Robust decision making with relative entropy

Andrew Lim
Department of Industrial Engineering
& Operations Research
U.C. Berkeley

Email: lim@ieor.berkeley.edu

July 2007.

Motivations

- Model uncertainty is unavoidable
  - erroneous statistical assumptions
  - calibration errors

  Moreover, real world performance of "optimal" decisions can be very bad if model error is ignored.

- Decision theory and economics: Ellsberg paradox suggests that decision makers are averse to model uncertainty.
  - Axiomatic justification for "worst case" models (Gilboa & Schmiedler (1989))
  - Explanations of the equity premium puzzle and home bias amongst investors, etc.

Worst case models based on "relative entropy"

(Nominal) probability space: $(\Omega, \mathcal{F}, \mathbb{P})$.

Decision variable: $\pi \in \mathcal{C}$ and payoff X_π.

Objective: $\max_{\pi \in \mathcal{C}} \mathbb{E}_{\mathbb{P}} X_\pi$

Model uncertainty $\Leftrightarrow$ we don't have a full specification of $\mathbb{P}$
$\Rightarrow$ family of possible models.

Describe family of alternative models via *relative entropy*

$$\mathcal{R}(\mathbb{Q} \,|\, \mathbb{P}) = \mathbb{E}_{\mathbb{Q}}[\,\ln \frac{d\mathbb{Q}}{d\mathbb{P}}]$$

i.e. $\mathcal{R}(\mathbb{Q} \,|\, \mathbb{P}) \equiv$ distance between models $\mathbb{P}$ and $\mathbb{Q}$.

Constrained formulation:

$$\max_{\pi\in\mathcal{C}} \min_{\mathbb{Q}} \mathbb{E}_{\mathbb{Q}} X_\pi \ \text{ subject to: } \mathcal{R}(\mathbb{Q}\,|\,\mathbb{P}) \le \gamma.$$

Set of alternative models

$$\mathcal{V} = \{\mathbb{Q}\,|\,\mathcal{R}(\mathbb{Q}\,|\,\mathbb{P}) \le \gamma\}$$

Penalty formulation:

$$\max_{\pi\in\mathcal{C}} \min_{\mathbb{Q}} \ \mathbb{E}_{\mathbb{Q}} X_\pi + \theta\mathcal{R}(\mathbb{Q}\,|\,\mathbb{P}) = \mathbb{E}_{\mathbb{Q}}\{X_\pi + \theta \ln \frac{d\mathbb{Q}}{d\mathbb{P}}\}$$

An interesting duality result

$$\underbrace{\max_{\pi\in\mathcal{C}} \min_{\mathbb{Q}} \ \mathbb{E}_{\mathbb{Q}}\{X_\pi + \theta \ln \frac{d\mathbb{Q}}{d\mathbb{P}}\}}_{\text{Robust problem}} = \underbrace{\max_{\pi\in\mathcal{C}} -\theta \ln \mathbb{E}_{\mathbb{P}} e^{-\frac{1}{\theta}X_\pi}}_{\text{“risk sensitive” problem}}$$

Duality between free energy and relative entropy from large deviations
e.g. Pra, Meneghini and Runggaldier (1996), Peterson, James & Dupuis (2000), etc.

Nominal model $\mathbb{P}$ is some multivariate distribution, and the standard "entropy approach" puts a "cloud" of models around the nominal and considers worst case optimization over this family.

Some interesting questions:

- how do we account for different ambiguity levels using RE?

- what are the implications for decision making

- what happens to the exponential utility representation?

- what are the implications for risk measures?

Overview

1. Intensity control problem: pricing perishable products.
   - Modelling different ambiguity levels
   - Robust control problem
   - Implications for solution
   - “Duality” & revenue sharing

2. Generalizations

3. Implications for risk measures

Multi-product pricing

Multi-product pricing

Typical setup:

- m products.
- Price $p = [p_1, p_2, \cdots, \cdots, p_m]'$
  $\Rightarrow$ instantaneous arrival rates $\lambda_1(p), \lambda_2(p), \cdots, \lambda_m(p)$.
- Pool of resources $x = [x_1, x_2, \cdots, x_n]'$.
- 1 unit of item i consumes A_{ji} of resource $j \Rightarrow A_i = [A_{1i}, \cdots, A_{ni}]'$.
- When there is no model ambiguity

$$\max_{p(t)} \mathbb{E} \int_0^T p(t)' dN(t) \text{ s.t. } \int_0^T A dN(t) \le c$$

Gallego and van Ryzin (1997).

Given $p(t, x)$, nominal model $\mathbb{P} \equiv$ intensity

$$\lambda(t) \equiv \lambda(p(t)) \triangleq [\lambda_1(p(t)), \cdots, \lambda_m(p(t))].$$

Likelihood ratio:

$$\frac{d\mathbb{Q}}{d\mathbb{P}} = \exp\left\{ \sum_{i=1}^{m} \left(\int_0^T \ln \kappa_i(s^-) dN_i(s) + \int_0^T (1 - \kappa_i(s))\lambda_i(s) ds \right) \right\}.$$

Non-negative predictable process $\kappa(t) \triangleq [\kappa_1(t), \cdots, \kappa_m(t)]$.

Under $\mathbb{Q}$, intensity $\beta_i(t) = \lambda_i(t)\kappa_i(t)$.

Relative Entropy

$$\begin{aligned}\mathcal{R}(\mathbb{Q}\,|\,\mathbb{P}) &= \mathbb{E}_{\mathbb{Q}}[\ln \frac{d\mathbb{Q}}{d\mathbb{P}}] \\ &= \mathbb{E}_{\mathbb{Q}} \int_0^T \sum_{i=1}^{m} \lambda_i(p(t))[\kappa_i(t) \ln \kappa_i(t) + 1 - \kappa_i(t)]dt\end{aligned}$$

$\equiv$ measure of deviation of $[\kappa_1(t), \cdots, \kappa_m(t)]$ from 1
$\equiv$ measure of deviation of $\mathbb{Q}$ from $\mathbb{P}$.

Observe:

- $\mathcal{R}(\mathbb{Q}\,|\,\mathbb{P}) \geq 0$
- $\mathcal{R}(\mathbb{Q}\,|\,\mathbb{P}) = 0$ iff $\mathbb{Q} = \mathbb{P}$ (equiv. $\kappa(t) = 1$)
- Convex in $\mathbb{Q}$ (equiv. $\kappa(t)$).

<u>“Classical” robust model</u>

“Robust problem”:

$$\begin{cases} \max_{p(t)} \min_{\kappa(t)} \mathbb{E}_{\mathbb{Q}} \Big[\sum_{i=1}^m \int_0^T p_i(t) dN_i(t) \\ \qquad\qquad + \theta \sum_{i=1}^m \int_0^T \lambda_i(p) \{\kappa_i(t) \ln \kappa_i(t) + 1 - \kappa_i(t)\} dt \Big] \\ \text{subject to:} \\ \sum_{i=1}^m \int_0^T A_i dN_i(t) \le c \end{cases}$$

This is equivalent to the “risk sensitive” problem:

$$\begin{cases} \max_{p(t)} -\theta \ln \mathbb{E}_{\mathbb{P}}\, e^{-\frac{1}{\theta} \int_0^T \sum_{i=1}^m p_i(t) dN_i(t)} \\ \text{subject to:} \\ \sum_{i=1}^m \int_0^T A_i dN_i(t) \le c \end{cases}$$

Robust pricing II: ambiguity levels

$$\max_{p(t)} \min_{\kappa(t)} \mathbb{E}_{\mathbb{Q}} \Bigg[\sum_{i=1}^{m} \int_0^T p_i(t) dN_i(t) + \sum_{i=1}^{m} \theta_i \underbrace{\int_0^T \lambda_i(p) \{\kappa_i(t) \ln \kappa_i(t) + 1 - \kappa_i(t)\} \, dt}_{\text{uncertainty for income stream } i} \Bigg]$$

subject to:

$$\sum_{i=1}^{m} \int_0^T A_i dN_i(t) \le c$$

θ_i "small" $\Rightarrow$ "high" level of ambiguity for item i.

Optimality equations:

$$0 = J_t(t, x) + \max_p \sum_{i=1}^{m} \lambda_i(p)\theta_i\{1 - e^{-\frac{1}{\theta_i}[p_i - (J(t,x) - J(t,x-A_i))]}\}$$

Intuition:

Sale of 1 unit of item i gives net profit of

$$p_i - [J(t, x) - J(t, x - A_i)].$$

If there is no ambiguity $\Rightarrow \theta_i \uparrow \infty$

$$0 = J_t(t, x) + \max_p \sum_{i=1}^{m} \lambda_i(p)[p_i - (J(t, x) - J(t, x - A_i))]$$

Max. instantaneous profit rate.

Plot of $U(x) = \theta[1 - \exp(-x/\theta)]$

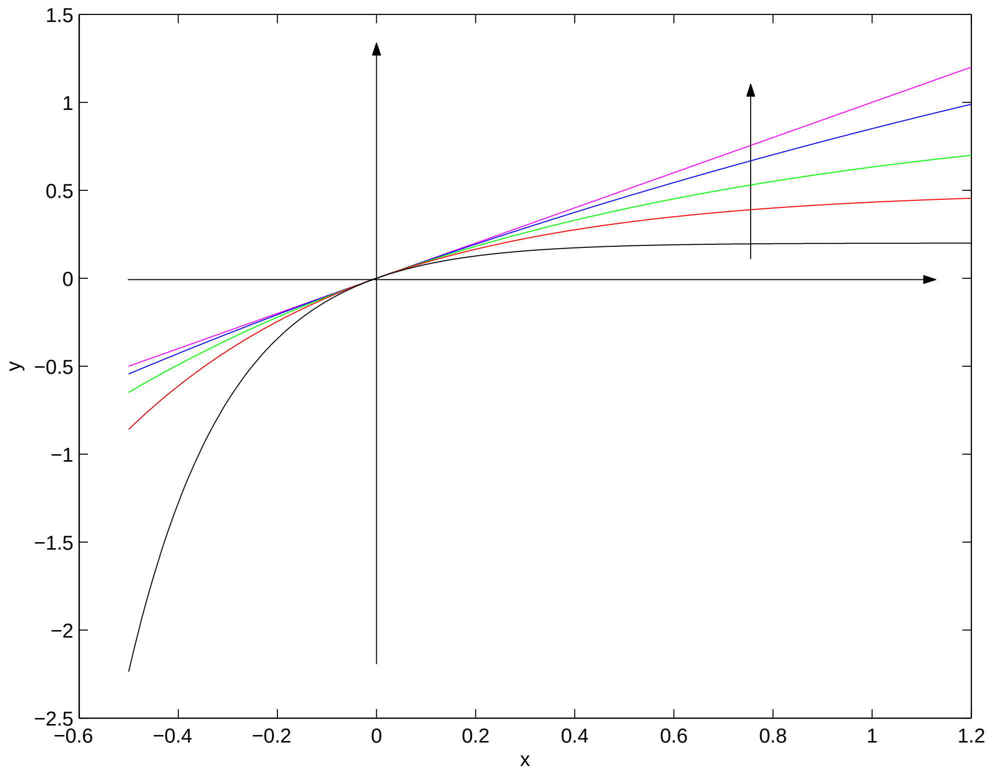

when $\theta \uparrow \infty$ (i.e. ambiguity decreases), $\theta[1 - \exp(-x/\theta)] \uparrow x$

$$0 = J_t(t, x) + \max_p \sum_{i=1}^{m} \lambda_i(p)\theta_i\{1 - e^{-\frac{1}{\theta_i}[p_i-(J(t,x)-J(t,x-A_i))]}\}$$

Ambiguity adjustment

$$U_i(p_i - [J(t, x) - J(t, x - A_i)]) = \theta_i\{1 - e^{-\frac{1}{\theta_i}[p_i-(J(t,x)-J(t,x-A_i))]}\}$$

Larger ambiguity (smaller θ_i) $\Rightarrow$ less value in profit from a sale (i.e. there is a "price for ambiguity")

Opt. equations $\equiv$ max. instantaneous *ambiguity adjusted* profit

"A dollar isn't equal to a dollar which isn't equal to a dollar"

Exponential utility representation ("Duality")

If $\theta = \theta_1 = \theta_2 = \cdots = \theta_m$ then dual problem is

$$\max_p -\theta \ln \mathbb{E} e^{-\frac{1}{\theta}\sum_{i=1}^m p_i(t)dN_i(t)}$$

If $\theta_1, \cdots, \theta_m$ are possibly different *but* all products are <u>independent</u> then dual problem is

$$\max_p \sum_{i=1}^m \underbrace{-\theta_i \ln \mathbb{E}_{\mathbb{P}}\, e^{-\frac{1}{\theta_i}\int_0^T p_i(t)dN_i(t)}}_{\text{CE for income stream } i}$$

If θ_i are different and products are not independent then neither is true.

Another multi-product pricing problem

Suppose that $p(t)$ is fixed.

The sale of one unit of product i at time t brings $p_i(t^-)$ revenue and uses A_i resources from the common pool.

The use of common resources affects/reduces the income that can be earned from other sources.
$\Rightarrow$ introduce "revenue sharing rule" $R_{ij}(t, x)$ to compensate for these losses.

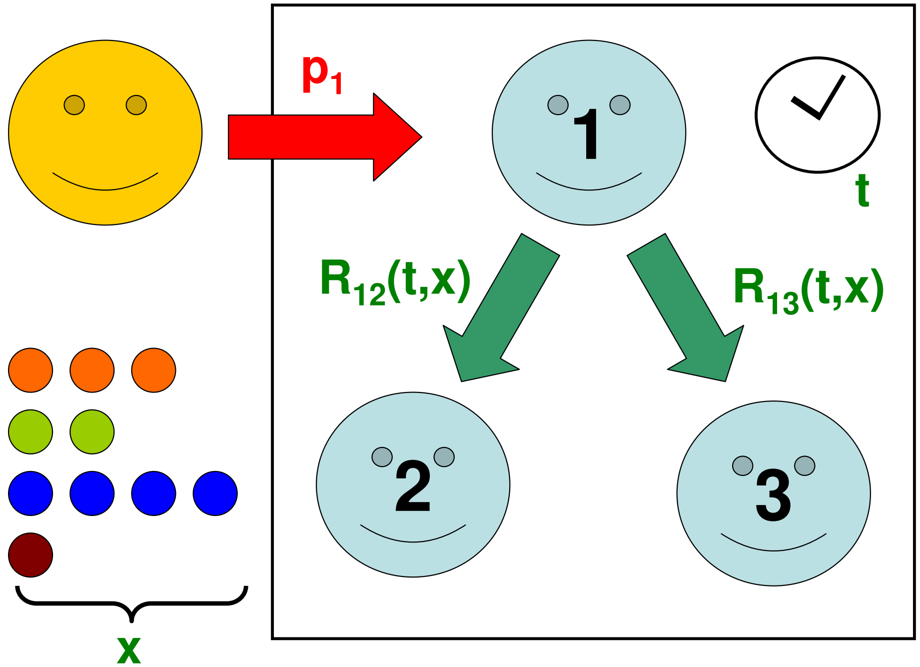

Income stream i now becomes:

$$\int_0^T \left[p_i(t) - \sum_{j \neq i} R_{ij}(t,\, x(t)) \right] dN_i(t) + \sum_{j \neq i} \int_0^T R_{ji}(t,\, x(t)) dN_j(t)$$

Multi-product pricing with revenue sharing

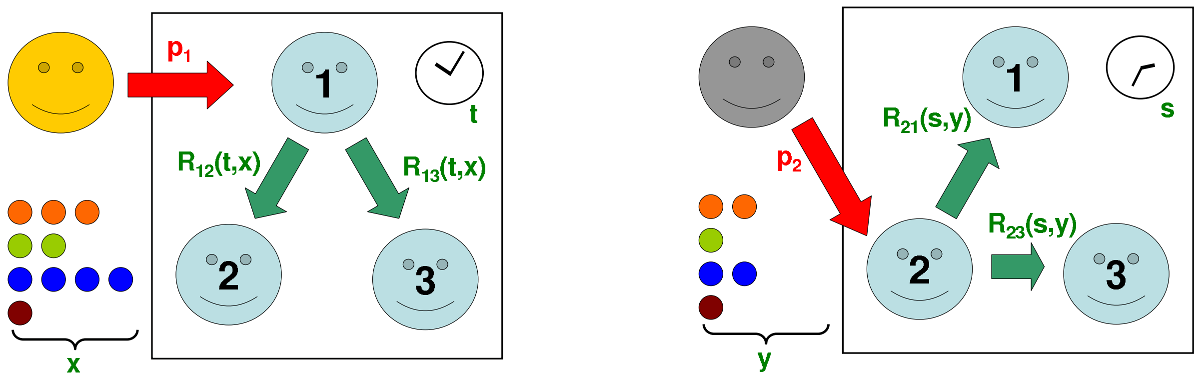

Pricing problem with revenue sharing:

$$\begin{cases} \max_p \sum_{i=1}^m -\theta_i \ln \mathbb{E}_{\mathbb{P}} \exp\left\{-\frac{1}{\theta_i}\left\{\int_0^T \left[p_i(t) - \sum_{j\neq i} R_{ij}(t,\, X(t^-))\right] dN_i \right.\right. \\ \qquad\qquad \left.\left. + \sum_{j\neq i} R_{ji}(t,\, X(t^-)) dN_j(t)\right\}\right\} \\ \text{Subject to:} \\ \int_0^T A\, dN(t) \le x \end{cases}$$

How should we set $R_{ij}(t,\, x)$?

Statement of main duality result

If we set

$$R_{i,j}(t, x; p) = \underbrace{W^j(t, x; p) - W^j(t, x - A_i; p)}_{\text{reduction in CE for stream } j \text{ from sale of product } i}$$

then "robust" and "revenue sharing" problems are equivalent.

Certainty equivalent for income stream j:

$$\begin{aligned} &W^i(t, x; p) \\ &\triangleq -\theta_i \ln \mathbb{E}^{\mathbb{P}}\left[\exp\left[-\frac{1}{\theta_i}\left\{\int_0^T \left[p_i(t) - \sum_{j\neq i} R_{ij}(t, x(t); p)\right] dN_i(t)\right.\right.\right. \\ &\qquad \left.\left.\left. + \sum_{j\neq i}\int_0^T R_{ji}(t, x(t); p) dN_j(t)\right\}\right]\right] \end{aligned}$$

As defined, $R_{ij}(t,\, x;\, p)$ and $W^i(t,\, x;\, p)$ are defined implicitly. However, it is possible to compute them:

$$
\begin{cases}
-\overline{W}_t(t,x;p) = \sum_{i=1}^{n} \lambda_i(p)\theta_i \left[1 - e^{-\frac{1}{\theta_i}\{p_i + \overline{W}(t,x-A_i;p) - \overline{W}(t,x;p)\}}\right] \\
\overline{W}(T,x;p) = 0, \quad \forall x \in \mathcal{C} \\
\overline{W}(t,x;p) = 0, \quad \forall t \in [0,T], x \in \{x \in \mathcal{C} \mid \forall i \ x \not\geq A_i\}.
\end{cases}
$$

and

$$
\begin{cases}
-W^i_t(t,x;p) = \lambda_i(p)\theta_i \left[1 - e^{-\frac{1}{\theta_i}\{p_i + \overline{W}(t,x-A_i;p) - \overline{W}(t,x;p)\}}\right] \\
W^i(T,x;p) = 0, \quad \forall x \in \mathcal{C} \\
W^i(t,x;p) = 0, \quad \forall t \in [0,T], x \in \{x \in \mathcal{C} \mid \forall j \ x \not\geq A_j\}.
\end{cases}
$$

We now have

$$R_{ij}(t,\, x;\, p) = W^j(t,\, x;\, p) - W^j(t,\, x - A_i;\, p).$$

Note: *revenue sharing rule depends on the entire pricing policy* $\Rightarrow$ "revenue sharing problem" has non-standard dynamics.

Recall the "revenue sharing" problem

$$\begin{cases} \max_p \sum_{i=1}^{m} -\theta_i \ln \mathbb{E}_{\mathbb{P}} \exp\left\{-\frac{1}{\theta_i}\left\{\int_0^T \left[p_i(t) - \sum_{j\neq i} R_{ij}(t,\, X(t^-))\right] dN_i \right.\right. \\ \qquad\qquad\qquad\qquad \left.\left. + \sum_{j\neq i} R_{ji}(t,\, X(t^-)) dN_j(t)\right\}\right\} \\ \text{Subject to:} \\ \int_0^T A\, dN(t) \leq x \end{cases}$$

Fix $p(t,\, x)$ and let $R_{ij}(t,\, x)$ be arbitrary.

Complication: Dynamic programming does not apply since the objective is non-separable.

<u>Step 1</u>: Given fixed $p(t, x)$ and arbitrary $R_{ij}(t, x)$, let

$$
\begin{aligned}
J(t, x; p) &= \min_{\kappa \geq 0} \sum_{i=1}^{n} \mathbb{E}_{\mathbb{Q}} \left[\int_0^T p_i(s) \lambda_i(p(s)) \kappa_i(s) ds \right] \\
&\quad + \sum_{i=1}^{n} \theta_i \mathbb{E}_{\mathbb{Q}} \left[\int_0^T \lambda_i(s)(\kappa_i(s) \ln \kappa_i(s) + 1 - \kappa_i(s)) ds \right] \\
&\equiv \text{“robust objective” for } p(t, x).
\end{aligned}
$$

and

$$
\begin{aligned}
& W^i(t, x; p) \\
&\triangleq -\theta_i \ln \mathbb{E}^{\mathbb{P}} \left[\exp \left[-\frac{1}{\theta_i} \left\{ \int_0^T \left[p_i(t) - \sum_{j \neq i} R_{ij}(t, x(t); p) \right] dN_i(t) \right. \right. \right. \\
&\qquad \left. \left. \left. + \sum_{j \neq i} \int_0^T R_{ji}(t, x(t); p) dN_j(t) \right\} \right] \right] \\
&\equiv \text{CE with revenue sharing}
\end{aligned}
$$

Calculate recursive equations for $J(t, x; p)$ and $W^i(t, x; p)$.

Step 2: Observe that if

$$R_{ij}(t, x; p) = W^j(t, x; p) - W^j(t, x - A_i; p)$$

then

$$W(t, x; p) = W^1(t, x; p) + \cdots + W^n(t, x; p) = J(t, x; p)$$

i.e. $J(t, x; p)$ and $W(t, x; p)$ coincide for each $p(t, x)$
$\Rightarrow$ they have the same opt. solution and optimality equations.

Example

K units of a single resource, m customer types, customer type i arrives with intensity $\lambda_i(p_i) = e^{-Bp_i}$ when price $p = [p_1, \cdots, p_m]$, and customer i pays p_i for one unit of the product.

Optimal pricing problem:

$$\begin{cases} \max_{p(t)} \min_{\kappa(t)} \mathbb{E}_{\mathbb{Q}} \Big[\sum_{i=1}^{m} \int_0^T p_i(t) dN_i(t) \\ \qquad\qquad + \theta \sum_{i=1}^{m} \int_0^T \lambda_i(p) \{\kappa_i(t) \ln \kappa_i(t) + 1 - \kappa_i(t)\} dt \Big] \\ \text{subject to:} \\ N_1(T) + \cdots + N_m(T) \le K \end{cases}$$

Value function

$$V(t,n) = \frac{1}{B} \ln \left\{ \sum_{j=0}^{n} \frac{[g(\theta)(T-t)]^j}{j!} \right\}$$

where

$$g(\theta) = \sum_{i=1}^{m} \left[\frac{B\theta_i}{1+B\theta_i} \right]^{1+B\theta_i}$$

Optimal pricing policy

$$p_i(t,n) = V(t,n) - V(t,n-1) - \theta_i \ln \left[\frac{1+B\theta_i}{B\theta_i} \right]$$

At optimality, revenue sharing rule is

$$R_{ij}(t,n) = \left(\frac{B\theta_j}{1+B\theta_j} \right)^{1+B\theta_j} \frac{1}{g(\theta)} \frac{1}{B} \ln \left\{ \frac{\sum_{k=0}^{n} \frac{[g(\theta)(T-t)]^k}{k!}}{\sum_{k=0}^{n-1} \frac{[g(\theta)(T-t)]^k}{k!}} \right\}$$

If item j is very ambiguous (θ_j is small) and l is quite unambiguous (θ_l large) then we price j higher than l, and the sale of product i compensates l more than j.

Summary

"a dollar isn't equal to a dollar which isn't equal to a dollar"

Implications for risk measures

Multi-dimensional point process $N_1(t), \cdots, N_m(t)$, filtration $\{\mathcal{F}_t\}$ $(0 \leq t \leq T)$, and $X \in \mathcal{F}_T$ the financial position at T.

$$\rho(X) \triangleq -\min_{\mathbb{Q}} \mathbb{E}_{\mathbb{Q}}\{X + \sum_{i=1}^{m} \theta_i \int_0^T \lambda_i(t)[\kappa_i(t) \ln \kappa_i(t) + 1 - \kappa_i(t)]dt\}$$

is a convex risk measure. (Foellmer & Schied (2002)).

Case 1: $N_1(t), \cdots, N_m(t)$ are independent and

$$X = \sum_{i=1}^{m} \int_0^T p_i(t) dN_i(t)$$

In this case:

$$\rho(X) = \sum_{i=1}^{m} \theta_i \ln \mathbb{E}_{\mathbb{P}} e^{-\frac{1}{\theta_i} \int_0^T p_i(t) dN_i(t)}$$

Case 2: $N_1(t), \cdots, N_m(t)$ dependent with intensity $\lambda_i(x(t))$, where

$$dx(t) = \sum_{j=1}^{m} \sigma_j(x(t))dN_j(t), \quad X = \sum_{i=1}^{m} \int_0^T p_i(t)dN_i(t)$$

$$\begin{aligned} \rho(X) &= \sum_{i=1}^{m} \theta_i \ln \mathbb{E}_{\mathbb{P}}\, e^{-\frac{1}{\theta_i}\left\{\int_0^T p_i(t)dN_i(t)+Y_i(T)\right\}} \\ Y_i(T) &= -\int_0^T \sum_{j\neq i} R_{ij}(t, x(t))dN_i(t) + \sum_{j\neq i}\int_0^T R_{ji}(t, x(t))dN_j(t) \end{aligned}$$

where

$$W^i(t, x(t)) \triangleq \theta_i \ln \mathbb{E}^{\mathbb{P}}_{x(t)} \exp\left[-\frac{1}{\theta_i}\left\{\int_t^T p_i(t)dN_i(t) \right.\right.$$
$$\left.\left. -\int_0^T \sum_{j\neq i} R_{ij}(t, x(t))dN_i(t) + \sum_{j\neq i}\int_0^T R_{ji}(t, x(t))dN_j(t)\right\}\right]$$

$$R_{ij}(t, x) = W^j(t, x+\sigma_i(x)) - W^j(t, x)$$

Capital and risk transfers through contracts

A standard application: Insurance company with m business units. Suppose each unit has end-of-year position

$$C_i = \int_0^T p_i(X(t))dN_i(t)$$

where

$$dX(t) = \alpha(X(t))dt + \sum_{i=1}^{m} \sigma_i(X(t))dN_i(t)$$

and intensity of $N_i(t)$, $\lambda_i(X(t))$, depends on $X(t)$.

Classical approach: the cash reserve is $\rho_\theta(C_1 + \cdots + C_m)$

Difficulty: Assumes full capital mobility between business units. If this does not hold, it suggests we should really be considering each business unit separately $\sum_{i=1}^{m} \rho_\theta(C_i)$. But this ignores possible diversification effects.

Another approach: Filipovic & Kupper

Basic idea: create capital mobility by writing contracts.

$$\min_{y_i^j} \sum_i \rho_\theta(C_i + \sum_{j=1}^m y_i^j Z^j) \text{ s.t. } \sum_{i=1}^m \sum_{j=1}^m y_i^j Z^j \le 0, \ \mathbb{P} - a.s.$$

e.g. For $Z^j = [C_j - K_j]^+$, if $y_1^1 = -1, y_1^2 = 0.1, \cdots, y_1^m = 0.1$, this means that unit 1 transfers all excess capital above K_1 to the common pool, and takes 10% of what other units put into the pool for itself.

In general, the resulting reserve is less than the "full mobility" case.

Our results show that

$$
\begin{aligned}
&\rho(X_1 + \cdots + X_m) \\
&\quad = -\theta \min_{\mathbb{Q}} \{\mathbb{E}_{\mathbb{Q}}(X_1 + \cdots + X_m) + \theta \mathbb{E}_{\mathbb{Q}} \ln \frac{d\mathbb{Q}}{d\mathbb{P}}\} \\
&\quad = \sum_{i=1}^{m} \rho_\theta(X_i + Y_i(T))
\end{aligned}
$$

i.e. We can achieve the reserve “full capital mobility” reserve if we write the appropriate contracts.