Existence of Risk-aware equilibrium solutions (distribution - risk-type tradeoffs)

NMEK615 - Stochastic programming and approximations

Lukáš Račko

20.11.2025

Contents

Games with random payoffs

Definition

Let $(\Omega, \mathcal{A}, \mathbb{P})$ *be a probability space. Let* I *be the set of players.* $\forall i \in I : X_i$ *be a set of strategies of the player* i *and* $u_i : X \times \Omega \to \mathbb{R}$ *be such that* $u_i(\boldsymbol{x})$ *is* $\mathcal{A}$*-measurable for every* $i \in I$ *and* $\boldsymbol{x} \in X$*. We say that* $G = (I, \{X_i\}_{i \in I}, \{u_i\}_{i \in I})$ *is a game with random payoff. For a given* $\omega \in \Omega$ *we will write* $G(\omega)$ *and mean* $(I, \{X_i\}_{i \in I}, \{u_i(\omega)\}_{i \in I})$ *the realization of game* G*.*

Solution concepts

How to characterize 'equilibrium' solutions in this model?

1. Make more solutions 'equilibrial' using probability constraints - $\boldsymbol{\alpha}$-Nash equilibria.
2. Introduce the risk aversion into the payoff (and hope for the best) - $\mathcal{R}$-equilibria.

Variational inequalities

For a standard model where we would like to maximize the expected payoff the equilibrium exists if and only if the following holds

$$\mathsf{E}\ u_i(x_i, \boldsymbol{x}_{-i}) \geq \mathsf{E}\ u_i(y, \boldsymbol{x}_{-i}), \forall y \in X_i, \forall i \in I. \tag{1}$$

Theorem (Nash, 1955)

If the set of players I is finite, $X_i = \mathcal{D}(P_i)$ and P_i is a finite set of pure strategies for each player$i \in I$ then there exists a solution $\boldsymbol{x}^* \in X$ that satisfies (1).

So in the 'risk-neutral' setting (or better in Von Neumann-Morgenstern utility setting) there is no restriction on the distribution of the payoff u_i and the only restriction in terms of strategies is that players must be able to diversify/randomize/mix their actions for the equilibrium to exist.

General variational inequalities

What happens when we replace E with general (player-dependent) risk measure $\mathcal{R}_i$?

$$-\mathcal{R}_i(-u_i(x_i, \boldsymbol{x}_{-i})) \geq -\mathcal{R}_i(-u_i(y, \boldsymbol{x}_{-i})), \forall y \in X_i, \forall i \in I. \tag{2}$$

In general, there may be no equilibrium, even when all the input data are 'nice'.

Contradiction

Let $\mathcal{R}_i = \mathrm{CV@R}_{1-\alpha}$ and $I = \{0, 1\}, P_i = \{0, 1\}$ and $(u_i(\boldsymbol{p}))_{\boldsymbol{p} \in P_0 \times P_1} \sim N_4((\mu_i(\boldsymbol{p}))_{\boldsymbol{p} \in P_0 \times P_1}, \Sigma_i)$. Then we receive a variational inequality

$$-\mathrm{CV@R}_{1-\alpha}(-u_0(x_0^*, x_1^*)) \geq -\mathrm{CV@R}_{1-\alpha}(-u_0(y, x_1^*)), \forall y \in X_0 \tag{3}$$

which can be reduced to

$$\begin{align} -\mathrm{CV@R}_{1-\alpha}(-u_0(x_0^*, x_1^*)) \geq -\mathrm{CV@R}_{1-\alpha}(-u_i(y, x_1^*)), \forall y \in X_0, \tag{4}\\ \mu_0(x_0^*, x_1^*) - \sigma_0(x_0^*, x_1^*)\,\mathrm{CV@R}_{1-\alpha}(Z) \geq \tag{5}\\ \geq \mu_0(y, x_1^*) - \sigma_0(y, x_1^*)\,\mathrm{CV@R}_{1-\alpha}(Z), \forall y \in X_0, \tag{6}\\ \mu_0(x_1^*, x_2^*) - \mu_0(y, x_2^*) \geq (\sigma_0(x_0^*, x_1^*) - \sigma_0(y, x_1^*))C(\alpha), \tag{7}\end{align}$$

where $Z \sim N(0, 1)$ and $C(\alpha) = \mathrm{CV@R}_{1-\alpha}(Z) \to \infty, \alpha \to 0$. And so the equilibrium may not exists for any α but at least we have a nice "closed form" characterization!

Sensible risk measures

We want to restrict ourselves into a sensible class of risk measures. Because each risk measure can be viewed as defining a linear ordering over a class of lotteries $L(\Sigma)$ defined as $X \succeq_{\mathcal{R}} Y$ if and only if $-\mathcal{R}(-X) \geq -\mathcal{R}(-Y)$, $X, Y \in L(\Sigma)$. We want to have at least some connection of $\mathcal{R}$ with the Von Neumann-Morgenstern utility theory therefore we want this ordering $\succeq_{\mathcal{R}}$ to be at least consistent with $\succeq_{FSD}$ because then if every Von Neumann-Morgenstern utility agent agrees that X is better than Y our agent should agree as well.

Distortion risk measures

Definition

Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space. Non-decreasing function $\psi : [0,1] \to [0,1]$ that satisfies $\psi(0) = 0$, $\psi(1) = 1$ is called a distortion function. By the dual distortion function corresponding to ψ we mean $\psi^*(x) := 1 - \psi(1-x)$, $x \in [0,1]$. The monotone, normalized set function defined on $(\Omega, \mathcal{A})$ as $\Psi(.) = (\psi \circ \mathbb{P})(.) = \psi(\mathbb{P}(.))$ is called the distorted version of $\mathbb{P}$ or ψ-distortion of $\mathbb{P}$.

Definition

Let $L(\Omega, \mathcal{A}, \mathbb{P})$ be the space of integrable random variables. Functional $\rho : L(\Omega, \mathcal{A}, \mathbb{P}) \to \mathbb{R}$ is called the distortion risk measure if $\rho(X) = \oint X d(\psi \circ \mathbb{P})$ for some distortion function ψ.

Distortion risk measures

Theorem (Basic properties of distortion risk measures ?)

Let $(\Omega, \mathcal{A}, \mathbb{P})$ *be a probability space and* Ψ *be a distorted version of* $\mathbb{P}$ *then*

1. *If* $X \succeq_{FSD} Y$ *then* $\oint X d\Psi \geq \oint Y d\Psi$
2. $-\oint X d\Psi = \oint -X d\Psi^*$
3. $\oint X d\Psi = \max_{Q \in \mathcal{Q}} E_Q X \iff \psi$ *is a concave function.*

Whatsmore, the distortion risk measures are exactly the Kahneman-Tversky's prospects!

The distorted expectation form

For non-negative random variables we receive just

$$\oint X d(\psi \circ \mathbb{P}) = \int_{(0,\infty)} \psi(\mathbb{P}(X > t))dt = \int_{(0,\infty)} \psi(S_X(t))dt, \tag{8}$$

where $S_X(t)$ is the survival function of X. Notice the similarity to the formula for the expectation of the non-negative random variable X as

$$\mathsf{E}\, X = \int_{(0,\infty)} S_X(t)dt$$

The integrated quantile function form

Another equivalent reformulation is the integrated quantile function form

$$\oint X d(\psi \circ \mathbb{P}) = \int_0^1 F_X^{-1}(\alpha) d\psi(\alpha). \tag{9}$$

For example

$$\mathrm{CV@R}_{1-\alpha}(X) = \frac{1}{\alpha} \int_{1-\alpha}^1 F_X^{-1}(u) du \tag{10}$$

would correspond to a distortion function

$$\psi(u) = \min\left(\frac{u}{1-\alpha}, 1\right) \tag{11}$$

Popular distortion risk measures

Measure Name	**Distortion Function** $g(p)$
Value at Risk (VaR)	$g(p) = \mathbf{1}_{\{p>1-\alpha\}}$
Conditional VaR	$g(p) = \min\left\{\frac{p}{1-\alpha}, 1\right\}$
Proportional Hazard Transform	$g(p) = p^{\gamma}, \quad \gamma > 0$
Wang Transform	$g(p) = \Phi\big(\Phi^{-1}(p) + \lambda\big), \quad \lambda \in \mathbb{R}$
Lookback / Minimax Distortion	$g(p) = 1 - (1-p)^{\gamma}, \quad \gamma > 0$

Table 1: Commonly used distortion risk measures and their distortion functions

Some examples

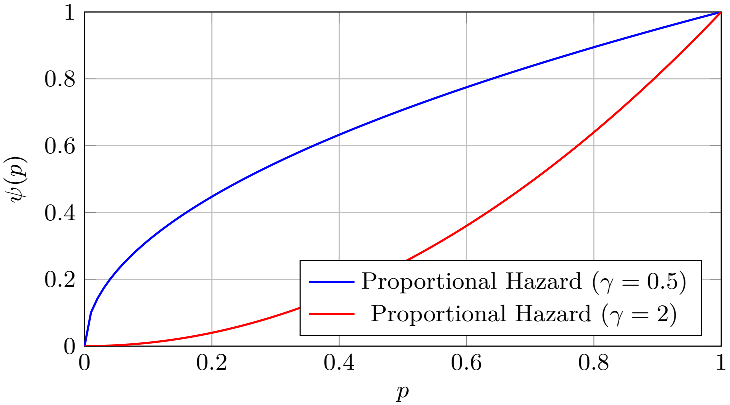

Finding a well suited distribution

To see for which distributions we may find a 'tracktable' variational inequalities lets consider the integrated quantile form of a distortion risk measure

$$\oint u(\boldsymbol{x})d(\psi \circ \mathbb{P}) = \int_0^1 F_{u(\boldsymbol{x})}^{-1}(\alpha)d\psi(\alpha). \tag{12}$$

and recall that $u(\boldsymbol{x}) = \sum_{\boldsymbol{p}\in P} x(\boldsymbol{p})u(\boldsymbol{p})$. And so the distribution of $u(\boldsymbol{x})$ can fully be characterized via the distribution of a random vector $(u(\boldsymbol{p}))_{\boldsymbol{p}\in P}$. And immediately we can see that we need distributions of random vectors that "behave nicely" with respect to positive transformation and summation.

Normal distribution

We know that for a normally distributed random vector $\boldsymbol{X} \sim N_d(\mu, \Sigma)$ we receive that

$$< \boldsymbol{a}, \boldsymbol{X} > \sim N_1(< \boldsymbol{a}, \mu >, \sigma^2(\boldsymbol{a})) \tag{13}$$

where $\sigma^2(\boldsymbol{a}) = \boldsymbol{a}^T \Sigma \boldsymbol{a}$ and in particular

$$< \boldsymbol{a}, \boldsymbol{X} > = < \boldsymbol{a}, \mu > + \sigma(\boldsymbol{a}) Z, \tag{14}$$

where $Z \sim N_1(0, 1)$.

Elliptically symmetric distributions

This can be generalized into the class of elliptically symmetric distributions where for $\boldsymbol{X} \sim Ell_d(\mu, \Sigma, g_d)$ we receive that

$$< \boldsymbol{a}, \boldsymbol{X} >=< \boldsymbol{a}, \mu > +\sigma(\boldsymbol{a})Z, \tag{15}$$

where $Z \sim Ell_d(0, 1, g_1)$.
Recall that the elliptical distribution is characterized by location μ, scale Σ and a d-dimensional density generator g_d such that

$$f(\boldsymbol{x}) \propto g_d((\boldsymbol{x} - \mu)^T \Sigma(\boldsymbol{x} - \mu)). \tag{16}$$

List of some Elliptical distributions

Table 2: Examples of d-dimensional Elliptically Symmetric Distributions

Distribution	**Description**
Multivariate Normal	$\exp\left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})\right)$
Multivariate t	$\left[1+\frac{1}{\nu}(\mathbf{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})\right]^{-(\nu+d)/2}$
Multivariate Laplace	$\exp\left(-\sqrt{(\mathbf{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}\right)$
Multivariate Pearson Type VII	$\left[1+(\mathbf{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})/m\right]^{-(m+d)/2}$
Multivariate Logistic	$\frac{\exp(-(\mathbf{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu}))}{(1+\exp(-(\mathbf{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})))^2}$

t-distribution

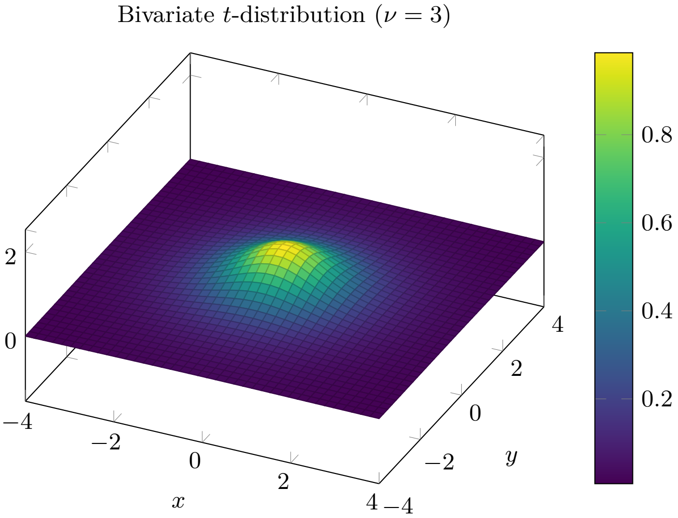

Multivariate Logistic distribution

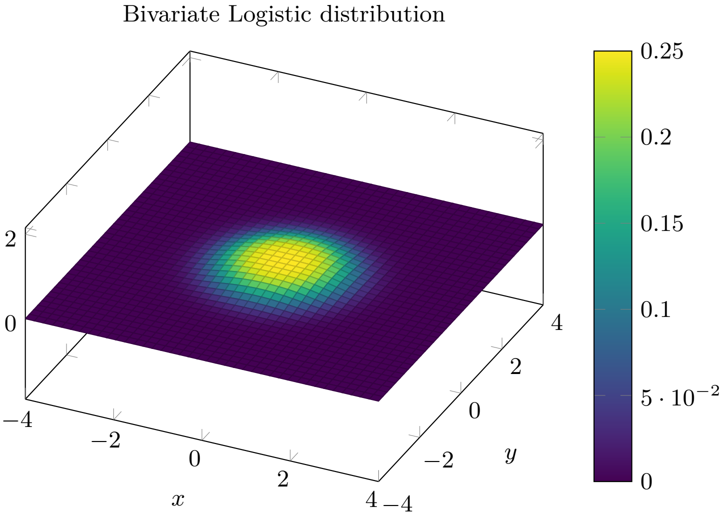

Why is this representation suitable?

Suppose that $(u_i(\boldsymbol{p}))_{\boldsymbol{p}\in P} \sim Ell_K((\mu_i(\boldsymbol{p}))_{\boldsymbol{p}\in P}, \Sigma_i, g_K)$ then $u_i(\boldsymbol{x}) =< \boldsymbol{x}, (u_i(\boldsymbol{p}))_{\boldsymbol{p}\in P} >\sim Ell_1(\mu_i(\boldsymbol{x}), \sigma(\boldsymbol{x}), g_1)$ and

$$F^{-1}_{u_i(\boldsymbol{x})}(p) = \mu_i(\boldsymbol{x}) + \sigma(\boldsymbol{x})F^{-1}_{Z_1}(p), \tag{17}$$

where $Z_1 \sim Ell_1(0, 1, g_1)$. Therefore for a given distortion function ψ_i we can express the distorted payoff as

$$\oint u_i(\boldsymbol{x})d\Psi = \int_0^1 \mu_i(\boldsymbol{x}) + \sigma(\boldsymbol{x})F^{-1}_{Z_1}(p)d\psi(p) = \tag{18}$$

$$= \mu_i(\boldsymbol{x}) + \sigma(\boldsymbol{x})\int_0^1 F^{-1}_{Z_1}(p)d\psi(p) = \tag{19}$$

$$= \mu_i(\boldsymbol{x}) + \sigma(\boldsymbol{x})\Psi(Z_1), \tag{20}$$

where $\Psi(Z_1) \in \mathbb{R}^*$ is a constant independent of the strategies $\boldsymbol{x} \in X$. This also yields one important condition for reasonable distributions that $|\Psi(Z_1)| < \infty$.

Going beyond...

Now let's reverse this process. Let us 'wish' for a nice 'tracktable' distribution and eventually maybe find it? For risk-aware Game theoretical applications nice distributions would have tractable quantiles in the following sense

$$F^{-1}_{<\boldsymbol{a},X>}(\alpha) =< \boldsymbol{a}, \mu > f(\alpha) + \sigma(\boldsymbol{a})g(\alpha), \alpha \in [0,1], \tag{21}$$

for some $f, g : [0,1] \to \mathbb{R}$ which are (at least) Lebesgue-Stieltjes ψ-integrable and $\forall \boldsymbol{a} \in [0,1]^d : \sum_{i=1}^{d} a_i = 1, a_i \geq 0$. If such random vector existed we can express the distorted payoff in terms of 'location' μ and a 'scale' σ and some constants $\Psi(f), \Psi(g)$ as

$$\oint < \boldsymbol{a}, \boldsymbol{X} > d\Psi =< \boldsymbol{a}, \boldsymbol{\mu} > \Psi(f) + \sigma(\boldsymbol{a})\Psi(g). \tag{22}$$

When such characterization is sensible?

Mostly never except the 'trivial' Elliptical case... But if:

1. $|f'(\alpha)| \leq \sqrt{\mu^T \Sigma^{-1} \mu}\, g'(\alpha), \forall \alpha \in [0, 1]$,
2. $\varphi(\boldsymbol{t}) = \int_0^1 \exp(i(<\boldsymbol{t}, \mu> f(\alpha) + \sigma(\boldsymbol{t}) g(\alpha))) d\alpha$ is positive definite.

then there is a well-defined random vector that has projections with the quantile function

$$F^{-1}_{<\boldsymbol{a}, X>}(\alpha) = <\boldsymbol{a}, \mu> f(\alpha) + \sigma(\boldsymbol{a}) g(\alpha).$$

Distorted variational inequalities

Suppose now that we have a sensible distribution class $D_i(\mu, \sigma)$ and $(u_i(\boldsymbol{p}))_{\boldsymbol{p} \in P} \sim D_i(\mu_i, \sigma_i)$ so that $F^{-1}_{u_i(\boldsymbol{x})}(\alpha) =< \boldsymbol{x}, \mu_i > f_i(\alpha) + \sigma_i(\boldsymbol{x}) g_i(\alpha)$ then we can formulate the variational inequality

$$\oint u_i(\boldsymbol{x}) d\Psi_i \geq \oint u_i(y, \boldsymbol{x}_{-i}) d\Psi_i, \forall y \in X_i, i \in I \tag{23}$$

in terms of the problem data as

$$\mu_i(\boldsymbol{x})\Psi_i(f_i) + \sigma_i(\boldsymbol{x})\Psi_i(g_i) \geq \mu_i(y, \boldsymbol{x}_{-i})\Psi_i(f_i) + \sigma_i(y, \boldsymbol{x}_{-i})\Psi_i(g_i), \tag{24}$$

$$(\mu_i(\boldsymbol{x}) - \mu_i(y, \boldsymbol{x}_{-i}))\Psi_i(f_i) \geq (\sigma_i(y, \boldsymbol{x}_{-i}) - \sigma_i(\boldsymbol{x}))\Psi_i(g_i), \tag{25}$$

$$\frac{\mu_i(\boldsymbol{x}) - \mu_i(y, \boldsymbol{x}_{-i})}{\sigma_i(y, \boldsymbol{x}_{-i}) - \sigma_i(\boldsymbol{x})} \geq \frac{\Psi_i(g_i)}{\Psi_i(f_i)}, \tag{26}$$

$\forall y \in X_i$, $i \in I$. Note that usually $\Psi_i(f_i) = 1$.

Interesting consequences

If the parameters μ_i represent the mean payoff of the player i and g_i is Ψ_i-orthogonal i.e. $\int_0^1 g_i(\alpha) d\psi(\alpha) = 0$ then the player i actually strategizes in the same way as if he was maximizing the expectation.

$$\frac{\mu_i(\boldsymbol{x}) - \mu_i(y, \boldsymbol{x}_{-i})}{\sigma_i(y, \boldsymbol{x}_{-i}) - \sigma_i(\boldsymbol{x})} \geq \frac{\Psi_i(g_i)}{\Psi_i(f_i)} \tag{27}$$

Thank you for your attention!

Its time for the discussion...