Strategic behavior of risk-averse agents under stochastic market clearing

Vincent Leclère, Andy Philpott

[a]CERMICS, ENPC, Paris, France
[b]Electric Power Optimization Centre, University of Auckland, New Zealand

Abstract

We discuss a model that uses stochastic programming to compute economic dispatch and system marginal prices in a single-settlement wholesale electricity pool. Agents are endowed with coherent risk measures, and maximize risk-adjusted profit in a market with complete trading of risk. We show that if agents' risk measures are known by the system operator then prices from a socially optimal dispatch are revenue adequate and recover agents' costs in risk-adjusted expectation. The paper then explores incentives on agents to truthfully reveal risk preferences to the system operator, and constructs a non-cooperative game to model this. We show by example that agents have incentives to misrepresent their risk measures to improve their risk-adjusted profit.

Keywords: electricity market, risk, Nash equilibrium

1. Introduction

Wholesale electricity markets in many jurisdictions are cleared using mathematical programming problems. These are often formulated as mixed integer programming problems to account for unit commitment effects. In the absence of such indivisibilities, the *economic dispatch problems* that optimize production are convex and yield optimal levels of electricity generation, transmission flows, and locational marginal prices that come from the dual variables of flow conservation constraints. Electricity consumers at any location pay the locational marginal price, and generation at a location is paid the locational marginal price.

The last decade has seen dramatic growth in renewable electricity generation from wind and solar energy. Wind and sunshine are intermittent and uncertain, and so backup thermal generation is often required to cover periods when wind and sunshine are unavailable. Most forms of thermal generation must be made available before wind strength and solar production are observed. If these are modeled as random variables with known probability distributions then the most efficient dispatch on average will come from the solution to a stochastic programming problem that maximizes expected welfare. Generally speaking, the solution to stochastic programs is computable only when the probability distributions are finite. This enables the stochastic programming problem to be written as a finite-dimensional optimization problem. When the distribution is continuous, sample average approximation, which constructs finite distributions by sampling, converges to the optimal solution set under mild conditions and is a popular approach to solving stochastic programs.

In the context of wholesale electricity markets we can formulate a stochastic economic dispatch problem with a finite number of scenarios, which yields optimal output levels for each generator, transmission flow, and locational marginal prices, all defined for each scenario. Stochastic dispatch and pricing mechanisms of this form have been studied for over ten years (see [1],[2],[3]).

Three properties would be desirable in such a pricing mechanism:

- Budget balance: is the pricing mechanism *revenue adequate*? This means that the amount paid by purchasers under the mechanism is at least sufficient to cover the amount paid to sellers. Revenue adequacy allows for a market surplus that the market maker can capture. If there were a deficit, external funds would need to be supplied to the market to ensure budget balance.

- Participation: does the pricing mechanism guarantee *cost recovery*? This condition asserts that prices must cover supply costs and guarantee some type of *individual rationality*. Indeed, if prices are below the cost of supply, then a rational supplier would refuse to participate in the market.

- Truth revealing: is the pricing mechanism *incentive compatible*? Incentive compatibility ensures that agents are incentivized to truthfully reveal their private information to the market maker. In other words, in the non-cooperative game, which is defined by supplying information to the market maker, truthfully revealing private information gives a Nash equilibrium.

As shown by [4], it is not possible for any single mechanism to have all three properties. Pritchard et al. [3] describes a pricing mechanism that is individually rational in every scenario while being revenue-adequate in expectation (only). In contrast, Zakeri et al. [5] present a model that is revenue-adequate in every scenario while being individually rational in expectation (only). In fact, as shown by [6], it is not possible to have a dispatch solution that maximizes expected welfare and yields prices that provide revenue adequacy and individual rationality in every scenario. This is not surprising as stochastic programming solutions hedge against future uncertainty, and so ex-ante dispatch and pricing outcomes may be suboptimal ex-post.

In this paper, we explore stochastic dispatch and pricing mechanisms that account for the risk aversion of agents when modeled using coherent risk measures under an assumption of complete risk markets. Our discussion of optimization and equilibrium in this setting draws heavily on the theory of coherent risk measures and risked competitive equilibrium studied in the papers ([7],[8],[9]).

The contributions of the paper are as follows:

1. We show how risk measures can be included in economic dispatch models to give risk-adjusted generation output and cost recovery in risk-adjusted expectation.

2. We describe a competitive market structure that requires agents to declare their level of risk aversion (and costs) to a system operator.

3. We show how incentive compatibility might fail in such a market by describing a game in which every Nash equilibrium involves some player misrepresenting their true level of risk aversion.

The rest of the paper is laid out as follows. In section 2 we define the risk-adjusted economic dispatch problem and prove that a standard pricing mechanism [5] is revenue adequate and has cost recovery in risk-adjusted expectation. The pricing mechanism defines a non-cooperative game in which agents each declare a level of risk aversion to the system operator, who then computes a socially optimal economic dispatch and prices based on this. Agents then trade contracts to reduce their actual risk as measured by their true risk measures. Some structural properties of this game are then studied in Section 3, which presents some characterizations of Nash equilibrium. Section 4 presents a small example to demonstrate that the mechanism defined by the game is not incentive-compatible. Agents can improve their ex-post risk-adjusted expected return by misrepresenting their level of risk aversion to the system operator. Section 5 then concludes the paper.

2. Risk-adjusted economic dispatch

In this section, we first recall the risk-neutral stochastic dispatch model and its extension to the risk-averse setting. We then introduce a strategic game where agents strategically display their risk aversion to the ISO. This game is a special case of the general model studied in the next section.

2.1. Stochastic dispatch model

We recall the stochastic electricity dispatch model first formulated in [3] and and studied in [6], [5], and [10]. These models are formulated as stochastic linear programs in a transmission network where generation and demand are located at nodes denoted $n \in \mathcal{N}$. This model can be easily extended to accommodate convex generator cost functions. We denote the agents by index a and define $\delta_{an} = 1$ when agent a is located at node n and 0 otherwise. We further use the following notation.

- x_a is the day-ahead setpoint level which generator a is advised to prepare to produce before the generation capacity of intermittent renewables is known. The cost of this is a convex function $c_a(x_a)$.

- $X_a(\omega)$ is the real-time dispatch produced by generator a in scenario $\omega \in \Omega$.

- $U_a(\omega)$ is the amount by which generator a deviates from x_a in scenario ω.

- r_a is a finite convex function such that $r_a\left(U(\omega)\right)$ is the cost incurred by generator a for deviating from its generation.

- $F(\omega) \in \mathcal{F}$ is the vector of branch flows in the network in scenario ω, where $\mathcal{F}$ is a set constraining the flows in the network to meet thermal limits and the DC-Load Flow constraints imposed by Kirchhoff's Laws. We assume that $0 \in \mathcal{F}$.

We denote by $\tau_n(F(\omega))$ the net amount of energy flowing from the transmission network into node n in scenario ω. We assume that τ_n is a concave function of F with $\tau_n(0) = 0, \ \forall n \in \mathcal{N}$. The other parameters in the model are:

- $G_a(\omega)$ the maximum output capacity of generator a in scenario ω.
- $D_n(\omega)$ the consumer demand at node n in scenario ω.

Given these parameters, the dispatch model is:

$$\begin{array}{llll} \text{DP:} \ \min & \sum_a c_a\left(x_a\right) + \sum_a \sum_\omega \mathbb{P}(\omega) r_a\left(U_a(\omega)\right) & & \\ \text{s.t.} & \sum_a \delta_{an} X_a(\omega) + \tau_n(F(\omega)) \geq D_n(\omega), & n \in \mathcal{N}, \omega \in \Omega, & [\mathbb{P}(\omega)\pi_n(\omega)] \\ & x + U(\omega) = X(\omega), & \omega \in \Omega, & \\ & F(\omega) \in \mathcal{F}, & \omega \in \Omega, & \\ & 0 \leq X(\omega) \leq G(\omega), x \geq 0, & \omega \in \Omega. & \end{array}$$

The first constraint ensures demand satisfaction in each node and scenario, the second defines real-time dispatch, and the two last equations ensure the physical acceptability of the flows. The term in square brackets is the probability-weighted dual variable for the flow balance constraint at node n.

After the optimal day-ahead setpoint x^* is found, the intermittent generation scenario $\omega = \hat{\omega}$ is realised, and the ISO follows the dispatch defined by $\left(X^*(\hat{\omega}), U^*(\hat{\omega}), F^*(\hat{\omega})\right)$. The optimal dispatch gives nodal prices $\pi_n(\omega)$ which are used to compensate generators for the dispatch. As discussed in [6] there are different ways of doing this: we consider the payment mechanism that pays $\pi_n(\omega)X_a(\omega)$ to all generators a at node n and charges $\pi_n(\omega)D_n(\omega)$ to demand at node n, which we denote the RA mechanism.

Definition 1. *A payment mechanism is revenue adequate if and only if in every scenario $\omega \in \Omega$, clearing the market does not leave the system operator in a financial deficit. As shown by [11], revenue adequacy is equivalent to the following statement:*

$$\sum_n \pi_n(\omega)\tau_n(F(\omega)) \geq 0, \ \forall \omega \in \Omega.$$

The following result was demonstrated by [5] for linear costs.

Proposition 1. *If* $(x^*, X^*(\omega), U^*(\omega), F^*(\omega))$ *solves DP, then the RA mechanism is revenue adequate.*

Definition 2. *A payment mechanism exhibits cost recovery if and only if, in every scenario* $\omega \in \Omega$*, all generators recover their short-run (fuel and deviation) costs. That is,*

$$R_a(\omega) - c_a(x_a) - r_a\left(U_a(\omega)\right) \geq 0, \quad \forall a, \quad \forall \omega \in \Omega,$$

where $R_a(\omega)$ *is generator* a*'s revenue in scenario* ω*.*

We say that a market clearing mechanism exhibits *expected cost recovery* if all generators recover their generation and ramping costs in expectation. This was shown to be the case for the RA mechanism with linear costs by [5].

2.2. Risk-adjusted dispatch

We now consider the case explored in [10] where each agent a is endowed with a coherent risk measure. A risk measure ρ is a mapping from a space of random variables measuring loss to real numbers (and $+\infty$). A *coherent* risk measure (see *e.g.*, [12] for definition) can be expressed as the worst-case expectation of loss Z over a convex set of probability distributions, *i.e.*,

$$\rho(Z) = \max_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}[Z]. \tag{1}$$

The set $\mathcal{Q}$ is called the *risk set* of the risk measure ρ. We denote the risk set of agent a by $\mathcal{Q}^a$.

The theory of risked equilibrium we use is presented in detail by [8]. Following their model, we assume that each $\mathcal{Q}^a$ is polyhedral and interior to the positive orthant with a reference probability distribution $\mathbb{P}_0 \subseteq \cap \mathcal{Q}^a$. We also assume that there is a complete market for trading risk using Arrow-Debreu securities. An *Arrow-Debreu* security for scenario ω returns a payoff of 1 if scenario ω occurs while requiring an advance payment of $\mu(\omega)$ in the first stage. A *risked equilibrium* is a set of prices for each scenario and a collection of generation, consumption, and Arrow-Debreu trading actions that minimize risk-adjusted disbenefit for each agent. By Theorem 4, Corollary 2 in [8], a perfectly competitive risked equilibrium can be found by the system operator by solving

$$\begin{array}{llll} \text{RADP:} & \min & \max_{\mathbb{Q}\in\cap\mathcal{Q}^a} & \sum_\omega \mathbb{Q}(\omega) \sum_a \left(c_a\left(x_a\right) + r_a\left(U_a(\omega)\right)\right) & \\ & & \text{s.t.} & \sum_a \delta_{an} X_a(\omega) + \tau_n(F(\omega)) \geq D_n(\omega), & n \in \mathcal{N}, \omega \in \Omega, \\ & & & x + U(\omega) = X(\omega), & \omega \in \Omega, \\ & & & F(\omega) \in \mathcal{F}, & \omega \in \Omega, \\ & & & 0 \leq X(\omega) \leq G(\omega), x \geq 0, & \omega \in \Omega. \end{array}$$

RADP solves a risk-averse social planning problem using a risk set that is the intersection of all the agent's risk sets.

Proposition 2. *RADP yields locational marginal energy prices that are revenue adequate in every scenario, and, after trading Arrow-Debreu securities yielding random returns $A_a(\omega)$, each agent a recovers costs in risk-adjusted expectation,* i.e.*,*

$$\min_{x_a, U_a, A_a} \max_{\mathbb{Q}\in\mathcal{Q}^a} \mathbb{E}_{\mathbb{Q}} \left[c_a(x_a) + r_a(U_a(\omega)) - \sum_n \delta_{an} \pi_n x_a - A_a(\omega) \right] \leq 0. \tag{2}$$

Proof. RADP yields a competitive equilibrium, where each agent (including the consumer) optimizes their risk adjusted expected welfare augmented with returns from trading Arrow-Debreu securities. In equilibrium, the returns for generators must recover risk-adjusted costs, and so (2) holds. To show revenue adequacy, consider the extreme points of $\cap\mathcal{Q}^a$ which are vectors p^k, $k \in \mathcal{K}$, and express the inner maximum using an epigraphical variable

$$\theta \geq \sum_\omega p^k(\omega) \sum_a r_a(U_a(\omega)), \quad k \in \mathcal{K}.$$

Then RADP has a solution that minimizes the Lagrangian

$$\begin{aligned} \mathcal{L}(x, X, U, F, \theta) \;\; = \;\; & \sum_a c_a(x_a) + \theta + \sum_k \mu^k (\sum_\omega p^k(\omega) \sum_a r_a(U_a(\omega)) - \theta) \\ & + \sum_n \sum_\omega \pi_n(\omega) \left(\sum_k \mu^k p^k(\omega) \right) (D_n(\omega) - \sum_a \delta_{an} X_a(\omega) - \tau_n(F(\omega))) \\ & + \sum_\omega \sum_a \sigma_a(\omega) \left(\sum_k \mu^k p^k(\omega) \right) (x_a + U_a(\omega) - X_a(\omega)) \end{aligned}$$

over the bound constraints on the variables, where the Lagrange multipliers (π and σ) on the other constraints of RADP are weighted by the positive quantity $\sum_k \mu^k p^k(\omega)$. In particular the solution maximizes

$$\sum_\omega \left(\sum_k \mu^k p^k(\omega)\right) \sum_n \pi_n(\omega)\tau_n(F(\omega))$$

for each ω independently over $F(\omega) \in \mathcal{F}$ and since $F = 0$ is feasible, we have for every ω,

$$\left(\sum_k \mu^k p^k(\omega)\right) \sum_n \pi_n(\omega)\tau_n(F(\omega)) \geq 0. \qquad \square$$

2.3. Strategic display of risk aversion

The mechanism above requires the truthful revelation of each agent's actual risk set $\mathcal{P}^a$ to optimize risk-adjusted social welfare. Each agent, however, might misrepresent their actual risk set as $\mathcal{Q}^a \neq \mathcal{P}^a$ to improve their risk-adjusted profit. This leads to a non-cooperative game between agents with the following sequence of steps.

1. Each agent declares a risk set $\mathcal{Q}^a$ and supplies this to the ISO. The ISO assumes this is truthful and solves an economic dispatch problem

$$\min_{x,U} \max_{\mathbb{Q}\in\cap\mathcal{Q}^a} \mathbb{E}_\mathbb{Q}\left[\sum_a (c_a(x_a) + r_a(U_a(\omega)))\right] \tag{3}$$

   where we define

$$\mathcal{Q}^\cap := \bigcap_{a\in\mathcal{A}} \mathcal{Q}^a. \tag{4}$$

2. Each agent fixes its dispatch but trades Arrow-Debreu securities to improve its risk-adjusted payoff when evaluated using its private risk set $\mathcal{P}^a$.

3. The improved risk-adjusted payoff for agent a can be computed by evaluating

$$\max_{\mathbb{Q}\in\mathcal{P}^\cap} \mathbb{E}_\mathbb{Q}\left[\sum_a (c_a(x_a) + r_a(U_a(\omega)))\right]$$

   where

$$\mathcal{P}^\cap := \bigcap_{a\in\mathcal{A}} \mathcal{P}^a. \tag{5}$$

and then evaluating each agent's payoff at the maximizing $\mathbb{Q}$.

Observe that the payoff for each agent depends on $(\mathcal{Q}^a, \mathcal{Q}^{-a})$ as a function of the intersection $\mathcal{Q}^\cap$ and depends on $(\mathcal{P}^a, \mathcal{P}^{-a})$ as a function of the intersection $\mathcal{P}^\cap$. This observation allows us to narrow down the form of the equilibrium, as discussed in the next section.

3. The risk-set game

The game described in the previous section has a special structure enabling us to derive general results. To do this, we let $(\Omega, \mathcal{F}, \mathbb{P}^0)$ be a discrete probability space and denote $\mathcal{M}$ the set of probability distributions on $(\Omega, \mathcal{F})$. We consider a set $\mathcal{R} \subset 2^{\mathcal{M}}$ of *admissible risk sets* satisfying[1]

i) if $\mathcal{Q}, \tilde{\mathcal{Q}} \in \mathcal{R}$, then $\mathcal{Q} \cap \tilde{\mathcal{Q}} \in \mathcal{R}$,

ii) for all $\mathcal{Q} \in \mathcal{R}$, $\mathbb{P}^0 \in \mathcal{Q}$.

Here, $\mathcal{R}$ defines the set of possible actions for each player in the game.

Example 1. *To fix ideas, we assume that $\mathbb{P}_0$ is the uniform measure over $\xi^1, \dots, \xi^N$. Then any probability measure in $\mathcal{M}$ is uniquely defined by a vector $p \in \Delta^{N-1}$ where $\Delta^{N-1} \subset \mathbb{R}^N_+$ is the $N-1$ dimensional simplex, that is such that $\sum_i p_i = 1$. We now discuss a few natural admissible risk sets $\mathcal{R}$ (uniquely defined as a subset of Δ^{N-1}) that satisfy the above condition.*

1. *If we assume that the agent can declare any polyhedral risk measure to the ISO, then $\mathcal{R}$ is the set of all polyhedrons in Δ^{N-1};*

2. *If the agent can declare any $AVAR_\alpha$ risk measure, then $\mathcal{R}$ is the set of polyhedron of the form $\{p \in \Delta^{N-1} \mid p_i \leq \frac{\alpha}{N}, \forall i\}$;*

3. *If we assume that we set a coherent risk measure ρ, and each agent can declare a risk measure of the form $t\rho + (1-t)\mathbb{E}_{\mathbb{P}_0}$, for $t \in [0, 1]$, then the admissible risk sets are of the form $(1-t)\{\mathbb{P}^0\} + t\mathcal{D}$, for some convex set $\mathcal{D}$. This specific case is further detailed in section 3.2.*

[1]The results of this section can be directly transposed to the case where $\mathcal{R}$ is a lower semilattice (that is, an ordered set such that for all $\mathcal{Q}^1, \mathcal{Q}^2 \in \mathcal{R}$, there exists a meet - *i.e.*, largest lower bound - $\mathcal{Q}^1 \wedge \mathcal{Q}^2$) by simply replacing the intersection $\cap$ by the meet $\wedge$.

3.1. Set intersection games

Formally, we consider a game with a finite number of agents, where each agent has the same set $\mathcal{R}$ of admissible actions. The actions of all agents are collected in a *strategy profile* $(\mathcal{Q}^1, \dots, \mathcal{Q}^{|\mathcal{A}|})$ denoted $(\mathcal{Q}^{a'})_{a' \in \mathcal{A}}$. The loss of any specific agent a is a function $L^a((\mathcal{Q}^{a'})_{a' \in \mathcal{A}})$ of the strategy profile collecting every agent's actions. To identify the effect of $\mathcal{Q}^a$ on a's loss we sometimes write $L^a((\mathcal{Q}^{a'})_{a' \in \mathcal{A}}) = L^a(\mathcal{Q}^a, \mathcal{Q}^{-a})$.

The key assumption in this section is that a strategy profile $((\mathcal{Q}^{a'})_{a' \in \mathcal{A}}$ impact the game only through the intersection of all actions $\cap_{a'} \mathcal{Q}^{a'}$. More precisely, using the following notation:

$$\mathcal{Q}^{\cap} = \bigcap_{a' \in \mathcal{A}} \mathcal{Q}^{a'}, \quad \mathcal{Q}^{\cap - a} = \bigcap_{a' \in \mathcal{A} \setminus \{a\}} \mathcal{Q}^{a'}, \tag{6}$$

we have that, for every agent a, the loss function L^a is such that

$$L^a((\mathcal{Q}^{a'})_{a' \in \mathcal{A}}) = L^a(\mathcal{Q}^{\cap}, \dots, \mathcal{Q}^{\cap}),$$

which will be simply denoted $\ell^a(\mathcal{Q}^{\cap})$. This assumption implies the following lemma.

Lemma 1. *Let $a \in \mathcal{A}$ be an agent. Given a strategy profiles $(\mathcal{Q}^{a'})_{a' \in \mathcal{A}}$, and an action $\tilde{\mathcal{Q}}^a \in \mathcal{R}$ we have,*

$$L^a(\tilde{\mathcal{Q}}^a, \mathcal{Q}^{-a}) = \ell^a(\tilde{\mathcal{Q}}^a \cap \mathcal{Q}^{\cap - a}).$$

From this lemma, we derive a first set of results.

Remark 1. *Consider a strategy profile $(\mathcal{Q}^{a'})_{a' \in \mathcal{A}}$, and an agent $a \in \mathcal{A}$. We have the following results:*

1. *as $\mathbb{P}_0 \in \mathcal{Q}^a$ for all possible action action, the symmetric strategy profile $(\{\mathbb{P}_0\})_{a \in \mathcal{A}}$ is a Nash equilibrium;*

2. *if at least two players play $\{\mathbb{P}^0\}$, then we have a Nash equilibrium.*

We now provide the necessary and sufficient conditions for a strategy profile to be a Nash Equilibrium. It states that no agent can profit from playing a set smaller than the intersection of the other player's action.

Proposition 3. *The strategy profile* $(\mathcal{Q}^{a'})_{a' \in \mathcal{A}}$ *is a Nash equilibrium if and only if, for each player* a*, any choice* $\mathcal{Q} \in \mathcal{R}$ *satisfies*

$$\mathcal{Q} \subseteq \mathcal{Q}^{\cap -a} \quad \Rightarrow \quad \ell^a(\mathcal{Q}^{\cap}) \leq \ell^a(\mathcal{Q}). \tag{7}$$

Proof. If $(\mathcal{Q}^{a'})_{a' \in \mathcal{A}}$ is a Nash equilibrium, then for all agents $a \in \mathcal{A}$, and all actions $\mathcal{Q} \subset \mathcal{Q}^{\cap -a}$ we have $\ell^a(\mathcal{Q}^{\cap}) = L^a((\mathcal{Q}^{a'})_{a' \in \mathcal{A}}) \leq L^a(\mathcal{Q}, \mathcal{Q}^{-a}) = \ell^a(\mathcal{Q})$.

Now suppose that $(\mathcal{Q}^{a'})_{a' \in \mathcal{A}}$ is not a Nash equilibrium. Let $a \in \mathcal{A}$ be an agent that can improve its payoff by playing $\tilde{\mathcal{Q}}^a \in \mathcal{R}$. Then, we have $L^a(\tilde{\mathcal{Q}}^a, \mathcal{Q}^{-a}) = \ell^a(\tilde{\mathcal{Q}}^a \cap \mathcal{Q}^{\cap -a}) < \ell^a(\mathcal{Q}^{\cap})$. If we let $\mathcal{Q} = \tilde{\mathcal{Q}}^a \cap \mathcal{Q}^{\cap -a}$ then $\mathcal{Q} \in \mathcal{R}$ but $\mathcal{Q}$ does not satisfy (7). □

A straightforward consequence of Proposition 3 is that we can restrict ourselves to the study of symmetric Nash equilibrium, as shown by the following corollary.

Corollary 1. *If* $(\mathcal{Q}^{a'})_{a' \in \mathcal{A}}$ *is a Nash equilibrium, then* $(\mathcal{Q}^{\cap})$ *is also a Nash equilibrium with the same payoffs.*

Further, a symmetric strategy (Q) *is a Nash equilibrium if and only if for all risk sets* $\tilde{Q} \subset Q$ *and all agents* $a \in \mathcal{A}$ *we have* $\ell^a(\tilde{Q}) \geq \ell^a(Q)$*.*

Proof. Since

$$\mathcal{Q} \subseteq \mathcal{Q}^{\cap} \quad \Rightarrow \quad \mathcal{Q} \subseteq \mathcal{Q}^{\cap -a} \quad \Rightarrow \quad \ell^a(\mathcal{Q}^{\cap}) \leq \ell^a(\mathcal{Q}),$$

because $(\mathcal{Q}^{a'})_{a' \in \mathcal{A}}$ is a Nash equilibrium, the result follows by Proposition 3.

The second part is a special case of Proposition 3 □

Another consequence of Proposition 3 is that the loss incurred at a (symmetric) Nash equilibrium is non-increasing with respect to inclusion, or, in other words, the best Nash-Equilibrium is the one resulting from the most risk-averse declaration. More precisely, we have the following result.

Corollary 2. *Let* (Q) *and* $(\tilde{Q})$ *be two symmetric Nash equilibria such that* $Q \subset \tilde{Q}$*. Then* $(\tilde{Q})$ *Pareto-dominates* (Q)*, that is, for every agent* a*,* $\ell^a(Q) \geq \ell^a(\tilde{Q})$*.*

Proof. If $Q \subset \tilde{Q}$, then applying Proposition 3 to the Nash equilibrium $(\tilde{Q})$ implies $\ell^a(\tilde{Q}) \leq \ell^a(Q)$ for every agent a. □

3.2. One parameter case

In some settings, the set of admissible risk sets $\mathcal{R}$ can be defined by a small number of parameters. We are especially interested in the case where $\mathcal{R}$ is defined by a single parameter $t \in [0,1]$, such that if $t \leq t'$, we have $\mathcal{Q}_t \subset \mathcal{Q}_{t'}$, with $\mathcal{Q}_0 = \{\mathbb{P}_0\}$. Such a case arises for example if we consider $\mathcal{D} = \partial\rho(0)$ where ρ is a coherent risk measure (*e.g.*, Average Value at Risk), and define

$$\mathcal{R} = \left\{(1-t)\{\mathbb{P}^0\} + t\mathcal{D} \mid t \in [0,1]\right\}.$$

Proposition 4. *In the one-parameter case, we have the following results.*

1. *$(\mathcal{Q}_t)$ is a Nash equilibrium if and only if, for all $t \leq \bar{t}$, and all $a \in \mathcal{A}$, we have $\ell^a(\mathcal{Q}_t) \geq \ell^a(\mathcal{Q}_{\bar{t}})$;*

2. *If $(\mathcal{Q}_t)$ and $(\mathcal{Q}_{\bar{t}})$ are two Nash equilibrium, with $\bar{t} \geq t$, then $(\mathcal{Q}_{\bar{t}})$ Pareto-dominates $(\mathcal{Q}_t)$;*

3. *Let $\mathcal{T}$ be the set of $t \in [0,1]$ such that, for all $a \in \mathcal{A}$, $\tau \mapsto \ell^a(\mathcal{Q}_\tau)$ is non-increasing on $[0,t]$. Then, $\mathcal{T}$ is a non-empty set of Nash equilibrium of the form $[0,t^*)$ or $[0,t^*]$.*

Nash equilibria in our one-dimensional framework are typically not unique. We say that Nash equilibrium $t^\sharp$ *dominates* Nash equilibrium t if every agent's payoff in $t^\sharp$ equals or exceeds its payoff in t, with at least one agent strictly better off. Nash equilibrium $t^\sharp$ is a (locally) *nondominated equilibrium* if (there exists a neighborhood in which) no other Nash equilibrium dominates it.

Proposition 5. *If $\mathcal{T} = [0,t^\sharp]$, which is the case if, for each agent a, $\sup_{t\in\mathcal{T}} \ell_a(\mathcal{Q}_t)$ is attained, then $t^\sharp$ defines a locally non-nondominated Nash Equilibrium while every $t \in [0,t^\sharp)$ defines a (locally) dominated Nash equilibrium.*

Remark 2. *Unfortunately, $\mathcal{T}$ does not describe all equilibrium. Indeed, in figure 1, $\mathcal{T} = [0,t_a]$, while the set of Nash Equilibrium is determined by $[0,t_a] \cup [t_b,t_c]$.*

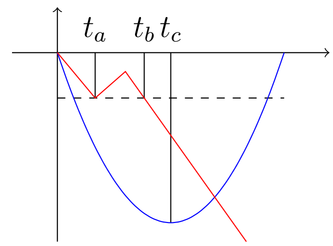

Figure 1: Example of Nash Equilibrium. We have a two-agent game, each curve representing the loss of one agent with respect to $\mathcal{Q}_t$, where $t = \min(t_1, t_2)$ is the minimum of the action of both players. We see that the set of Nash equilibrium is $[0, t_a] \cup [t_b, t_c]$.

4. Stochastic dispatch example

In the previous section, we studied the structure of the Nash Equilibrium associated with the risk-set game. We have seen that there always exists a Nash Equilibrium, which consists of all agents declaring to be risk neutral. Still, it is also the worst possible Nash-Equilibrium, in the sense that any other would Pareto-dominate it. Nothing was said, however, about truly revealing its risk set. The monotonicity property shown would lend to the idea that the agent should declare the largest set possible to find a larger Nash Equilibrium, resulting in lower cost. This idea is, unfortunately, false. Indeed, the following toy example showcases how a player might benefit from declaring a smaller risk set and another a larger risk set.

4.1. Toy example setting

Consider a situation with two agents: an electricity generator operating a thermal plant, a wind farm, and an electricity consumer. We consider a single period with two decision stages. In the first stage, the generator chooses a forward dispatch x for its thermal plant. In the second stage, a wind generation outcome $\xi(\omega)$ is observed, and the generator then dispatches extra energy $U(\omega)$ from the thermal plant. The total generation is $x + U(\omega) + \xi(\omega)$ which equals the consumer demand. This is an example of RADP where V is absent.

Recall that the inverse demand function $P(z)$ gives the marginal welfare that the consumer receives from consuming z. Thus the consumer's welfare when consuming $X(\omega)$ is expressed as $\int_0^{X(\omega)} P(z)dz$ while the associated cost is $X(\omega)P(X(\omega))$, resulting in a consumer disbenefit of $Z_c(\omega) =$

$\int_0^{X(\omega)} P(z)dz - X(\omega)P(X(\omega))$. On the other hand, the generator disbenefit Z_g is the cost of generation minus the revenue from selling the energy.

From now on, we assume that the inverse demand curve is $P(z) = a - bz$, the generator has generation cost $\frac{1}{2}cx^2$ in stage 1, and $\frac{1}{2}dU(\omega)^2$ in stage 2. Straightforward computation shows that the total disbenefit is

$$Z_s(\omega) = \frac{1}{2}cx^2 + \frac{1}{2}dU(\omega)^2 + \frac{1}{2}b(x+U(\omega))^2 - (a-b\xi)(x+U(\omega)) + \frac{1}{2}b\xi^2 - a\xi. \quad (8)$$

Note that, given x, and realization of $\xi(\omega)$ in scenario ω the ISO will choose $U(\omega)$ to minimize $Z_s(\omega)$ giving $U(\omega) = (a - b\xi - bx)/(b + d)$.

For the risk sensitivity, we assume that both agents consider a risk measure of the form

$$\rho_\lambda(Z) = (1-\lambda)\mathbb{E}[Z] + \lambda\mathbb{W}[Z] \quad (9)$$

where $\mathbb{W}$ is the worst-case risk measure, and $\lambda \in [0, 1]$. This falls into the setting of section 3.2. We denote λ_g and λ_c as the true risk parameters of the generator and the consumer, respectively. Following the setting of section 2.3, the agents do not have to reveal their risk parameters truthfully but instead declare t_g and t_c, respectively. Assuming truthful revelation, the system operator then chooses $t = \min\{t_g, t_c\}$ and computes a dispatch minimizing $\rho_t(Z_s)$.

To further simplify the example, we assume that the wind generation outcome ξ is either g or h with equal probability and that $g < h$. We also choose the following numerical values: $a = 200$, $b = 1$, $c = 2$, $d = 1$, $g = 10$, $h = 90$.

4.2. System operator dispatch

Based on the declared risk parameter t, the system operator computes the dispatch that minimizes the risk-adjusted disbenefit $\rho_t(Z_s)$. In particular, we have $U(\xi) = 100 - \xi - x$, and the social disbenefit is

$$Z_s(\omega) = \frac{5}{4}x^2 - 100(\xi + x) + \frac{1}{4}\xi^2 + \frac{1}{2}x\xi - 10000.$$

We can then see that, for all $x \leq 200$, $Z_s(\xi = g) \geq Z_s(\xi = h)$, meaning that the worst scenario is the low wind outcome $\xi = g$, *i.e.*, $\mathbb{W}[Z_s] = Z_s(g)$. We can, therefore, deduce the risk-adjusted disbenefit for the system operator as

$$\rho_t(Z_s) = \frac{5}{4}x^2 - 20tx - 75x + 3000t - 13975,$$

which has minimizer $x = 8t + 30$ yielding in turn

$$U(\omega) = \begin{cases} 80 - 4t, & \xi = g \\ 40 - 4t, & \xi = h. \end{cases}$$

Using $p(\omega) = P(x + U(\omega) + \xi)$, we get the prices in each scenario

$$p(\omega) = \begin{cases} 80 - 4t, & \xi = g \\ 40 - 4t, & \xi = h. \end{cases}$$

4.3. Agent response to dispatch

In this example of a risk set game, the agent declares a risk parameter t_c (resp. t_g) to the system operator, which deduces a dispatch as seen above. The agent can then trade risk minimizing disbenefit using their true risk-measure, *i.e.*, $\rho_{\lambda_g}(Z_g)$ for the generator (resp. $\rho_{\lambda_c}(Z_c)$) for the consumer. Note that if the market for risk is complete, then this leads to a risked equilibrium with effective risk measure parameter $\lambda = \min\{\lambda_g, \lambda_c\}$.

In the following sections, we show that, in this example, the generator has an incentive to declare $t_g = 0$ to the system operator, while the consumer has an incentive to declare $t_c = 1$.

4.3.1. Generator

Recall that in scenario ω, the generator disbenefit is

$$Z_g(\omega) = \frac{1}{2}cx^2 + \frac{1}{2}dU(\omega)^2 - P(\omega)(x + U(\omega) + \xi(\omega))$$

which simplifies to

$$\begin{aligned} Z_g(g) &= 88t^2 + 280t - 5500 \\ Z_g(h) &= 88t^2 + 440t - 4700 \end{aligned}$$

after substituting $x = 8t + 30$. Thus, for all values of t, the worst case disbenefit is $Z_g(h)$. The risk-adjusted disbenefit for the generator is then

$$\rho_{\lambda_g}(Z_g) = \frac{1}{2}(1 - \lambda_g)(88t^2 + 280t - 5500) + \frac{1}{2}(1 + \lambda_g)(88t^2 + 440t - 4700).$$

Irrespective of λ_g, this function is increasing in t, so the generator would prefer the system operator to use $t_g = 0$. It has an incentive to misrepresent itself as less risk-averse than it really is.

4.3.2. Consumer

Recall that in scenario ω, the consumer disbenefit is

$$Z_c(\omega) = -W(\omega) + P(\omega)(x + U(\omega) + \xi(\omega))$$

Using that $x = 8t + 30$ and $U(\omega) = \begin{cases} 80 - 4t, & \xi = g \\ 40 - 4t, & \xi = h \end{cases}$ we get

$$\begin{aligned} Z_c(\omega) &= \frac{1}{2}b(x + U(\omega))^2 - (a - b\xi)(x + U(\omega)) + \frac{1}{2}b\xi^2 - a\xi \\ &\quad + (a - b(x + U(\omega) + \xi(\omega)))\,(x + U(\omega) + \xi(\omega)) \\ &= \begin{cases} -8\,(t + 30)^2\,, & \xi = g \\ -8\,(t + 40)^2\,, & \xi = h. \end{cases} \end{aligned}$$

The worst case scenario is then the low wind scenario $\xi = g$, is in the low wind scenario. Thus, the consumer risk-adjusted disbenefit is given by

$$\frac{1}{2}(1+\lambda)(-8\,(t+30)^2)+\frac{1}{2}(1-\lambda)(-8\,(t+40)^2) = 2800\lambda+(80\lambda-560)t-8t^2-10\,000.$$

In particular, for all $\lambda \in [0, 1]$, the lowest disbenefit for the consumer is obtained by choosing $t_c = 1$.

However, since $t = \min\{t_g, t_c\}$, the consumer has to accept that the system operator will use $t = 0$ from the generator's offer.

4.4. Payoffs

Finally, we explore the trade in Arrow-Debreu securities that emerges in equilibrium. It is convenient to express these in terms of payoffs $\Pi_a(\omega)$ rather than disbenefits. If the system operator dispatches assuming $t = 0$, then the payoffs to agents (generator and consumer) are:

$$\begin{aligned} \Pi_g(g) &= 5500 \\ \Pi_g(h) &= 4700 \end{aligned}$$

$$\begin{aligned} \Pi_c(g) &= 7200 \\ \Pi_c(h) &= 12800 \end{aligned}$$

Given these payoffs, consider the following trades in Arrow-Debreu securities. The generator pays $\frac{1}{2}(1 + \lambda)(-400) + \frac{1}{2}(1 - \lambda)(400) = -400\lambda$, *i.e.*, he sells 400 Arrow-Debreu securities at price $\frac{1}{2}(1+\lambda)$ giving payoff -1 when $\omega = g$ and buys 400 Arrow-Debreu securities at price $\frac{1}{2}(1-\lambda)$ paying off 1 when $\omega = h$.

The consumer pays $\frac{1}{2}(1+\lambda)(400) + \frac{1}{2}(1-\lambda)(-400)$, *i.e.*, she buys 400 Arrow-Debreu securities at price $\frac{1}{2}(1+\lambda)$ paying off 1 when $\omega = g$ and sells 400 Arrow-Debreu securities at price $\frac{1}{2}(1-\lambda)$ paying off 1 when $\omega = h$.

This example illustrates how it is optimal for the generator to offer $t_g = 0$ to the system operator, even if $\lambda_g > 0$. They then evaluate their payoff using $\min\{\lambda_a\}$. In general, this assumes that they know the risk sets of the other agents. If not, agent a might optimize the choice of t_g, assuming a probability distribution for λ_j, $j \neq a$.

5. Conclusions

In this paper, we describe a pricing mechanism that guarantees both revenue adequacy in all scenarios and cost recovery in risk-adjusted expected cost. The application of this mechanism assumes the existence of a complete market for trading risk using Arrow-Debreu securities, which will not be the case in practice. Nevertheless, we expect that markets with deep contract markets might exhibit some of the behavior we describe. Unfortunately, our mechanism does not ensure the truthful revelation of risk sensitivity by the agents. Further, we showed by example that some agents might have incentives to under-declare their risk aversion, while others have incentives to exaggerate their risk aversion. It is interesting to speculate on how market regulators might detect this misrepresentation. For example, one could use observed contract trades to estimate the risk profile of the agents or at least statistically check consistency between risk declaration and actual trades.

Acknowledgments.

This research was performed while the authors were participating in the Architecture of Green Energy Systems Program hosted by the Institute for Mathematical and Statistical Innovation (IMSI), which is supported by the National Science Foundation (Grant No. DMS-1929348). Andy Philpott acknowledges support from UOCX2117 MBIE Catalyst Fund New Zealand German Platform for Green Hydrogen Integration (HINT).

References

[1] F. Bouffard, F. Galiana, A. Conejo, Market-clearing with stochastic security-part i: formulation, IEEE Transactions on Power Systems 20 (4) (2005) 1818–1826.

[2] S. Wong, J. Fuller, Pricing energy and reserves using stochastic optimization in an alternative electricity market, IEEE Transactions on Power Systems 22 (2) (2007) 631–638.

[3] G. Pritchard, G. Zakeri, A. Philpott, A single-settlement, energy-only electric power market for unpredictable and intermittent participants, Operations Research 58 (4-part-2) (2010) 1210–1219.

[4] R. Myerson, M. Satterthwaite, Efficient mechanisms for bilateral trading, Journal of Economic Theory 29 (2) (1983) 265–281.

[5] G. Zakeri, G. Pritchard, M. Bjorndal, E. Bjorndal, Pricing wind: a revenue adequate, cost recovering uniform price auction for electricity markets with intermittent generation, INFORMS Journal on Optimization 1 (1) (2019) 35–48.

[6] R. Cory-Wright, A. Philpott, G. Zakeri, Payment mechanisms for electricity markets with uncertain supply, Operations Research Letters 46 (1) (2018) 116–121.

[7] D. Ralph, Y. Smeers, Risk trading and endogenous probabilities in investment equilibria, SIAM Journal on Optimization 25 (4) (2015) 2589–2611.

[8] M. Ferris, A. Philpott, Dynamic risked equilibrium, Operations Research 70 (3) (2022) 1933–1952.

[9] H. Gérard, V. Leclère, A. Philpott, On risk averse competitive equilibrium, Operations Research Letters 46 (1) (2018) 19–26.

[10] R. Cory-Wright, G. Zakeri, On stochastic auctions in risk-averse electricity markets with uncertain supply, Operations Research Letters 48 (3) (2020) 376–384.

[11] A. Philpott, G. Pritchard, Financial transmission rights in convex pool markets, Operations Research Letters 32 (2) (2004) 109–113.