to be small
  (*i.e.*, an objective or cost)
- $\mathbf{E}\, f$ gives average or mean value
- many ways to quantify **risk** (of a large value of f)
  - $\mathbf{Prob}(f \geq f^{\mathrm{bad}})$ (value-at-risk, VAR)
  - $\mathbf{E}(f - f^{\mathrm{bad}})_+$ (conditional value-at-risk, CVAR)
  - $\mathbf{var}\, f = \mathbf{E}(f - \mathbf{E}\, f)^2$ (variance)
  - $\mathbf{E}(f - \mathbf{E}\, f)_+^2$ (downside variance)
  - $\mathbf{E}\, \phi(f)$, where ϕ is increasing and convex
    (when large f is good: expected utility $\mathbf{E}\, U(f)$ with increasing concave *utility function* U)
- **risk aversion:** we want $\mathbf{E}\, f$ small **and** low risk

Risk averse optimization

- now suppose random cost $f(x, \omega)$ is a function of a decision variable x and a random variable ω
- different choices of x lead to different values of mean cost $\mathbf{E} f(x, \omega)$ and risk $R(f(x, \omega))$
- there is typically a trade-off between minimizing mean cost and risk
- standard approach: minimize $\mathbf{E} f(x, \omega) + \lambda R(f(x, \omega))$
  - $\mathbf{E} f(x, \omega) + \lambda R(f(x, \omega))$ is the **risk-adjusted mean cost**
  - $\lambda > 0$ is called the **risk aversion parameter**
  - varying λ over $(0, \infty)$ gives trade-off of mean cost and risk
- **mean-variance optimization**: choose x to minimize $\mathbf{E} f(x, \omega) + \lambda \,\mathbf{var}\, f(x, \omega)$

Example: Stochastic shortest path

- find path in directed graph from vertex A to vertex B
- edge weights are independent random variables with known distributions
- commit to path beforehand, with no knowledge of weight values
- path length L is random variable
- minimize $\mathbf{E}\, L + \lambda \,\mathbf{var}\, L$, with $\lambda \geq 0$
- for fixed λ, reduces to deterministic shortest path problem with edge weights $\mathbf{E}\, w_e + \lambda \,\mathbf{var}\, w_e$

Stochastic shortest path

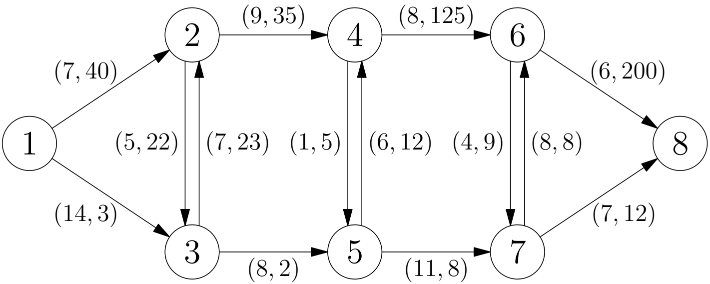

- find path from vertex $A = 1$ to vertex $B = 8$
- edge weights are lognormally distributed
- edges labeled with mean and variance: $(\mathbf{E}\, w_e, \mathbf{var}\, w_e)$

Stochastic shortest path

$\lambda = 0$: $\mathbf{E}\, L = 30$, $\mathbf{var}\, L = 400$

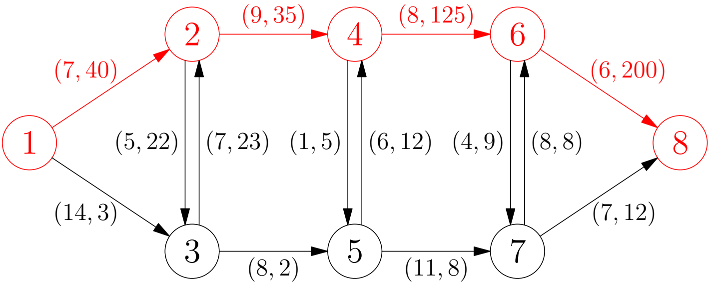

Stochastic shortest path

$\lambda = 0.05$: $\mathbf{E}\, L = 35$, $\mathbf{var}\, L = 100$

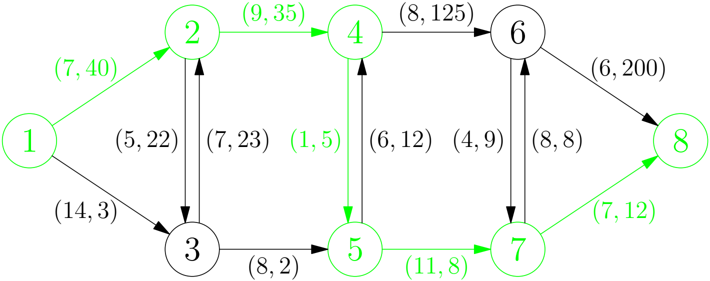

Stochastic shortest path

$\lambda = 10$: $\mathbf{E}\,L = 40$, $\mathbf{var}\,L = 25$

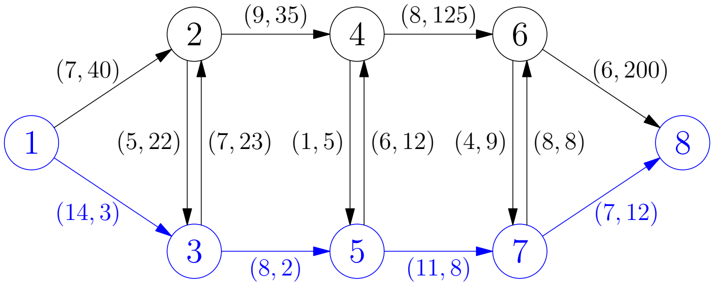

Stochastic shortest path

trade-off curve: $\lambda = 0$, $\lambda = 0.05$, $\lambda = 10$

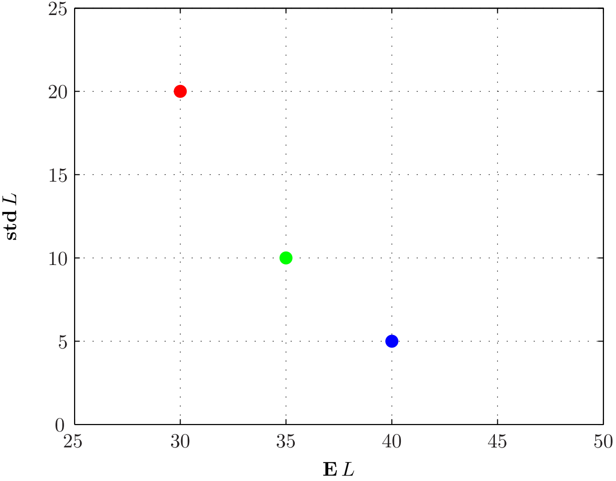

Stochastic shortest path

distribution of L: $\lambda = 0$, $\lambda = 0.05$, $\lambda = 10$

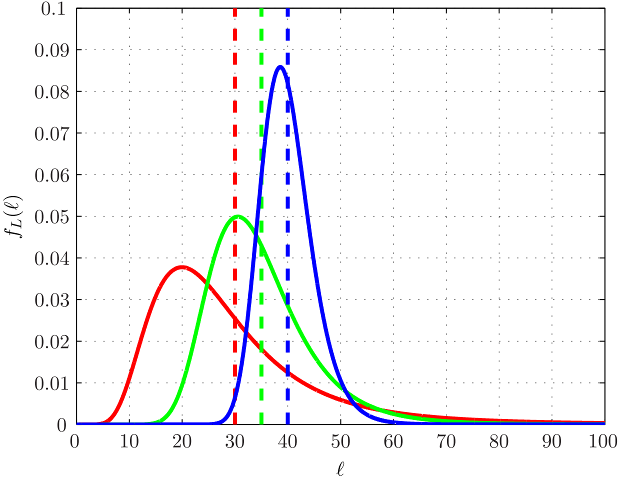

Example: Mean-variance (Markowitz) portfolio optimization

- choose portfolio $x \in \mathbf{R}^n$
  - x_i is amount of asset i held (short position when $x_i < 0$)
- (random) asset return $r \in \mathbf{R}^n$ has known mean $\mathbf{E}\, r = \mu$, covariance $\mathbf{E}\,(r-\mu)(r-\mu)^T = \Sigma$
- portfolio return is (random variable) $R = r^T x$
  - mean return is $\mathbf{E}\, R = \mu^T x$
  - return variance is $\mathbf{var}\, R = x^T \Sigma x$
- maximize $\mathbf{E}\, R - \gamma\, \mathbf{var}\, R = \mu^T x - \gamma x^T \Sigma x$, 'risk adjusted (mean) return'
- $\gamma > 0$ is risk aversion parameter

Example: Mean-variance (Markowitz) portfolio optimization

- can add constraints such as
  - $\mathbf{1}^T x = 1$ (budget constraint)
  - $x \geq 0$ (long positions only)
- can be solved as a (convex) quadratic program (QP)

$$\begin{array}{ll} \text{maximize} & \mu^T x - \gamma x^T \Sigma x \\ \text{subject to} & \mathbf{1}^T x = 1, \quad x \geq 0 \end{array}$$

  (or analytically without long-only constraint)
- varying γ gives trade-off of mean return and risk

Example: Mean-variance (Markowitz) portfolio optimization

numerical example: $n = 30$, $r \sim \mathcal{N}(\mu, \Sigma)$
trade-off curve: $\gamma = 10^{-2}$, $\gamma = 10^{-1}$, $\gamma = 1$

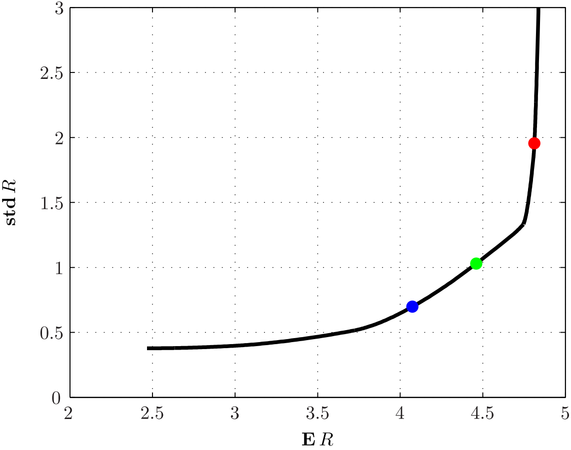

Example: Mean-variance (Markowitz) portfolio optimization

numerical example: $n = 30$, $r \sim \mathcal{N}(\mu, \Sigma)$
distribution of portfolio return: $\gamma = 10^{-2}$, $\gamma = 10^{-1}$, $\gamma = 1$

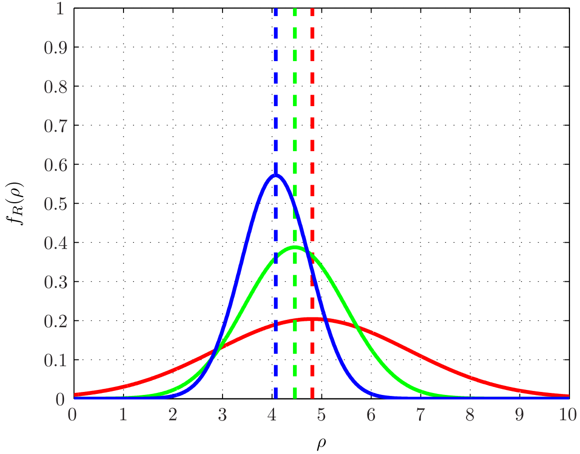

Outline

Exponential risk aversion

- suppose f is a random variable
- exponential risk measure, with parameter $\gamma > 0$, is given by

$$R_\gamma(f) = \frac{1}{\gamma} \log \left(\mathbf{E} \exp(\gamma f) \right)$$

  ($R_\gamma(f) = \infty$ if f is heavy-tailed)
- $\exp(\gamma f)$ term emphasizes large values of f
- $R_\gamma(f)$ is (up to a factor of γ) the cumulant generating function of f
- we have

$$R_\gamma(f) = \mathbf{E}\, f + (\gamma/2)\, \mathbf{var}\, f + o(\gamma)$$

- so minimizing exponential risk is (approximately) mean-variance optimization, with risk aversion parameter $\gamma/2$

Exponential risk expansion

- use $\exp u = 1 + u + u^2/2 + \cdots$ to write

$$\mathbf{E}\exp(\gamma f) = 1 + \gamma\,\mathbf{E} f + (\gamma^2/2)\,\mathbf{E} f^2 + \cdots$$

- use $\log(1+u) = u - u^2/2 + \cdots$ to write

$$\log \mathbf{E}\exp(\gamma f) = \gamma\,\mathbf{E} f + (\gamma^2/2)\,\mathbf{E} f^2 - (1/2)\left(\gamma\,\mathbf{E} f + (\gamma^2/2)\right)^2 + \cdots$$

- expand square, drop γ^3 and higher order terms to get

$$\log \mathbf{E}\exp(\gamma f) = \gamma\,\mathbf{E} f + (\gamma^2/2)\,\mathbf{E} f^2 - (\gamma^2/2)(\mathbf{E} f)^2 + \cdots$$

- divide by γ to get

$$R_\gamma(f) = \mathbf{E} f + (\gamma/2)\,\mathbf{var}\, f + o(\gamma)$$

Properties

- $R_\gamma(f) = \mathbf{E} f + (\gamma/2)\,\mathbf{var}\, f$ for f normal
- $R_\gamma(a + f) = a + R_\gamma(f)$ for deterministic a
- $R_\gamma(f)$ can be thought of as a variance adjusted mean, but in fact it's probably closer to what you really want (*e.g.*, it penalizes deviations above the mean more than deviations below)
- monotonicity: if $f \leq g$, then $R_\gamma(f) \leq R_\gamma(g)$
- can extend idea to conditional expectation:

$$R_\gamma(f \mid g) = \frac{1}{\gamma} \log \mathbf{E}(\exp(\gamma f) \mid g)$$

Value at risk bound

- exponential risk gives an upper bound on VaR (value at risk)
- indicator function of $f \geq f^{\rm bad}$ is $I^{\rm bad}(f) = \begin{cases} 0 & f < f^{\rm bad} \\ 1 & f \geq f^{\rm bad} \end{cases}$
- $\mathbf{E}\, I^{\rm bad}(f) = \mathbf{Prob}(f \geq f^{\rm bad})$
- for $\gamma > 0$, $\exp \gamma(f - f^{\rm bad}) \geq I^{\rm bad}(f)$ (for all f)
- so $\mathbf{E} \exp \gamma(f - f^{\rm bad}) \geq \mathbf{E}\, I^{\rm bad}(f)$
- hence
$$\mathbf{Prob}(f \geq f^{\rm bad}) \leq \exp \gamma(R_\gamma(f) - f^{\rm bad})$$

Outline

Risk averse optimization

Exponential risk aversion

Risk averse control

Risk averse stochastic control

- dynamics: $x_{t+1} = f_t(x_t, u_t, w_t)$, with $x_0, w_0, w_1, \ldots$ independent
- state feedback policy: $u_t = \phi_t(x_t)$, $t = 0, \ldots, T-1$
- risk averse objective:

$$J = \frac{1}{\gamma} \log \mathbf{E} \exp \gamma \left(\sum_{t=0}^{T-1} g_t(x_t, u_t) + g_T(x_T) \right)$$

  - g_t is stage cost; g_T is terminal cost
  - $\gamma > 0$ is risk aversion parameter
- **risk averse stochastic control problem:**
  find policy $\mu = (\mu_0, \ldots, \mu_{T-1})$ that minimizes J

Interpretation

- total cost is random variable

$$C = \sum_{t=0}^{T-1} g_t(x_t, u_t) + g_T(x_T)$$

- standard stochastic control minimizes $\mathbf{E}\, C$
- risk averse control minimizes $R_\gamma(C)$
- risk averse policy yields larger expected total cost than standard policy, but smaller risk

Risk averse value function

- we are to minimize

$$J = R_\gamma\left(\sum_{t=0}^{T-1} g_t(x_t, u_t) + g_T(x_T)\right)$$

  over policies $\mu = (\mu_0, \ldots, \mu_{T-1})$

- define value function

$$V_t(x) = \min_{\mu_t,\ldots,\mu_{T-1}} R_\gamma\left(\sum_{\tau=t}^{T-1} g_\tau(x_\tau, u_\tau) + g_T(x_T) \;\middle|\; x_t = x\right)$$

  - $V_T(x) = g_T(x)$
  - could minimize over *input* u_t, *policies* $\mu_{t+1}, \ldots, \mu_{T-1}$
  - same as usual value function, but replace $\mathbf{E}$ with R_γ

Risk averse dynamic programming

- optimal policy $\mu^\star$ is

$$\mu_t^\star(x) \in \operatorname*{argmin}_u \left(g_t(x,u) + R_\gamma\, V_{t+1}(f_t(x,u,w_t))\right)$$

  where expectation in R_γ is over w_t

- (backward) recursion for V_t:

$$V_t(x) = \min_u \left(g_t(x,u) + R_\gamma\, V_{t+1}(f_t(x,u,w_t))\right)$$

- same as usual DP, but replace $\mathbf{E}$ with R_γ (both over w_t)

Multiplicative version

- precompute $h_t(x, u) = \exp \gamma g_t(x, u)$
- instead of V_t, change variables $W_t(x) = \exp \gamma V_t(x)$
- DP recursion is

$$W_t(x) = \min_u \big(h_t(x, u) \, \mathbf{E} \; W_{t+1}(f_t(x, u, w_t)) \big)$$

- optimal policy is

$$\mu_t^\star(x) \in \underset{u}{\operatorname{argmin}} \big(h_t(x, u) \, \mathbf{E} \; W_{t+1}(f_t(x, u, w_t)) \big)$$

Example: Optimal purchase

- must buy an item in one of $T = 4$ time periods
- prices are IID with $p_t \in \{1, \ldots, 10\}$, $\mathbf{Prob}(p_t = p) \propto 0.95^p$
- in each time period, the price is revealed and you choose to buy or wait
  - once you've bought the item, your only option is to wait
  - in the last period, you must buy the item if you haven't already

Example: Optimal purchase

optimal policy: wait, buy

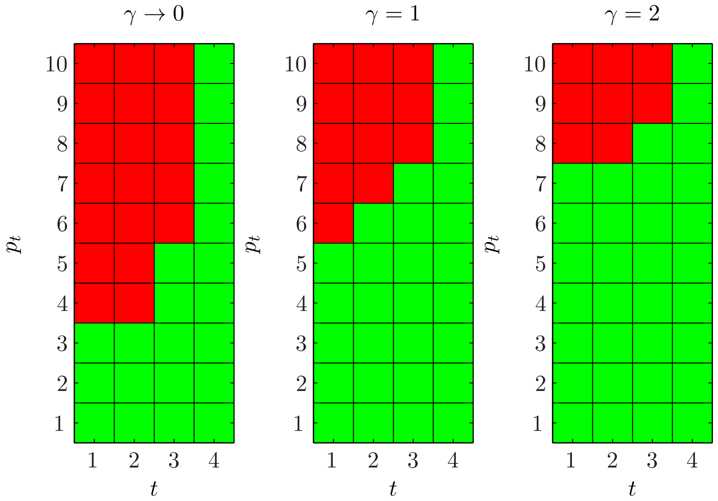

Example: Optimal purchase

purchase if price is below threshold: $\gamma \to 0$, $\gamma = 1$, $\gamma = 2$

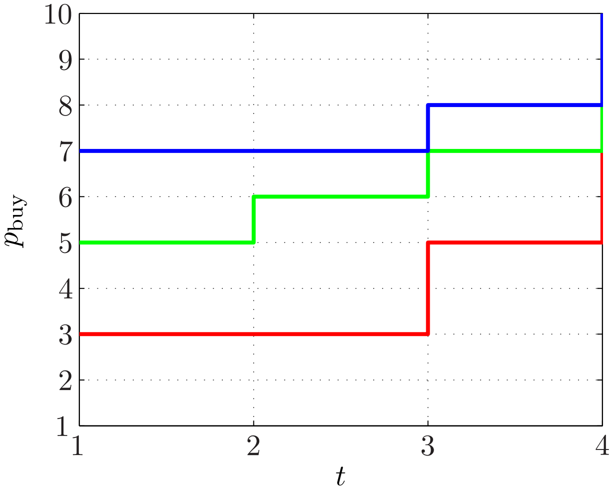

Example: Optimal purchase

trade-off curve: $\gamma \to 0$, $\gamma = 1$, $\gamma = 2$

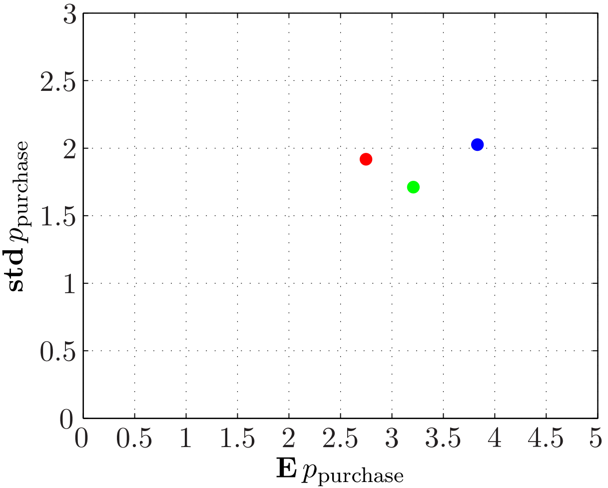

Example: Optimal purchase

distribution of purchase price: $\gamma \to 0$, $\gamma = 1$, $\gamma = 2$