Risk-Aware Market Clearing in Electricity Markets

HRITHIK NAMBIAR

CONTENTS

Acknowledgement

I would like to thank my advisor, Professor Amy Greenwald, for her guidance, support and encouragement throughout my time at Brown University. I am glad that I discovered my interest in the area of Algorithmic Game Theory through Professor Greenwald's course and through discussions with her and other friends in the E-GLAMOR research group.

I would also like to thank my collaborator, Wirattawut (Max) Boonbandansook, for his friendship, invaluable discussions and contribution to this thesis. Max helped immensely with all the experimental results that are presented in this work. Thank you to all my other collaborators, Gianluca Brero, Ben Lubin and Ryan Edmonds, for the stimulating discussions and ideas presented in this work.

To all my friends at Brown University, thank you for being my family away from home. These last two years have been some of the best days of my life.

Finally, my eternal gratitude to my family for always being there for me and for their constant support and encouragement. This thesis is dedicated to you.

1 Introduction

In this work, we study the design of a centralised market under uncertainty. Our focus is on two-stage markets, where a decision in the first stage is made based on an uncertain second stage. Electricity markets serve as a canonical example of such a two-stage market. These markets, although structurally different in the United States and Europe, operate in two stages: the day-ahead and the real-time market. What makes the design of this market challenging and consequently interesting to study is that, unlike other markets, electricity is a commodity that cannot be economically stored at scale; it must be delivered over a physical grid network that is governed by Kirchhoff's laws. Consequently, supply must equal demand at every node in the electricity grid. The two-stage electricity market is a direct consequence of these features. In the Day-Ahead (DA) market, generators and consumers submit bids and offers, and the Independent System Operator (ISO) clears the market to publish allocations (tentative commitment and dispatch schedule) and (day-ahead) prices. The bids submitted are based on each individual market agent's forecast of load and uncertain generation sources (such as renewable energy sources). Closer to the time of delivery, the real-time (RT) market clears any imbalances that arise from forecast errors, generator outages, or contingencies.

Among many considerations which make market design in this environment challenging, the most central to this thesis is the temporal separation between commitment and delivery. This means that any DA decision is made under a probability distribution over future states rather than under perfect information, and the welfare-optimal benchmark itself becomes a stochastic optimization problem rather than a deterministic one. This raises an important question: how can we generalize existing market designs to account for the risk preferences of the agents in the market?

These challenges have been amplified by the rapid penetration of variable renewable resources. Wind and solar generators now supply a substantial and growing share of energy in most major systems, and their output is intrinsically stochastic, as it is weather-dependent. As a result, the forecast errors that the RT market must absorb have grown both in magnitude and in tail risk. Designing a market that performs well in this regime is not merely a refinement of the classical problem; it requires an explicitly stochastic, multi-stage and importantly risk-aware problem formulation in which the market's allocation rule, payment rule, and equilibrium concept are all defined relative to risk-aware market agents; that is, agents, hedging against uncertainty in the market.

Main Contributions. This paper develops a competitive-equilibrium theory of two-stage electricity markets in which agents have heterogeneous risk preferences. We use Conditional-Value-at-Risk (CVaR) as a risk measure. We work within the framework of risk-averse equilibrium with Arrow–Debreu securities developed by Philpott et al.[11] and Gerard et al. [6], and specialise their results to the institutionally relevant setting of day-ahead and real-time markets with multiple consumers and producers, mean–CVaR preferences, and individual recourse penalties. Our main contributions are: (i) we give closed-form expressions for the equilibrium day-ahead price, real-time price and Arrow–Debreu state-contingent price (ii) establish First and Second Welfare Theorems for the risk-aware market (iii) we characterise decentralisation under heterogeneous risk preferences (iv) we show emperical result on the advantages of risk-aware clearing on data collected from western intercontinental grid in the United States.

2 Electricity Market Design (US vs EU)

Electricity market design differs considerably between the United States and Europe, particularly in the degree of centralisation, the treatment of network constraints, and the division of responsibilities between the market operators and the system operators.

In the United States, Independent System Operators (ISOs) or Regional Transmission Organisations (RTOs) are responsible for both market clearing and system operations. They centrally run both day-ahead and real-time markets. To do this, they solve two interconnected optimisation problems. First, they solve a unit commitment problem, referred to as the Security Constrained Unit Commitment (SCUD), which determines which generators should be online in advance, accounting for nonconvex operational constraints such as ramp-up rates, minimum up and down times, startup costs, and planned outages. This is solved as a Mixed-Integer Programming problem. Second, they solve an economic dispatch problem that determines how much each committed generator should produce to meet demand at minimum cost. This dispatch problem is fundamentally a market-clearing problem. It matches supply and demand, subject to network feasibility. Generators submit supply offers, typically in the form of stepwise price-quantity bids, and demand is represented through ISO demand forecasts informed by various Load-Serving Entities (LSEs). In most US formulations, demand is treated as inelastic, meaning it is not modeled as responding continuously to market prices. As a result, the day-ahead ISO objective is typically written as cost minimisation subject to market-clearing and transmission constraints. In the real-time market, SCED is solved as adjusting to realised system conditions in short intervals, often every 5 to 15 minutes. The US market-clearing model is also distinguished by its use of nodal pricing. A node is a location in the transmission system, typically a bus or an aggregate injection or withdrawal point, at which power can be injected by generators, withdrawn by loads, or transferred through transmission lines. The transmission network is represented explicitly through nodes and lines, and power balance is ensured at each of these nodes. The market enforces power balance at individual nodes and computes Locational Marginal Prices (LMPs) that vary across locations.

The European market design is organised differently. The day-ahead market is a single pan-European cross-zonal day-ahead electricity clearing problem that allocates cross-border transmission capacity through a common algorithm that maximises social welfare. The output of this process includes clearing prices, matched trades, scheduled exchanges, and the net positions of bidding areas. A bidding area is the largest geographical area within which bids and offers can be matched without attributing cross-zonal capacity, and the bidding zones in Europe are mostly defined by national borders. Consequently, European short-term pricing is predominantly zonal rather than nodal. After the day-ahead auction, market participants may continue to trade in intraday markets, allowing schedules to adjust as new information arrives before delivery. This trading can continue close to delivery time and is beneficial because it allows participants to rebalance positions after the day-ahead market has closed. Physical feasibility and security are then supported by Transmission System Operators (TSOs) through redispatching and balancing actions.

2.1 Motivation for Market Design

We acknowledge that each of these market designs has its own merits. The U.S. approach yields strong locational signals and integrates network feasibility directly into market clearing through nodal balance and transmission constraints. The European approach, by contrast, is naturally framed as a welfare-maximising exchange and supports broad cross-border participation through a simpler zonal trading architecture.

In this work, our objective is to combine the strongest elements of each paradigm. We therefore study a welfare-maximising electricity market with nodal power-balance constraints. This retains the welfare-based perspective that is natural in exchange design, while also preserving the nodal network representation needed to capture congestion and locational scarcity. In this sense, our framework can be interpreted as a general benchmark that combines the economic surplus logic emphasised in European market coupled with the nodal feasibility and pricing discipline characteristic of U.S. markets.

3 Literature Review

The classical electricity market in the United States is cleared by a deterministic economic dispatch whose dual variables yield locational marginal prices (LMPs) supporting a competitive equilibrium. The growing penetration of variable renewables has prompted a sequence of generalisations of this design that internalise uncertainty within the clearing problem itself. In an early work, Wong and Fuller [17] proposed a two-stage stochastic linear program for joint energy and reserve dispatch under contingency uncertainty. Pritchard et al. [12] introduced a stochastic clearing approach with a single-settlement energy-only market. Zavala et al. [18] extended this work by focusing on pricing consistency; deterministic clearing can create arbitrary distortions between day-ahead prices and expected real-time prices, while the stochastic formulation bounds these distortions and can eliminate them under appropriate penalty structures. Inspired by the aforementioned work, we formulate market clearing as a two-stage stochastic program with first-stage as day-ahead commitments and second-stage as real-time recourse. However, we differ by applying Conditional-Value-at-Risk (CVaR) as a risk measure to the distribution of welfare losses.

Recent pricing work also studies uncertainty in multi-period dispatch. Guo et al. [7] analyse pricing multi-interval dispatch under uncertainty, focusing on dispatch-following incentives for generators in settings with intertemporal operational constraints. Werner et al. [16] also study pricing uncertainty in stochastic multi-stage electricity markets, developing prices for uncertainty across a multi-stage setting rather than only for deterministic energy quantities. Extending our framework to multi-stage settings, integrating the explicit price identifications and heterogeneity results developed here with the multi-stage pricing structures of is a natural direction for future work.

Another branch of relevant literature studies risk-sensitive dispatch models. Dvorkin et al. [5] and Meith et al. [10] developed a chance-constrained stochastic market in which uncertainty is internalised through chance constraints on generator and network limits, with energy and reserve prices recovered from the dual of the resulting program. They also introduce Arrow–Debreu securities into a chance-constrained stochastic electricity market to enable risk trading. Madavan et al. [9] propose a risk-sensitive security-constrained economic dispatch model in which resilience to line failures is represented using CVaR. The main difference between their setting and ours is that theirs is security-constrained dispatch with contingency risk, whereas ours is a two-stage stochastic market-clearing model in which the risk measure is applied to social welfare loss. Tapia et al. [15] propose market clearing for extreme events using large-deviation-theory chance constraints and weighted chance constraints. This paper is relevant because it treats extreme events as a distinct object of market design rather than as ordinary low-probability scenarios. Our approach differs by using risk measures to implicitly tilt the welfare clearing objective to prioritise the tail events.

The theoretical foundation for risk trading goes back to Arrow's theory of securities [3] and optimal risk-bearing. Arrow showed that state-contingent securities can support efficient allocation of risk under uncertainty, establishing the conceptual basis for completing markets through claims contingent on realised states. Gerard et al. [6] study risk-averse competitive partial equilibrium in stochastic electricity markets where agents are endowed with coherent risk measures. They distinguish risk-neutral equilibrium, risk-averse equilibrium, and risk-trading equilibrium with

Arrow–Debreu securities. Their main message is directly relevant to our setting: with coherent risk measures, a complete risk market can support an efficient risk-adjusted social optimum, but when the risk market is incomplete, risk-averse competitive equilibria need not be unique and may fail to coincide with the social planner solution. Our work specialises this abstract framework to a two-stage day-ahead and real-time electricity market with multiple agents and mean–CVaR preferences, deriving explicit price formulas and a constructive transfer rule. We furrther develop the heterogeneous-risk-preference case in detail, showing that more risk-averse agents are endogenously fully insured while less risk-averse agents bear all residual aggregate risk.

4 Risk Measures

A risk measure assigns a real number to a random variable representing loss, summarising in a single value the riskiness of holding that random loss. In our setting, the random loss in question is the negative of second-stage welfare, which depends on the realised scenario. Different choices of risk measure correspond to different attitudes toward uncertainty, and the structural properties of the risk measure determine whether it can be decentralised through prices.

This section reviews the framework of coherent risk measures [4], introduces Conditional Value-at-Risk (CVaR) as a coherent risk measure, and motivates the convex combination of expectation and CVaR that we adopt in Section 8.1.

4.1 Coherent Risk Measures

Let Ω be a finite scenario set with probability measure P assigning $p_\omega > 0$ to each ω. Let $\mathcal{X}$ denote the space of real-valued random variables on Ω, interpreted as losses (larger values are worse).

Definition 1 (Coherent Risk Measure, 4). *A function $\rho : \mathcal{X} \to \mathbb{R}$ is a* coherent risk measure *if it satisfies the following four axioms for all $X, Y \in \mathcal{X}$ and $a \in \mathbb{R}$, $\alpha \geq 0$:*

- **(M)** Monotonicity*: if $X \leq Y$ pointwise, then $\rho(X) \leq \rho(Y)$.*
- **(T)** Translation equivariance*: $\rho(X + a) = \rho(X) + a$.*
- **(P)** Positive homogeneity*: $\rho(\alpha X) = \alpha\, \rho(X)$.*
- **(S)** Subadditivity*: $\rho(X + Y) \leq \rho(X) + \rho(Y)$.*

Coherent risk measures admit a fundamental dual characterisation.

Lemma 1 (Dual Representation, 4). *A function $\rho : \mathcal{X} \to \mathbb{R}$ is a coherent risk measure if and only if there exists a closed convex set $\mathcal{Q}$ of probability measures on Ω, each absolutely continuous with respect to P, such that*

$$\rho(X) \;=\; \sup_{Q \in \mathcal{Q}} \mathbb{E}_Q[X], \qquad \forall X \in \mathcal{X}. \tag{1}$$

The set $\mathcal{Q}$ is called the risk set *of ρ, and any $Q \in \mathcal{Q}$ attaining the supremum at X is called a* worst-case measure *for X.*

A useful property is that convex combinations of coherent risk measures are coherent.

Lemma 2 (Convex combinations preserve coherence). *Let ρ_1, ρ_2 be coherent risk measures with risk sets $\mathcal{Q}_1, \mathcal{Q}_2$. For any $\alpha \in [0, 1]$, the function $\rho := \alpha\, \rho_1 + (1 - \alpha)\, \rho_2$ is a coherent risk measure with risk set*

$$\mathcal{Q} \;=\; \{\alpha Q_1 + (1 - \alpha) Q_2 : Q_1 \in \mathcal{Q}_1,\ Q_2 \in \mathcal{Q}_2\}.$$

Proof. Each axiom (M), (T), (P), (S) is preserved under convex combinations: (M), (T), (P) are linear in ρ, and (S) is convex (a convex combination of subadditive inequalities is subadditive). The risk set follows from $\sup_{Q_1,Q_2}(\alpha\, \mathbb{E}_{Q_1}[X] + (1 - \alpha)\, \mathbb{E}_{Q_2}[X]) = \sup_{Q_1,Q_2} \mathbb{E}_{\alpha Q_1 + (1-\alpha) Q_2}[X]$. □

4.2 Conditional Value-at-Risk

Let $f(x, y)$ be loss associated with $x \in X \subseteq \mathbb{R}^n$, $y \in \mathbb{R}^m$. For example, x is the social welfare, y is the uncertainties involved, for example, the market parameters. Let,

$$p(y) = \mathbb{P}(Y \leq y)$$

Let's define the probability of $f(x, y)$ not exceeding α as,

$$\Psi(x, \alpha) = \int_{\{f(x,y) \leq \alpha\}} p(y)\, dy = \mathbb{P}(f(x, y) \leq \alpha)$$

Assumption: The probability distribution is such that no jumps occur, i.e., $\Psi(x, \alpha)$ is everywhere continuous w.r.t. α. Then,

$$\beta\text{-VaR is defined as} : \alpha_\beta(x) = \min\{\alpha \in \mathbb{R} : \Psi(x, \alpha) \geq \beta\}$$

$$\beta\text{-CVaR is defined as} : \Phi_\beta(x) = \frac{1}{1-\beta} \int_{\{f(x,y) \geq \alpha_\beta(x)\}} f(x, y)\rho(y)\, dy$$

The important work by [13] is the characterization of $\Phi_\beta(x)$ and $\alpha_\beta(x)$ in terms of a new function,

$$\Rightarrow F_\beta(x, \alpha) = \alpha + \frac{1}{1-\beta} \int_{\mathbb{R}^m} [f(x, y) - \alpha]^+ \rho(y)\, dy$$

Theorem 01: As a function of α, $F_\beta(x, \alpha)$ is convex & continuously differentiable.
β-CVaR of $f(x, y)$ is $\Phi_\beta(x) = \min_{\alpha \in \mathbb{R}} F_\beta(x, \alpha)$, and β-VaR is $\alpha_\beta(x) = \inf \arg\min_{\alpha \in \mathbb{R}} F_\beta(x, \alpha)$

Proof of Theorem 1:

$$F_\beta(x, \alpha) = \alpha + \frac{1}{1-\beta} \int_{\mathbb{R}^m} [f(x, y) - \alpha]^+ \rho(y)\, dy$$

We assume that $\Psi(x, \alpha)$ is continuous w.r.t. α, i.e., regardless of the choice of x, the set of y with $f(x, y) = \alpha$ has probability 0, i.e.,

$$\int_{\{f(x,y)=\alpha\}} \rho(y)\, dy = 0$$

Lemma 01: With fixed x, let

$$G(\alpha) = \int_{\mathbb{R}^m} g(\alpha, y)\rho(y)\, dy \text{ where } g(\alpha, y) = [f(x, y) - \alpha]^+$$

Then G is a convex, continuously differentiable function, with

$$G'(\alpha) = \Psi(x, \alpha) - 1$$

[14]

Proof of Lemma 01: Consider a function $F(\theta)$ that has the form

$$F(\theta) = \mathbb{E}\{f(\theta)\} = \int_\Omega f(\theta, \omega)P(d\omega)$$

where $f(\theta)$ is a measurable random function (whose realization is denoted by $f(\theta, \omega)$ as in the above integral) defined on a probability space $(\Omega, \mathcal{F}, P)$. Let the domain of f be a convex open set $D \subset \mathbb{R}^m$ and suppose that, for every $\theta \in D$, $\mathbb{E}\{|f(\theta)|\} < \infty$. We now give sufficient conditions for directional differentiability of F at a point $\theta_0 \in D$.

Assumption 2.1: There exists a positive-valued random variable $K = K(\omega)$ such that $\mathbb{E}\{K\}$ is finite and

$$|f(\theta_1, \omega) - f(\theta_2, \omega)| \leq K(\omega)\|\theta_1 - \theta_2\| \tag{2.2}$$

for almost all $\omega \in \Omega$ and for all $\theta_1, \theta_2 \in D$.

Assumption 2.2: With probability one (w.p.1) the function $f(\theta)$ is directionally differentiable at θ_0.

Proposition 2.1: Suppose that either i) Assumptions 2.1 and 2.2 hold, or ii) the function $f(\theta)$ is convex with probability 1; Then the expected-value function $F(\theta)$ is directionally differentiable at $\theta_0 \in D$ and

$$F'(\theta_0, d) = \mathbb{E}\{f'(\theta_0, d)\}. \tag{2.3}$$

Proof: Under Assumptions 2.1 and 2.2, formula (2.3) follows easily from the Lebesgue Dominated Convergence Theorem. In the convex case, formula (2.3) is implied by the Monotone Convergence Theorem. Indeed, consider $\theta_0 \in D, d \in \mathbb{R}^m$, and a monotone-decreasing sequence $t_n \to 0^+$. It follows then from convexity of $f(\cdot, \omega)$, that the sequence $\psi_n = \psi_n(\omega) = t_n^{-1}[f(\theta_0 + t_n d, \omega) - f(\theta_0, \omega)]$ is monotone decreasing for almost every $\omega \in \Omega$, and that $\psi_n \to f'(\theta_0, d)$ w.p.1. Therefore by the Monotone Convergence Theorem

$$\lim_{n\to\infty} \mathbb{E}\{\psi_n\} = \mathbb{E}\left\{\lim_{n\to\infty} \psi_n\right\} = \mathbb{E}\{f'(\theta_0, d)\}.$$

Since the expected-value function $F(\theta) = \mathbb{E}\{f(\theta)\}$ is then convex, we have that the limit $\lim_{n\to\infty} \mathbb{E}\{\psi_n\}$ is finite and is equal to $F'(\theta_0, d)$. This shows that, in the convex case, if $\mathbb{E}\{|f(\theta)|\} < \infty$ for all $\theta \in D$, then the interchangeability formula (2.3) holds.

Considering,

$$G(\alpha) = \mathbb{E}\{g(\alpha, y)\} \implies G'(\alpha) = \mathbb{E}\{g'(\alpha, y)\} = \Psi(x, \alpha) - 1$$

completes the proof of Lemma 01.

Proof of Theorem 1 (Continued): With $F_\beta(\mathbf{x}, \alpha)$ as defined above, it is immediate from the Lemma 01 that $F_\beta(\mathbf{x}, \alpha)$ is convex and continuously differentiable with derivative

$$\frac{\partial}{\partial\alpha} F_\beta(\mathbf{x}, \alpha) = 1 + (1-\beta)^{-1}[\Psi(\mathbf{x}, \alpha) - 1] = (1-\beta)^{-1}[\Psi(\mathbf{x}, \alpha) - \beta].$$

Therefore, the values of α that furnish the minimum of $F_\beta(\mathbf{x}, \alpha)$, i.e., the ones comprising the set $A_\beta(\mathbf{x}) = \operatorname{argmin} F_\beta$, are precisely those for which $\Psi(\mathbf{x}, \alpha) - \beta = 0$. They form a nonempty closed interval, inasmuch as $\Psi(\mathbf{x}, \alpha)$ is continuous and nondecreasing in α with limit 1 as $\alpha \to \infty$ and limit 0 as $\alpha \to -\infty$. This further yields the validity of the β-VaR formula. In particular, then, we have

$$\min_{\alpha\in\mathbb{R}} F_\beta(\mathbf{x}, \alpha) = F_\beta(\mathbf{x}, \alpha_\beta(\mathbf{x})) = \alpha_\beta(\mathbf{x}) + (1-\beta)^{-1} \int_{y\in\mathbb{R}^m} [f(\mathbf{x}, \mathbf{y}) - \alpha_\beta(\mathbf{x})]^+ p(\mathbf{y})\, d\mathbf{y}.$$

But the integral here equals

$$\int_{f(\mathbf{x},\mathbf{y})\geq\alpha_\beta(\mathbf{x})} [f(\mathbf{x}, \mathbf{y}) - \alpha_\beta(\mathbf{x})] p(\mathbf{y})\, d\mathbf{y} = \int_{f(\mathbf{x},\mathbf{y})\geq\alpha_\beta(\mathbf{x})} f(\mathbf{x}, \mathbf{y}) p(\mathbf{y})\, d\mathbf{y} - \alpha_\beta(\mathbf{x}) \int_{f(\mathbf{x},\mathbf{y})\geq\alpha_\beta(\mathbf{x})} p(\mathbf{y})\, d\mathbf{y},$$

where the first integral on the right is by definition $(1-\beta)\phi_\beta(\mathbf{x})$ and the second is $1 - \Psi(\mathbf{x}, \alpha_\beta(\mathbf{x}))$. Moreover $\Psi(\mathbf{x}, \alpha_\beta(\mathbf{x})) = \beta$. Thus,

$$\min_{\alpha\in\mathbb{R}} F_\beta(\mathbf{x}, \alpha) = \alpha_\beta(\mathbf{x}) + (1-\beta)^{-1}[(1-\beta)\phi_\beta(\mathbf{x}) - \alpha_\beta(\mathbf{x})(1-\beta)] = \phi_\beta(\mathbf{x}).$$

This confirms for β-CVaR formula and finishes the proof of Theorem 1. $\diamond$

Theorem 2. Minimizing the β-CVaR of the loss associated with $\mathbf{x}$ over all $\mathbf{x} \in X$ is equivalent to minimizing $F_\beta(\mathbf{x}, \alpha)$ over all $(\mathbf{x}, \alpha) \in X \times \mathbb{R}$, in the sense that

$$\min_{\mathbf{x} \in X} \phi_\beta(\mathbf{x}) = \min_{(\mathbf{x},\alpha) \in X \times \mathbb{R}} F_\beta(\mathbf{x}, \alpha),$$

where moreover a pair $(\mathbf{x}^*, \alpha^*)$ achieves the second minimum if and only if $\mathbf{x}^*$ achieves the first minimum and $\alpha^* \in A_\beta(\mathbf{x}^*)$. In particular, therefore, in circumstances where the interval $A_\beta(\mathbf{x}^*)$ reduces to a single point (as is typical), the minimization of $F_\beta(\mathbf{x}, \alpha)$ over $(\mathbf{x}, \alpha) \in X \times \mathbb{R}$ produces a pair $(\mathbf{x}^*, \alpha^*)$, not necessarily unique, such that $\mathbf{x}^*$ minimizes the β-CVaR and α^* gives the corresponding β-VaR.

Furthermore, $F_\beta(\mathbf{x}, \alpha)$ is convex with respect to $(\mathbf{x}, \alpha)$, and $\phi_\beta(\mathbf{x})$ is convex with respect to $\mathbf{x}$, when $f(\mathbf{x}, \mathbf{y})$ is convex with respect to $\mathbf{x}$, in which case, if the constraints are such that X is a convex set, the joint minimization is an instance of convex programming.

It should be noted that the integral in the definition of $F_\beta(\mathbf{x}, \alpha)$ can be approximated in various ways. For example, this can be done by sampling the probability distribution of y according to its density $p(\mathbf{y})$. If the sampling generates a collection of vectors $\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_q$, then the corresponding approximation to $F_\beta(\mathbf{x}, \alpha)$ is $\tilde{F}_\beta(\mathbf{x}, \alpha) = \alpha + \frac{1}{q(1-\beta)} \sum_{k=1}^{q} [f(\mathbf{x}, \mathbf{y}_k) - \alpha]^+$.

The expression $\tilde{F}_\beta(\mathbf{x}, \alpha)$ is convex and piecewise linear with respect to α. Although it is not differentiable with respect to α, it can readily be minimised, either by line search techniques or by representation in terms of an elementary linear programming problem.

CVaR is a coherent risk measure for any confidence level $\beta \in [0, 1)$. It hence has a dual representation.

Lemma 3 (Dual representation of CVaR, [2]). *For any $\beta \in [0, 1)$,*

$$CVaR_\beta[X] = \sup_{Q \in \mathcal{Q}_\beta} \mathbb{E}_Q[X], \qquad \mathcal{Q}_\beta = \left\{ Q \ll P : \frac{dQ}{dP}(\omega) \le \frac{1}{1-\beta} \ \forall \omega \right\}. \tag{2}$$

The supremum is attained, and a worst-case measure Q^* places mass $\frac{1}{1-\beta}\rho_\omega$ on each scenario in the upper β-tail of X.

We will use the fact that γ_ω^*, the multiplier on the CVaR shortfall constraint in the epigraph reformulation (62), identifies the worst-case measure: the optimal multipliers from the planner's KKT system encode precisely the probability distortion induced by tail-aversion.

5 Stochastic Day Ahead Clearing

In this section, we formulate the stochastic day-ahead clearing problem as a two-stage programme. In the first stage, tentative allocations are declared before uncertainty is realised. In the second stage, recourse actions adjust these allocations once the uncertain parameters have been realised. We present the social welfare maximisation problem solved by a central market planner (the *optimisation problem*), followed by the individual utility maximisation problems solved independently by each market participant (the *equilibrium problem*).

5.1 The Optimisation Problem

Consider a two-stage market with n consumers and m producers, indexed by i and j respectively, who wish to commit to trading quantities of a single divisible good in the first stage, for which their values and costs, and hence their optimal quantities, will be realised in the second stage. The uncertainty in this environment is modelled by a discrete random variable with outcome space Ω (e.g., in the case of solar energy, the degree of cloud cover), whose values $\omega \in \Omega$, so-called *scenarios*, are drawn from the probability distribution $\mathcal{P}$ and realised with probability ρ_ω. These scenarios parameterise the consumers' values and the producers' costs, so that each consumer i (resp. producer j) realises a per-scenario value $v_{i\omega}$ (resp. cost $c_{j\omega}$), along with per-scenario lower and upper bounds on consumption $\underline{x}_{i\omega}$ and $\overline{x}_{i\omega}$ (resp. production $\underline{y}_{j\omega}$ and $\overline{y}_{j\omega}$).

In the first stage of the market, before any uncertainty is resolved, a pair of tentative allocations $(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{R}^n \times \mathbb{R}^m$ for the consumers and producers, respectively, is declared. In the second stage, after the scenario has been realised, recourse actions are taken via positive and negative deviations $\boldsymbol{X} = (\boldsymbol{X}^+, \boldsymbol{X}^-)$ with $\boldsymbol{X}^+, \boldsymbol{X}^- \in \mathbb{R}^{n \times |\Omega|}$ and $\boldsymbol{Y} = (\boldsymbol{Y}^+, \boldsymbol{Y}^-)$ with $\boldsymbol{Y}^+, \boldsymbol{Y}^- \in \mathbb{R}^{m \times |\Omega|}$ from the tentative allocation, that is, agents can increase or decrease their consumption or production relative to the tentative allocations. Each consumer i (resp. producer j) then suffers an exogenous marginal penalty α_i^+ (resp. α_j^+) for positive deviations and an exogenous marginal penalty α_i^- (resp. α_j^-) for negative deviations, with $\alpha_k^+, \alpha_k^- \geq 0$ for all $k \in \{i, j\}$. This asymmetry allows the market to penalise undersupply or overdemand more heavily than the opposite.

We model producers and consumers with *quasi-linear utilities*, meaning consumers subtract their expenses (the price of their allocations) from their consumption values, while producers subtract their production costs from their revenue. Let $p_0 \in \mathbb{R}$ be the posted price in the first stage, and let $\boldsymbol{p} \in \mathbb{R}^{|\Omega|}$ so that $p_\omega \in \mathbb{R}$ is the posted price in the second stage under scenario ω. Agent i's (quasi-linear) utility under scenario ω is:

$$
\begin{aligned}
u_i(x_i, \boldsymbol{X}_i, p_0, \boldsymbol{p}; \omega) = {} & (v_{i\omega} - p_0)x_i + (v_{i\omega} - p_\omega)(x_{i\omega}^+ - x_{i\omega}^-) \\
& - (\alpha_i^+ \, x_{i\omega}^+ + \alpha_i^- \, x_{i\omega}^-)
\end{aligned} \tag{3}
$$

Similarly, producer j's utility is:

$$
\begin{aligned}
u_j(y_j, \boldsymbol{Y}_j, p_0, \boldsymbol{p}; \omega) = {} & (p_0 - c_{j\omega})y_j + (p_\omega - c_{j\omega})(y_{j\omega}^+ - y_{j\omega}^-) \\
& - (\alpha_j^+ \, y_{j\omega}^+ + \alpha_j^- \, y_{j\omega}^-)
\end{aligned} \tag{4}
$$

For both producers and consumers, the first term represents day-ahead trades, the second represents real-time trades, and the third represents penalties for deviation from the tentative allocation.

We have thus defined a *two-stage stochastic market* $\mathcal{M} = (m, n, \Omega, \mathcal{P})$ for a single divisible good, where agents' values and costs are non-negative real values, and their utilities are quasi-linear.

5.1.1 Second-Stage Welfare Problem. Given such a stochastic market $\mathcal{M}$, together with a tentative first-stage allocation $(\boldsymbol{x}, \boldsymbol{y})$, the welfare maximisation problem $\mathrm{W2P}(\boldsymbol{x}, \boldsymbol{y}; \omega)$ in the second stage assuming scenario $\omega \in \Omega$ is:

$$
\begin{aligned}
\mathrm{W2P}(\boldsymbol{x}, \boldsymbol{y}; \omega) \ := \ \arg\max_{\boldsymbol{X}, \boldsymbol{Y}} \quad & \sum_{i \in [n]} \Big[v_{i\omega}(x_i + x_{i\omega}^+ - x_{i\omega}^-) - (\alpha_i^+ \, x_{i\omega}^+ + \alpha_i^- \, x_{i\omega}^-) \Big] \\
& - \sum_{j \in [m]} \Big[c_{j\omega}(y_j + y_{j\omega}^+ - y_{j\omega}^-) + (\alpha_j^+ \, y_{j\omega}^+ + \alpha_j^- \, y_{j\omega}^-) \Big]
\end{aligned} \tag{5}
$$

subject to

$$
\sum_{i \in [n]} (x_{i\omega}^+ - x_{i\omega}^-) = \sum_{j \in [m]} (y_{j\omega}^+ - y_{j\omega}^-) \tag{6}
$$

$$
x_i + x_{i\omega}^+ - x_{i\omega}^- \in [\underline{x}_{i\omega}, \ \overline{x}_{i\omega}] \quad \forall \, i \in [n] \tag{7}
$$

$$
y_j + y_{j\omega}^+ - y_{j\omega}^- \in [\underline{y}_{j\omega}, \ \overline{y}_{j\omega}] \quad \forall \, j \in [m] \tag{8}
$$

$$
\boldsymbol{X}, \boldsymbol{Y} \geq 0 \tag{9}
$$

Constraint (6) ensures supply–demand balance in the second stage. Constraints (7)–(8) ensure feasibility of the second-stage allocations. Constraint (9) precludes negative adjustments; decreasing production, for example, requires a positive $y^-_{j\omega}$ rather than a negative $y^+_{j\omega}$.

Remark 1. *Our formulation enforces the equality constraint 6. In practice, it may be desirable to permit over-supply, for instance, to accommodate energy storage via battery systems. This can be achieved within our framework by introducing auxiliary slack variables representing stored or curtailed energy, thereby relaxing the balance to an inequality while preserving the dual structure from which equilibrium prices are derived.*

We write $\mathrm{W2}(\boldsymbol{x}, \boldsymbol{y}, X_\omega, Y_\omega; \omega)$ for the welfare function (the objective evaluated at the given arguments), and $\mathrm{W2}(\boldsymbol{x}, \boldsymbol{y}; \omega)$ for the optimal welfare in scenario ω.

5.1.2 First-Stage Welfare Problem. With the per-scenario second-stage problem thus defined, the first-stage welfare maximisation problem $\mathrm{W1P}(\mathcal{P})$ is to maximise the expected welfare in the second stage:

$$\mathrm{W1P}(\mathcal{P}) \;:=\; \arg\max_{\boldsymbol{x},\boldsymbol{y}} \quad \mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[\mathrm{W2}(\boldsymbol{x}, \boldsymbol{y}; \omega)\right] \tag{10}$$

$$\text{subject to} \quad \sum_{i\in[n]} x_i = \sum_{j\in[m]} y_j \tag{11}$$

$$\boldsymbol{x}, \boldsymbol{y} \geq 0 \tag{12}$$

Constraint (11) ensures feasibility of the first-stage allocation, while constraint (12) ensures non-negative allocations. We write $\mathrm{W1}(\boldsymbol{x}, \boldsymbol{y}; \mathcal{P})$ for the objective of $\mathrm{W1P}(\mathcal{P})$ evaluated at the given arguments, and $\mathrm{W1}(\mathcal{P})$ for the optimal welfare.

Observation 1. *There exists an optimal solution to W1P($\mathcal{P}$) in which, for any scenario ω, at most one of $x^+_{i\omega}$ and $x^-_{i\omega}$ is nonzero for all consumers i, and similarly for producers.*

Proof. Suppose for contradiction that $x^+_{i\omega} > 0$ and $x^-_{i\omega} > 0$ for some consumer in some scenario. Let $\delta = \min(x^+_{i\omega}, x^-_{i\omega})$. Replacing $x^+_{i\omega}$ by $x^+_{i\omega} - \delta$ and $x^-_{i\omega}$ by $x^-_{i\omega} - \delta$ preserves feasibility (all constraints remain satisfied) and improves the objective by $\rho_\omega(\alpha^+_i + \alpha^-_i)\delta \geq 0$. Thus, any optimal solution can be weakly improved until at most one deviation is nonzero per agent per scenario. □

5.1.3 Closed-Form Reformulation. Since Ω is finite, $\mathbb{E}_{\omega\sim\mathcal{P}}[\mathrm{W2}(\boldsymbol{x}, \boldsymbol{y}; \omega)] = \sum_{\omega\in\Omega} \rho_\omega \mathrm{W2}(\boldsymbol{x}, \boldsymbol{y}; \omega)$. Because the second-stage decision variables are independent across scenarios, we can swap the outer maximisation with the inner sum, yielding the following equivalent closed-form programme:

$$\mathrm{W1P}(\mathcal{P}) \;:=\; \arg\max_{\boldsymbol{x},\boldsymbol{y},X,Y} \quad \sum_{\omega\in\Omega} \rho_\omega \left[\mathrm{W2}(\boldsymbol{x}, \boldsymbol{y}, X_\omega, Y_\omega; \omega)\right] \tag{13}$$

subject to

$$\sum_{i\in[n]} x_i = \sum_{j\in[m]} y_j \tag{14}$$

$$\sum_{i\in[n]} (x^+_{i\omega} - x^-_{i\omega}) = \sum_{j\in[m]} (y^+_{j\omega} - y^-_{j\omega}) \quad \forall\, \omega \in \Omega \tag{15}$$

$$x_i + x^+_{i\omega} - x^-_{i\omega} \in [\underline{x}_{i\omega},\ \overline{x}_{i\omega}] \quad \forall\, i \in [n],\ \forall\, \omega \in \Omega \tag{16}$$

$$y_j + y^+_{j\omega} - y^-_{j\omega} \in [\underline{y}_{j\omega},\ \overline{y}_{j\omega}] \quad \forall\, j \in [m],\ \forall\, \omega \in \Omega \tag{17}$$

$$\boldsymbol{x}, \boldsymbol{y}, X, Y \geq 0 \tag{18}$$

5.1.4 Dual of the First-Stage Welfare Problem. We now take the dual of the closed-form W1P($\mathcal{P}$). We assign dual variables $p_0 \geq 0$ to constraint (14), $p_\omega \in \mathbb{R}$ to each constraint (15), and $\boldsymbol{A}_i, \boldsymbol{B}_j \geq 0$ to the bound constraints. After scaling each dual variable by its corresponding scenario probability (i.e., $a^+_{i\omega} \to \rho_\omega a^+_{i\omega}$, etc.), we obtain the equivalent dual W1D($\mathcal{P}$):

$$\mathrm{W1D}(\mathcal{P}) \;:=\; \arg\min_{p_0,\boldsymbol{p},\boldsymbol{A},\boldsymbol{B}} \quad \mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[\sum_{i\in[n]} \left(\overline{x}_{i\omega} a^+_{i\omega} - \underline{x}_{i\omega} a^-_{i\omega}\right) + \sum_{j\in[m]} \left(\overline{y}_{j\omega} b^+_{j\omega} - \underline{y}_{j\omega} b^-_{j\omega}\right) \right] \tag{19}$$

subject to

$$\begin{aligned}
\mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[a^+_{i\omega} - a^-_{i\omega}\right] &\geq \mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[v_{i\omega} - (p_0 + p_\omega)\right] && \forall\, i \in [n] && (20)\\
\mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[b^+_{j\omega} - b^-_{j\omega}\right] &\geq \mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[(p_0 + p_\omega) - c_{j\omega}\right] && \forall\, j \in [m] && (21)\\
a^+_{i\omega} - a^-_{i\omega} &\geq (v_{i\omega} - \alpha^+_i) - p_\omega && \forall\, i,\ \forall\, \omega && (22)\\
a^-_{i\omega} - a^+_{i\omega} &\geq p_\omega - (v_{i\omega} + \alpha^-_i) && \forall\, i,\ \forall\, \omega && (23)\\
b^+_{j\omega} - b^-_{j\omega} &\geq p_\omega - (c_{j\omega} + \alpha^+_j) && \forall\, j,\ \forall\, \omega && (24)\\
b^-_{j\omega} - b^+_{j\omega} &\geq (c_{j\omega} - \alpha^-_j) - p_\omega && \forall\, j,\ \forall\, \omega && (25)\\
\boldsymbol{A}, \boldsymbol{B} &\geq 0 && && (26)
\end{aligned}$$

We also write W1D($\mathcal{P}; p_0, \boldsymbol{p}$) for the programme with the dual prices p_0 and $\boldsymbol{p}$ fixed.

Remark 2. *The dual variable p_0 corresponds to the first-stage balance constraint and p_ω to each second-stage balance constraint. As we show in Section 6, equilibrium prices can be read off from these dual variables.*

5.2 The Equilibrium Problem

Assuming that all agents in the market are price-takers and given prices $(p_0, \boldsymbol{p}) \in \mathbb{R}\times\mathbb{R}^{|\Omega|}$, scenario ω, and day-ahead allocations (x_i, y_j), we now define the first and second-stage utility maximisation problems for individual agents.

$$\begin{aligned}
\mathrm{U2P}_i(x_i, p_0, \boldsymbol{p}; \omega) \;:=\; \arg\max_{\boldsymbol{X}_i} \quad & u_i(x_i, \boldsymbol{X}_i, p_0, \boldsymbol{p}; \omega) \\
\text{subject to} \quad & x_i + x^+_{i\omega} - x^-_{i\omega} \in [\underline{x}_{i\omega},\ \overline{x}_{i\omega}] \\
& \boldsymbol{X}_i \geq 0
\end{aligned} \tag{27}$$

$$\begin{aligned}
\mathrm{U2P}_j(y_j, p_0, \boldsymbol{p}; \omega) \;:=\; \arg\max_{Y_j} \quad & u_j(y_j, Y_j, p_0, \boldsymbol{p}; \omega) \\
\text{subject to} \quad & y_j + y^+_{j\omega} - y^-_{j\omega} \in [\underline{y}_{j\omega},\ \overline{y}_{j\omega}] \\
& Y_j \geq 0
\end{aligned} \tag{28}$$

We write $\mathrm{U2}_i(x_i, p_0, \boldsymbol{p}; \omega)$ and $\mathrm{U2}_j(y_j, p_0, \boldsymbol{p}; \omega)$ for the corresponding optimal values.

5.2.1 First-Stage Utility Problems. Given these, we construct the first-stage utility problems, in which each agent chooses their day-ahead allocation to maximise expected utility, anticipating optimal adjustment in the second stage:

$$\begin{aligned}
\mathrm{U1P}_i(\mathcal{P}, p_0, \boldsymbol{p}) \;:=\; \arg\max_{x_i} \quad & \mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[\mathrm{U2}_i(x_i, p_0, \boldsymbol{p}; \omega)\right] \\
\text{subject to} \quad & x_i \geq 0
\end{aligned} \tag{29}$$

$$\mathrm{U1P}_j(\mathcal{P}, p_0, \boldsymbol{p}) \;:=\; \arg\max_{y_j} \quad \mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[\mathrm{U2}_j(y_j, p_0, \boldsymbol{p}; \omega)\right] \tag{30}$$
$$\text{subject to} \quad y_j \geq 0$$

We denote the optimal values by $\mathrm{U1}_i(\mathcal{P}, p_0, \boldsymbol{p})$ and $\mathrm{U1}_j(\mathcal{P}, p_0, \boldsymbol{p})$.

5.2.2 Duals of the First-Stage Utility Problems. By writing the closed-form primal (combining first and second stages) and scaling dual variables in the same manner as for W1D, we obtain the following dual programmes:

$$\mathrm{U1D}_i(\mathcal{P}, p_0, \boldsymbol{p}) \;:=\; \arg\min_{\boldsymbol{A}_i} \quad \mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[\overline{x}_{i\omega} a^+_{i\omega} - \underline{x}_{i\omega} a^-_{i\omega}\right] \tag{31}$$

subject to

$$\mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[a^+_{i\omega} - a^-_{i\omega}\right] \geq \mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[v_{i\omega} - p_0\right] \tag{32}$$
$$a^+_{i\omega} - a^-_{i\omega} \geq (v_{i\omega} - \alpha^+_i) - p_\omega \quad \forall\, \omega \in \Omega \tag{33}$$
$$a^-_{i\omega} - a^+_{i\omega} \geq p_\omega - (v_{i\omega} + \alpha^-_i) \quad \forall\, \omega \in \Omega \tag{34}$$
$$\boldsymbol{A}_i \geq 0 \tag{35}$$

$$\mathrm{U1D}_j(\mathcal{P}, p_0, \boldsymbol{p}) \;:=\; \arg\min_{\boldsymbol{B}_j} \quad \mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[\overline{y}_{j\omega} b^+_{j\omega} - \underline{y}_{j\omega} b^-_{j\omega}\right] \tag{36}$$

subject to

$$\mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[b^+_{j\omega} - b^-_{j\omega}\right] \geq \mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \left[p_0 - c_{j\omega}\right] \tag{37}$$
$$b^+_{j\omega} - b^-_{j\omega} \geq p_\omega - (c_{j\omega} + \alpha^+_j) \quad \forall\, \omega \in \Omega \tag{38}$$
$$b^-_{j\omega} - b^+_{j\omega} \geq (c_{j\omega} - \alpha^-_j) - p_\omega \quad \forall\, \omega \in \Omega \tag{39}$$
$$\boldsymbol{B}_j \geq 0 \tag{40}$$

6 Stochastic Competitive Equilibrium

We now formally define the two-stage stochastic market $\mathcal{M}$. In the classical Arrow–Debreu framework, a competitive equilibrium is characterised by three properties: feasibility of allocations, utility maximisation by each agent at the posted prices, and market clearing (Walras' law). We generalise this notion to the stochastic two-stage setting, where Walras' law must hold across both stages. We refer to a tuple $(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{X}, Y, p_0, \boldsymbol{p})$ as a *market outcome.*

Definition 2 (Stochastic Competitive Equilibrium with Recourse). *Given a stochastic market $\mathcal{M} = (m, n, \Omega, \mathcal{P})$, a market outcome $(\boldsymbol{x}^*, \boldsymbol{y}^*, \boldsymbol{X}^*, Y^*, p_0^*, \boldsymbol{p}^*)$ is a* stochastic competitive equilibrium with recourse *(SCER) if the following three conditions hold:*

- **(F)** ***Feasibility.*** *The allocations satisfy the constraints of W1P($\mathcal{P}$): first-stage supply equals demand* (14)*, second-stage adjustments balance per scenario* (15)*, all bound constraints* (16)–(17) *hold, and all quantities are non-negative* (18)*.*
- **(UM)** ***Utility Maximisation.*** *Each agent's first- and second-stage allocation is utility maximising at the given prices $(p_0^*, \boldsymbol{p}^*)$. Specifically, for each consumer $i \in [n]$:*

$$(x_i^*,\ \boldsymbol{X}_i^*) \in \mathit{U1P}_i(\mathcal{P}, p_0^*, \boldsymbol{p}^*),$$

and for each producer $j \in [m]$:

$$(y_j^*,\ Y_j^*) \in \mathit{U1P}_j(\mathcal{P}, p_0^*, \boldsymbol{p}^*).$$

(WL) ***Walras' Law.*** *Let z_0 and z_ω denote the aggregate excess demand in the first and second stages, respectively:*

$$z_0 \;:=\; \sum_{i\in[n]} x_i^* - \sum_{j\in[m]} y_j^*, \qquad z_\omega \;:=\; \sum_{i\in[n]} (x_{i\omega}^{*+} - x_{i\omega}^{*-}) - \sum_{j\in[m]} (y_{j\omega}^{*+} - y_{j\omega}^{*-}).$$

The market value of aggregate excess demand is zero in each stage:

$$p_0^* \cdot z_0 = 0, \tag{41}$$
$$p_\omega^* \cdot z_\omega = 0, \quad \forall\, \omega \in \Omega. \tag{42}$$

Remark 3. *Conditions* (41)–(42) *are the natural generalisation of the classical Walras' Law* $\mathbf{p}\cdot\mathbf{z}(\mathbf{p}) = 0$ *to the two-stage stochastic setting, applied independently to each market: the day-ahead market (a single good priced at p_0^*) and each scenario-contingent real-time market (priced at p_ω^*). Each condition asserts that when there is excess supply in a given market ($z < 0$), the corresponding price must be zero; equivalently, a strictly positive price implies exact market clearing.*

We now establish the existence of SCER via a strong-duality argument. The key insight is that prices supporting a competitive equilibrium can be read off from the dual of the welfare maximisation problem.

6.1 Existence of SCER

Theorem 1 (Existence of SCER). *Let $(\boldsymbol{x}^*, \boldsymbol{y}^*, \boldsymbol{X}^*, \boldsymbol{Y}^*)$ be a solution to W1P($\mathcal{P}$). Then there exist prices $(p_0^*, \boldsymbol{p}^*)$ such that the market outcome $(\boldsymbol{x}^*, \boldsymbol{y}^*, \boldsymbol{X}^*, \boldsymbol{Y}^*, p_0^*, \boldsymbol{p}^*)$ is an SCER.*

Proof. Since W1P($\mathcal{P}$) is a finite-dimensional convex programme with affine constraints and a feasible point, strong duality holds. We write the KKT system for the closed-form first-stage formulation (13)–(18), assigning multipliers as follows:

- Day-ahead balance (14): multiplier $\lambda \geq 0$.
- For each $\omega \in \Omega$, second-stage balance (15): multiplier $\nu_\omega \in \mathbb{R}$.
- Consumer bounds (16): multipliers $\overline{\mu}^{\mathrm{c}}_{i,\omega} \geq 0$ (upper) and $\underline{\mu}^{\mathrm{c}}_{i,\omega} \geq 0$ (lower).
- Producer bounds (17): multipliers $\overline{\mu}^{\mathrm{g}}_{j,\omega} \geq 0$ (upper) and $\underline{\mu}^{\mathrm{g}}_{j,\omega} \geq 0$ (lower).
- Non-negativity (18): multipliers σ_i^{c}, σ_j^{g}, $\sigma_{i,\omega}^{\mathrm{c}+}$, $\sigma_{i,\omega}^{\mathrm{c}-}$, $\sigma_{j,\omega}^{\mathrm{g}+}$, $\sigma_{j,\omega}^{\mathrm{g}-} \geq 0$.

We define prices by:

$$p_0^* := \lambda, \qquad p_\omega^* := \frac{\nu_\omega}{\rho_\omega}, \quad \forall\, \omega \in \Omega. \tag{43}$$

Proof of (F): Feasibility. This holds directly, since $(\boldsymbol{x}^*, \boldsymbol{y}^*, \boldsymbol{X}^*, \boldsymbol{Y}^*)$ is a feasible solution to W1P($\mathcal{P}$).

Proof of (UM): Utility Maximisation. We show that the planner's KKT conditions coincide with those of each individual agent's utility problem. Let $x_{i,\omega} := x_i + x_{i\omega}^+ - x_{i\omega}^-$ and $y_{j,\omega} := y_j + y_{j\omega}^+ - y_{j\omega}^-$.

Planner stationarity conditions. For each $i \in [n]$, $j \in [m]$, $\omega \in \Omega$:

$$\partial_{x_i} \mathcal{L}: \quad \textstyle\sum_\omega \rho_\omega v_{i\omega} - \lambda + \sum_\omega (\mu^{c}_{i,\omega} - \underline{\mu}^{c}_{i,\omega}) - \sigma^{c}_{i} = 0, \tag{44}$$

$$\partial_{x^+_{i\omega}} \mathcal{L}: \quad \rho_\omega (v_{i\omega} - \alpha^+_i) - \nu_\omega + (\mu^{c}_{i,\omega} - \underline{\mu}^{c}_{i,\omega}) - \sigma^{c+}_{i,\omega} = 0, \tag{45}$$

$$\partial_{x^-_{i\omega}} \mathcal{L}: \quad \rho_\omega (-v_{i\omega} - \alpha^-_i) + \nu_\omega - (\mu^{c}_{i,\omega} - \underline{\mu}^{c}_{i,\omega}) - \sigma^{c-}_{i,\omega} = 0, \tag{46}$$

$$\partial_{y_j} \mathcal{L}: \quad -\textstyle\sum_\omega \rho_\omega c_{j\omega} + \lambda + \sum_\omega (\mu^{g}_{j,\omega} - \underline{\mu}^{g}_{j,\omega}) - \sigma^{g}_{j} = 0, \tag{47}$$

$$\partial_{y^+_{j\omega}} \mathcal{L}: \quad \rho_\omega (-c_{j\omega} - \alpha^+_j) + \nu_\omega + (\mu^{g}_{j,\omega} - \underline{\mu}^{g}_{j,\omega}) - \sigma^{g+}_{j,\omega} = 0, \tag{48}$$

$$\partial_{y^-_{j\omega}} \mathcal{L}: \quad \rho_\omega (c_{j\omega} - \alpha^-_j) - \nu_\omega - (\mu^{g}_{j,\omega} - \underline{\mu}^{g}_{j,\omega}) - \sigma^{g-}_{j,\omega} = 0. \tag{49}$$

All primal constraints (14)–(18) hold, dual feasibility requires all multipliers to be non-negative, and the standard complementary slackness relations hold for each constraint.

Agent stationarity conditions. Fix prices $(p^*_0, \boldsymbol{p}^*)$ as defined in (43) and consider consumer i maximising expected utility (3) over $(x_i, \boldsymbol{X}_i)$ subject to the per-scenario feasibility $x_{i,\omega} \in [\underline{x}_{i\omega}, \overline{x}_{i\omega}]$ and non-negativity. Let the consumer's bound and non-negativity multipliers be exactly the same symbols as above. The consumer's KKT stationarity conditions are:

$$\textstyle\sum_\omega \rho_\omega v_{i\omega} - \sum_\omega \rho_\omega p^*_0 + \sum_\omega (\mu^{c}_{i,\omega} - \underline{\mu}^{c}_{i,\omega}) - \sigma^{c}_{i} = 0, \tag{50}$$

$$\rho_\omega (v_{i\omega} - p^*_\omega - \alpha^+_i) + (\mu^{c}_{i,\omega} - \underline{\mu}^{c}_{i,\omega}) - \sigma^{c+}_{i,\omega} = 0, \tag{51}$$

$$\rho_\omega (-v_{i\omega} + p^*_\omega - \alpha^-_i) - (\mu^{c}_{i,\omega} - \underline{\mu}^{c}_{i,\omega}) - \sigma^{c-}_{i,\omega} = 0. \tag{52}$$

Similarly, producer j maximises expected utility (4) with the same per-scenario feasibility, yielding:

$$-\textstyle\sum_\omega \rho_\omega c_{j\omega} + \sum_\omega \rho_\omega p^*_0 + \sum_\omega (\mu^{g}_{j,\omega} - \underline{\mu}^{g}_{j,\omega}) - \sigma^{g}_{j} = 0, \tag{53}$$

$$\rho_\omega (p^*_\omega - c_{j\omega} - \alpha^+_j) + (\mu^{g}_{j,\omega} - \underline{\mu}^{g}_{j,\omega}) - \sigma^{g+}_{j,\omega} = 0, \tag{54}$$

$$\rho_\omega (-p^*_\omega + c_{j\omega} - \alpha^-_j) - (\mu^{g}_{j,\omega} - \underline{\mu}^{g}_{j,\omega}) - \sigma^{g-}_{j,\omega} = 0. \tag{55}$$

Matching the two systems. By construction, $\lambda = p^*_0$ and $\nu_\omega = \rho_\omega p^*_\omega$. Comparing (44) with (50): the planner has $-\lambda$ where the consumer has $-\sum_\omega \rho_\omega p^*_0 = -p^*_0$, and since $\lambda = p^*_0$ these coincide. Comparing (45) with (51): the planner has $-\nu_\omega = -\rho_\omega p^*_\omega$ where the consumer has $-\rho_\omega p^*_\omega$, so these match identically. Likewise (46) matches (52). The producer conditions (47)–(49) match (53)–(55) by the same argument. All complementary slackness relations also coincide, since the feasible sets are identical.

Each agent's problem is convex with the same feasible region that the planner imposes, so the KKT conditions are necessary and sufficient. Therefore, at prices $(p^*_0, \boldsymbol{p}^*)$, every agent maximises their utility at allocation $(\boldsymbol{x}^*, \boldsymbol{X}^*, \boldsymbol{y}^*, \boldsymbol{Y}^*)$.

Proof of (WL): Walras' Law. At the optimal primal–dual pair, the complementary slackness conditions on the balance constraints yield:

$$p^*_0 \left(\sum_{j \in [m]} y^*_j - \sum_{i \in [n]} x^*_i \right) = 0, \tag{56}$$

$$p^*_\omega \left(\sum_{j \in [m]} (y^{*+}_{j\omega} - y^{*-}_{j\omega}) - \sum_{i \in [n]} (x^{*+}_{i\omega} - x^{*-}_{i\omega}) \right) = 0, \quad \forall\, \omega \in \Omega. \tag{57}$$

Condition (56) is precisely Walras' law in the first stage (41). Condition (57) is precisely Walras' law in the second stage (42), holding for each scenario $\omega \in \Omega$ individually. □

Remark 4 (Price Interpretation). *The price formulas* (43) *admit a natural economic interpretation. Complementary slackness on the day-ahead balance* (14) *implies that if the constraint is slack then $p_0^* = 0$, and if it binds then p_0^* is the common marginal value/cost of day-ahead quantities. When no bound constraints are binding, the stationarity conditions* (45)–(49) *yield the familiar marginal-equals-price relations. In particular, for any scenario ω:*

$$p_\omega^* = v_{i\omega} - \alpha_i^+ \;\; (\text{if } x_{i\omega}^{*+} > 0), \qquad p_\omega^* = c_{j\omega} + \alpha_j^+ \;\; (\text{if } y_{j\omega}^{*+} > 0).$$

When bound constraints bind, these equalities become the corresponding one-sided conditions; equivalently, p_ω^* lies in the interval consistent with the active set at ω.

7 Fundamental Welfare Theorems of Economics

Having established the existence of SCER via Theorem 1, we now state the two fundamental welfare theorems for our stochastic two-stage market. The first welfare theorem states that any SCER allocation is welfare maximising. The second welfare theorem asserts the converse: any welfare-maximising allocation can be supported as an SCER by suitable prices. Together, they establish the equivalence between competitive equilibria and Pareto (social welfare) optima.

Theorem 2 (First Welfare Theorem for SCER). *If $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*, p_0^*, \boldsymbol{p}^*)$ is an SCER, then $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*)$ is a welfare-maximising allocation; i.e., it solves W1P($\mathcal{P}$).*

Proof. Since the market outcome is an SCER, condition **(UM)** gives that each consumer i and producer j is individually utility maximising at prices $(p_0^*, \boldsymbol{p}^*)$. Let $\mathcal{A}$ be the programme obtained by summing all agents' utility-maximisation objectives and aggregating their individual constraints. Since the agents' problems are independent (prices are fixed), $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*)$ is an optimal solution to $\mathcal{A}$.

We show that $\mathcal{A}$ and W1P($\mathcal{P}$) have the same solutions. By Feasibility **(F)**, the price-dependent terms in the *objective* cancel when summing agent utilities, so the objective of $\mathcal{A}$ reduces to the welfare function W1. For the *constraints*: the individual bound and non-negativity constraints are exactly those of W1P($\mathcal{P}$), while the market-clearing constraints (14)–(15) hold by condition **(F)**.

Since $\mathcal{A}$ and W1P($\mathcal{P}$) share the same objective and feasible set, any optimal solution to $\mathcal{A}$ is also optimal for W1P($\mathcal{P}$). □

Theorem 3 (Second Welfare Theorem for SCER). *If $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*)$ is a feasible welfare-maximising allocation (i.e., it solves W1P($\mathcal{P}$)), then there exist prices $(p_0^*, \boldsymbol{p}^*)$ such that $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*, p_0^*, \boldsymbol{p}^*)$ is an SCER.*

Proof. This is a direct consequence of Theorem 1. Equilibrium prices are constructed from the dual multipliers of W1P($\mathcal{P}$) via the identification (43). Theorem 1 establishes that all three SCER conditions, that is, feasibility, utility maximisation (via KKT matching), and Walras' law (via complementary slackness) are satisfied at these prices. □

8 Risk-Aware Day Ahead Clearing

The stochastic clearing formulation of Section 5.1 maximises *expected* welfare, which treats all scenarios symmetrically according to their probabilities. Previous work has shown that renewable generation profiles often follow a heavy-tailed distribution. Hence, in practice, market operators may wish to hedge against tail events-scenarios in which welfare is particularly poor (for example, due to renewable generation shortfalls). Conditional Value-at-Risk (CVaR), introduced in Section 4, provides a coherent risk measure that quantifies the expected welfare in the worst-case tail of the distribution.

In this section, we incorporate CVaR into the day-ahead clearing problem. We first formulate the risk-aware optimisation problem as a mean–CVaR objective. We then address the equilibrium problem, where a key challenge arises: the CVaR risk measure introduces scenario-specific dual multipliers that, in general, differ across agents. Following previous work, we address this by introducing an insurance (Arrow-Debreu) securities market that completes the market and aligns agents' risk assessments with the planner's.

8.1 The Optimization Problem

We consider a social planner who seeks to maximise a convex combination of expected welfare and its CVaR at confidence level $\beta \in [0, 1)$, with the trade-off controlled by a risk-aversion parameter $\varepsilon \in [0, 1]$. Denoting the per-scenario loss by $\ell(\omega) := -\mathrm{W2}(\boldsymbol{x}, \boldsymbol{y}; \omega)$, the social planner optimises:

$$\min_{\boldsymbol{x},\boldsymbol{y},X,Y} \quad (1-\varepsilon) \mathop{\mathbb{E}}_{\omega\sim\mathcal{P}} \big[\ell(\omega)\big] + \varepsilon \cdot \mathrm{CVaR}_\beta\big[\ell(\omega)\big]. \tag{58}$$

When $\varepsilon = 0$, this reduces to the risk-neutral problem W1P($\mathcal{P}$). When $\varepsilon = 1$, the planner optimises pure CVaR, hedging only against the worst $(1-\beta)$-fraction of scenarios. For intermediate values, the planner trades off expected performance against tail protection.

Using the CVaR characterisation from Theorem 2 of Section 4.1 (due to [13]), we can reformulate the CVaR term via an auxiliary variable. Introducing a scalar $t \in \mathbb{R}$ and slack variables $u_\omega \geq 0$ for each $\omega \in \Omega$, we obtain the following epigraph reformulation:

$$\min_{\boldsymbol{x},\boldsymbol{y},X,Y,t,\{u_\omega\}} \quad -(1-\varepsilon)\sum_{\omega\in\Omega} \rho_\omega\, \mathrm{W2}(\boldsymbol{x},\boldsymbol{y},X_\omega,Y_\omega;\omega) + \varepsilon\left(t + \frac{1}{1-\beta}\sum_{\omega\in\Omega}\rho_\omega\, u_\omega\right) \tag{59}$$

subject to

$$\sum_{i\in[n]} x_i = \sum_{j\in[m]} y_j \tag{60}$$

$$\sum_{i\in[n]} (x^+_{i\omega} - x^-_{i\omega}) = \sum_{j\in[m]} (y^+_{j\omega} - y^-_{j\omega}) \quad \forall\, \omega \in \Omega \tag{61}$$

$$u_\omega \geq -t - \mathrm{W2}(\boldsymbol{x},\boldsymbol{y},X_\omega,Y_\omega;\omega) \quad \forall\, \omega \in \Omega \tag{62}$$

$$x_i + x^+_{i\omega} - x^-_{i\omega} \in [\underline{x}_{i\omega},\ \overline{x}_{i\omega}] \quad \forall\, i, \forall\, \omega \tag{63}$$

$$y_j + y^+_{j\omega} - y^-_{j\omega} \in [\underline{y}_{j\omega},\ \overline{y}_{j\omega}] \quad \forall\, j, \forall\, \omega \tag{64}$$

$$\boldsymbol{x},\boldsymbol{y},X,Y,u_\omega \geq 0 \tag{65}$$

8.1.1 Lagrangian and KKT Conditions. Before deriving the optimality conditions, we rewrite the minimisation problem (59)–(65) as the equivalent maximisation:

$$\max_{\boldsymbol{x},\boldsymbol{y},X,Y,t,\{u_\omega\}} \quad (1-\varepsilon)\sum_{\omega\in\Omega} \rho_\omega\, \mathrm{W2}(\boldsymbol{x},\boldsymbol{y},X_\omega,Y_\omega;\omega) - \varepsilon\left(t + \frac{1}{1-\beta}\sum_{\omega\in\Omega}\rho_\omega\, u_\omega\right) \tag{66}$$

subject to the same constraints (60)–(65). We assign dual variables as follows:

- Day-ahead balance (60): multiplier $\lambda \geq 0$.
- For each $\omega \in \Omega$, real-time balance (61): multiplier $\nu_\omega \in \mathbb{R}$.
- For each $\omega \in \Omega$, CVaR shortfall (62): multiplier $\gamma_\omega \geq 0$.
- For each $\omega \in \Omega$, non-negativity $u_\omega \geq 0$: multiplier $\eta_\omega \geq 0$.
- Consumer bounds (63): multipliers $\overline{\mu}^{\mathrm{c}}_{i,\omega} \geq 0$ (upper) and $\underline{\mu}^{\mathrm{c}}_{i,\omega} \geq 0$ (lower).

- Producer bounds (64): multipliers $\overline{\mu}^{\mathrm{g}}_{j,\omega} \geq 0$ (upper) and $\underline{\mu}^{\mathrm{g}}_{j,\omega} \geq 0$ (lower).
- Non-negativity (65): multipliers $\sigma^{\mathrm{c}}_i, \sigma^{\mathrm{g}}_j, \sigma^{\mathrm{c}+}_{i,\omega}, \sigma^{\mathrm{c}-}_{i,\omega}, \sigma^{\mathrm{g}+}_{j,\omega}, \sigma^{\mathrm{g}-}_{j,\omega} \geq 0$.

We define the *CVaR-tilted probability weight*

$$\tilde{\rho}_\omega \;:=\; (1-\varepsilon)\,\rho_\omega + \gamma_\omega, \qquad \forall\,\omega \in \Omega. \tag{67}$$

Let $x_{i,\omega} := x_i + x^+_{i\omega} - x^-_{i\omega}$ and $y_{j,\omega} := y_j + y^+_{j\omega} - y^-_{j\omega}$. The Lagrangian is:

$$\begin{aligned}
\mathcal{L} = \sum_{\omega\in\Omega} \tilde{\rho}_\omega \Bigg[& \sum_{i\in[n]} \Big(v_{i\omega}\, x_{i,\omega} - \alpha^+_i\, x^+_{i\omega} - \alpha^-_i\, x^-_{i\omega} \Big) \\
& - \sum_{j\in[m]} \Big(c_{j\omega}\, y_{j,\omega} + \alpha^+_j\, y^+_{j\omega} + \alpha^-_j\, y^-_{j\omega} \Big) \Bigg] \\
& + \Big(\sum_{\omega\in\Omega} \gamma_\omega - \varepsilon \Big)\, t + \sum_{\omega\in\Omega} \Big(\gamma_\omega + \eta_\omega - \frac{\varepsilon\,\rho_\omega}{1-\beta} \Big)\, u_\omega \\
& + \lambda \Big(\sum_{j\in[m]} y_j - \sum_{i\in[n]} x_i \Big) + \sum_{\omega\in\Omega} \nu_\omega \Big(\sum_{j\in[m]} (y^+_{j\omega} - y^-_{j\omega}) - \sum_{i\in[n]} (x^+_{i\omega} - x^-_{i\omega}) \Big) \\
& + \sum_{\omega\in\Omega} \sum_{i\in[n]} \Big[\overline{\mu}^{\mathrm{c}}_{i,\omega}\, (x_{i,\omega} - \overline{x}_{i\omega}) - \underline{\mu}^{\mathrm{c}}_{i,\omega}\, (x_{i,\omega} - \underline{x}_{i\omega}) \Big] \\
& + \sum_{\omega\in\Omega} \sum_{j\in[m]} \Big[\overline{\mu}^{\mathrm{g}}_{j,\omega}\, (y_{j,\omega} - \overline{y}_{j\omega}) - \underline{\mu}^{\mathrm{g}}_{j,\omega}\, (y_{j,\omega} - \underline{y}_{j\omega}) \Big] \\
& - \sum_{i\in[n]} \sigma^{\mathrm{c}}_i\, x_i - \sum_{j\in[m]} \sigma^{\mathrm{g}}_j\, y_j \\
& - \sum_{\omega\in\Omega} \sum_{i\in[n]} \Big(\sigma^{\mathrm{c}+}_{i,\omega}\, x^+_{i\omega} + \sigma^{\mathrm{c}-}_{i,\omega}\, x^-_{i\omega} \Big) - \sum_{\omega\in\Omega} \sum_{j\in[m]} \Big(\sigma^{\mathrm{g}+}_{j,\omega}\, y^+_{j\omega} + \sigma^{\mathrm{g}-}_{j,\omega}\, y^-_{j\omega} \Big).
\end{aligned} \tag{68}$$

8.1.2 KKT Stationarity Conditions. Setting $\partial_t \mathcal{L} = 0$ and $\partial_{u_\omega} \mathcal{L} = 0$ yields:

$$\partial_t \mathcal{L}: \quad \sum_{\omega\in\Omega} \gamma_\omega = \varepsilon, \tag{69}$$

$$\partial_{u_\omega} \mathcal{L}: \quad \gamma_\omega + \eta_\omega = \frac{\varepsilon\,\rho_\omega}{1-\beta}, \quad \forall\,\omega \in \Omega. \tag{70}$$

Since $\eta_\omega \geq 0$, it follows that $0 \leq \gamma_\omega \leq \varepsilon\,\rho_\omega/(1-\beta)$ for all ω. Moreover, $\sum_\omega \tilde{\rho}_\omega = 1+\varepsilon$, so the normalised weights $\tilde{\rho}_\omega/(1+\varepsilon)$ define a probability measure that places additional mass on tail scenarios (those with $\gamma_\omega > 0$).

The stationarity conditions for the physical variables are identical in structure to the risk-neutral conditions (44)–(49) with ρ_ω replaced everywhere by $\tilde{\rho}_\omega$:

$$\partial_{x_i}\mathcal{L}: \quad \textstyle\sum_\omega \tilde{\rho}_\omega\, v_{i\omega} - \lambda + \sum_\omega(\mu^{\mathrm{c}}_{i,\omega} - \underline{\mu}^{\mathrm{c}}_{i,\omega}) - \sigma^{\mathrm{c}}_i = 0, \tag{71}$$

$$\partial_{x^+_{i\omega}}\mathcal{L}: \quad \tilde{\rho}_\omega(v_{i\omega} - \alpha^+_i) - \nu_\omega + (\mu^{\mathrm{c}}_{i,\omega} - \underline{\mu}^{\mathrm{c}}_{i,\omega}) - \sigma^{\mathrm{c}+}_{i,\omega} = 0, \tag{72}$$

$$\partial_{x^-_{i\omega}}\mathcal{L}: \quad \tilde{\rho}_\omega(-v_{i\omega} - \alpha^-_i) + \nu_\omega - (\mu^{\mathrm{c}}_{i,\omega} - \underline{\mu}^{\mathrm{c}}_{i,\omega}) - \sigma^{\mathrm{c}-}_{i,\omega} = 0, \tag{73}$$

$$\partial_{y_j}\mathcal{L}: \quad -\textstyle\sum_\omega \tilde{\rho}_\omega\, c_{j\omega} + \lambda + \sum_\omega(\mu^{\mathrm{g}}_{j,\omega} - \underline{\mu}^{\mathrm{g}}_{j,\omega}) - \sigma^{\mathrm{g}}_j = 0, \tag{74}$$

$$\partial_{y^+_{j\omega}}\mathcal{L}: \quad \tilde{\rho}_\omega(-c_{j\omega} - \alpha^+_j) + \nu_\omega + (\mu^{\mathrm{g}}_{j,\omega} - \underline{\mu}^{\mathrm{g}}_{j,\omega}) - \sigma^{\mathrm{g}+}_{j,\omega} = 0, \tag{75}$$

$$\partial_{y^-_{j\omega}}\mathcal{L}: \quad \tilde{\rho}_\omega(c_{j\omega} - \alpha^-_j) - \nu_\omega - (\mu^{\mathrm{g}}_{j,\omega} - \underline{\mu}^{\mathrm{g}}_{j,\omega}) - \sigma^{\mathrm{g}-}_{j,\omega} = 0. \tag{76}$$

All primal constraints hold, dual feasibility requires all multipliers to be non-negative, and the standard complementary slackness relations hold for each constraint, together with

$$\gamma_\omega\big(u_\omega + \mathrm{W2}(\cdot\,;\omega) + t\big) = 0, \qquad \eta_\omega\, u_\omega = 0, \qquad \forall\, \omega \in \Omega. \tag{77}$$

8.2 The Equilibrium Problem

We now define the risk-aware utility maximisation problem for individual agents. Each agent evaluates their utility using the same mean–CVaR criterion as the social planner, with common risk-aversion parameter ε and confidence level β. Given posted prices $(p_0, \boldsymbol{p})$, consumer i solves the following first-stage problem:

$$\begin{aligned}\mathrm{CVaRP}_i(\mathcal{P}, p_0, \boldsymbol{p}) \;:=\; \arg \max_{x_i, \boldsymbol{X}_i, t^c_i, \{\zeta^c_{i,\omega}\}} \quad & (1-\varepsilon)\sum_{\omega\in\Omega}\rho_\omega\, u_i(x_i, \boldsymbol{X}_i, p_0, \boldsymbol{p}; \omega) \\ & - \varepsilon\left(t^c_i + \frac{1}{1-\beta}\sum_{\omega\in\Omega}\rho_\omega\, \zeta^c_{i,\omega}\right)\end{aligned} \tag{78}$$

subject to

$$\begin{aligned} \zeta^c_{i,\omega} &\geq -u_i(x_i, \boldsymbol{X}_i, p_0, \boldsymbol{p}; \omega) - t^c_i, \quad \forall\, \omega \in \Omega & (79)\\ x_i + x^+_{i\omega} - x^-_{i\omega} &\in [\underline{x}_{i\omega},\ \overline{x}_{i\omega}], \quad \forall\, \omega \in \Omega & (80)\\ x_i, \boldsymbol{X}_i, \zeta^c_{i,\omega} &\geq 0 & (81)\end{aligned}$$

where u_i is the quasi-linear utility defined in (3). The agent maximises expected utility less ε times the CVaR of their negative utility (loss), hedging against scenarios in which their utility is particularly poor.

Similarly, producer j solves:

$$\begin{aligned}\mathrm{CVaRP}_j(\mathcal{P}, p_0, \boldsymbol{p}) \;:=\; \arg \max_{y_j, Y_j, t^g_j, \{\zeta^g_{j,\omega}\}} \quad & (1-\varepsilon)\sum_{\omega\in\Omega}\rho_\omega\, u_j(y_j, \boldsymbol{Y}_j, p_0, \boldsymbol{p}; \omega) \\ & - \varepsilon\left(t^g_j + \frac{1}{1-\beta}\sum_{\omega\in\Omega}\rho_\omega\, \zeta^g_{j,\omega}\right)\end{aligned} \tag{82}$$

subject to

$$\zeta^g_{j,\omega} \geq -u_j(y_j, Y_j, p_0, \boldsymbol{p}; \omega) - t^g_j, \quad \forall\, \omega \in \Omega \tag{83}$$

$$y_j + y^+_{j\omega} - y^-_{j\omega} \in [\underline{y}_{j\omega}, \ \overline{y}_{j\omega}], \quad \forall\, \omega \in \Omega \tag{84}$$

$$y_j, Y_j, \zeta^g_{j,\omega} \geq 0 \tag{85}$$

8.2.1 Lagrangian of the Equilibrium Problem. We now write the Lagrangian for the consumer's risk-aware problem (78)–(81). Introduce dual variables $\gamma^c_{i,\omega} \geq 0$ for each CVaR shortfall constraint (79), $\eta^c_{i,\omega} \geq 0$ for non-negativity of $\zeta^c_{i,\omega}$, and the usual bound/non-negativity multipliers $\mu^{\mathrm{c}}_{i,\omega}$, $\underline{\mu}^{\mathrm{c}}_{i,\omega}$, σ^{c}_i, $\sigma^{\mathrm{c}+}_{i,\omega}$, $\sigma^{\mathrm{c}-}_{i,\omega} \geq 0$. The Lagrangian is:

$$\begin{aligned}
\mathcal{L}^{\mathrm{c}}_i = & \sum_{\omega\in\Omega} \left((1-\varepsilon)\,\rho_\omega + \gamma^c_{i,\omega}\right) u_i(x_i, \boldsymbol{X}_i, p_0, \boldsymbol{p}; \omega) \\
& + \Big(\sum_{\omega\in\Omega} \gamma^c_{i,\omega} - \varepsilon\Big)\, t^c_i + \sum_{\omega\in\Omega} \Big(\gamma^c_{i,\omega} + \eta^c_{i,\omega} - \frac{\varepsilon\,\rho_\omega}{1-\beta}\Big)\, \zeta^c_{i,\omega} \\
& + \sum_{\omega\in\Omega} \Big[\mu^{\mathrm{c}}_{i,\omega}\,(x_{i,\omega} - \overline{x}_{i\omega}) - \underline{\mu}^{\mathrm{c}}_{i,\omega}\,(x_{i,\omega} - \underline{x}_{i\omega})\Big] \\
& - \sigma^{\mathrm{c}}_i\, x_i - \sum_{\omega\in\Omega} \Big(\sigma^{\mathrm{c}+}_{i,\omega}\, x^+_{i\omega} + \sigma^{\mathrm{c}-}_{i,\omega}\, x^-_{i\omega}\Big).
\end{aligned} \tag{86}$$

Setting $\partial_{t^c_i}\mathcal{L}^{\mathrm{c}}_i = 0$ and $\partial_{\zeta^c_{i,\omega}}\mathcal{L}^{\mathrm{c}}_i = 0$ gives the agent-specific analogues of (69)–(70):

$$\sum_{\omega\in\Omega} \gamma^c_{i,\omega} = \varepsilon, \qquad \gamma^c_{i,\omega} + \eta^c_{i,\omega} = \frac{\varepsilon\,\rho_\omega}{1-\beta}, \quad \forall\, \omega \in \Omega. \tag{87}$$

Defining the agent's tilted weight $\tilde{\rho}^c_{i,\omega} := (1-\varepsilon)\,\rho_\omega + \gamma^c_{i,\omega}$, we have $\sum_\omega \tilde{\rho}^c_{i,\omega} = (1-\varepsilon) + \varepsilon = 1$, so the agent's tilted weights also form a probability measure. The stationarity conditions for x_i, $x^+_{i\omega}$, and $x^-_{i\omega}$ become:

$$\partial_{x_i}\mathcal{L}^{\mathrm{c}}_i: \quad \sum_\omega \tilde{\rho}^c_{i,\omega}\,(v_{i\omega} - p_0) + \sum_\omega (\mu^{\mathrm{c}}_{i,\omega} - \underline{\mu}^{\mathrm{c}}_{i,\omega}) - \sigma^{\mathrm{c}}_i = 0, \tag{88}$$

$$\partial_{x^+_{i\omega}}\mathcal{L}^{\mathrm{c}}_i: \quad \tilde{\rho}^c_{i,\omega}\,(v_{i\omega} - p_\omega - \alpha^+_i) + (\mu^{\mathrm{c}}_{i,\omega} - \underline{\mu}^{\mathrm{c}}_{i,\omega}) - \sigma^{\mathrm{c}+}_{i,\omega} = 0, \tag{89}$$

$$\partial_{x^-_{i\omega}}\mathcal{L}^{\mathrm{c}}_i: \quad \tilde{\rho}^c_{i,\omega}\,(-v_{i\omega} + p_\omega - \alpha^-_i) - (\mu^{\mathrm{c}}_{i,\omega} - \underline{\mu}^{\mathrm{c}}_{i,\omega}) - \sigma^{\mathrm{c}-}_{i,\omega} = 0. \tag{90}$$

The producer's Lagrangian and KKT conditions are analogous, with tilted weight $\tilde{\rho}^g_{j,\omega} := (1-\varepsilon)\,\rho_\omega + \gamma^g_{j,\omega}$ (also summing to 1) and stationarity:

$$\partial_{y_j}\mathcal{L}^{\mathrm{g}}_j: \quad -\sum_\omega \tilde{\rho}^g_{j,\omega}\,(c_{j\omega} - p_0) + \sum_\omega (\mu^{\mathrm{g}}_{j,\omega} - \underline{\mu}^{\mathrm{g}}_{j,\omega}) - \sigma^{\mathrm{g}}_j = 0, \tag{91}$$

$$\partial_{y^+_{j\omega}}\mathcal{L}^{\mathrm{g}}_j: \quad \tilde{\rho}^g_{j,\omega}\,(p_\omega - c_{j\omega} - \alpha^+_j) + (\mu^{\mathrm{g}}_{j,\omega} - \underline{\mu}^{\mathrm{g}}_{j,\omega}) - \sigma^{\mathrm{g}+}_{j,\omega} = 0, \tag{92}$$

$$\partial_{y^-_{j\omega}}\mathcal{L}^{\mathrm{g}}_j: \quad \tilde{\rho}^g_{j,\omega}\,(-p_\omega + c_{j\omega} - \alpha^-_j) - (\mu^{\mathrm{g}}_{j,\omega} - \underline{\mu}^{\mathrm{g}}_{j,\omega}) - \sigma^{\mathrm{g}-}_{j,\omega} = 0. \tag{93}$$

8.2.2 Challenge for decentralisation. We now compare the planner's KKT (71)–(76) with the agents' KKT (88)–(93). Consider the stationarity condition for x_i. The planner (71) has:

$$\sum_{\omega} \tilde{\rho}_\omega \, v_{i\omega} - \lambda + \cdots = 0,$$

while the agent (88) has:

$$\sum_{\omega} \tilde{\rho}^c_{i,\omega} \, v_{i\omega} - p_0 + \cdots = 0,$$

since $\sum_\omega \tilde{\rho}^c_{i,\omega} = 1$. For the real-time variables, comparing (72) with (89) requires $\tilde{\rho}_\omega = \tilde{\rho}^c_{i,\omega}$ and $v_\omega = \tilde{\rho}_\omega \, p_\omega$. Therefore, the planner's and agents' conditions coincide if and only if:

$$\gamma^c_{i,\omega} = \gamma^g_{j,\omega} = \gamma_\omega, \qquad \forall\, i \in [n],\; j \in [m],\; \omega \in \Omega. \tag{94}$$

That is, all agents' CVaR tail multipliers must equal the planner's. In general, there is no reason for this to hold: each agent's CVaR multiplier $\gamma^c_{i,\omega}$ is determined by the tail structure of their *individual* utility, which varies across agents. Two agents may experience very different utility levels in the same scenario (e.g., a solar generator and a gas generator face different cost structures under the same weather outcome), leading to different sets of tail scenarios and hence different γ-vectors.

Without a mechanism to align these multipliers, the risk-averse planner's solution *cannot* in general be decentralised as a competitive equilibrium. This observation is consistent with the findings of [6], who demonstrate that in incomplete risk markets, risk-averse equilibria may be non-unique and need not coincide with the social planner's optimum. To resolve this, we introduce a market for state-contingent claims that allows agents to trade risk.

8.3 Insurance Market

We complete the market by introducing *Arrow–Debreu securities* [?]. For each scenario $\omega \in \Omega$, an Arrow–Debreu security is a contract that charges a price q_ω in the first stage and delivers one unit of payment in scenario ω. For each consumer i and producer j, let

$$r^c_{i\omega} \in \mathbb{R}, \qquad r^g_{j\omega} \in \mathbb{R}$$

denote the net units of the Arrow–Debreu security for scenario ω purchased by consumer i and producer j, respectively. Positive values represent insurance indemnities received; negative values represent premiums paid.

The insurance market clears state-by-state:

$$\sum_{i \in [n]} r^c_{i\omega} + \sum_{j \in [m]} r^g_{j\omega} = 0, \qquad \forall\, \omega \in \Omega. \tag{95}$$

8.4 The Equilibrium Problem with Insurance

With the insurance market available, each agent's problem is augmented with insurance decisions. Consumer i solves:

$$\begin{aligned}\max_{x_i, \boldsymbol{X}_i, \{r^c_{i\omega}\}, t^c_i, \{\zeta^c_{i,\omega}\}} \quad & (1-\varepsilon) \sum_{\omega \in \Omega} \rho_\omega \left[u_i(x_i, \boldsymbol{X}_i, p_0, \boldsymbol{p}; \omega) + r^c_{i\omega} \right] - \sum_{\omega \in \Omega} q_\omega \, r^c_{i\omega} \\ & - \varepsilon \left(t^c_i + \frac{1}{1-\beta} \sum_{\omega \in \Omega} \rho_\omega \, \zeta^c_{i,\omega} \right) \end{aligned} \tag{96}$$

subject to

$$\zeta^c_{i,\omega} \geq -u_i(x_i, \boldsymbol{X}_i, p_0, \boldsymbol{p}; \omega) - r^c_{i\omega} - t^c_i, \quad \forall\, \omega \in \Omega \tag{97}$$

$$x_i + x^+_{i\omega} - x^-_{i\omega} \in [\underline{x}_{i\omega},\ \overline{x}_{i\omega}], \quad \forall\, \omega \in \Omega \tag{98}$$

$$x_i, \boldsymbol{X}_i, \zeta^c_{i,\omega} \geq 0. \tag{99}$$

The agent's stochastic per-scenario payoff is $u_i(\cdot\,;\omega) + r^c_{i\omega}$, that is, the energy utility plus the contingent insurance payout in scenario ω. Risk preferences are applied to this stochastic payoff, yielding the $(1-\varepsilon)$-weighted expectation of $u_i + r^c_{i\omega}$ together with the ε-weighted CVaR of its negative (i.e., the insured loss $-u_i - r^c_{i\omega}$). The insurance premium $\sum_\omega q_\omega\, r^c_{i\omega}$, paid in the first stage before any uncertainty resolves, is a deterministic cost that sits outside the risk measure.

Similarly, producer j solves:

$$\max_{y_j, Y_j, \{r^g_{j\omega}\}, t^g_j, \{\zeta^g_{j,\omega}\}} \quad (1-\varepsilon) \sum_{\omega\in\Omega} \rho_\omega \left[u_j(y_j, Y_j, p_0, \boldsymbol{p}; \omega) + r^g_{j\omega}\right] - \sum_{\omega\in\Omega} q_\omega\, r^g_{j\omega} \\ - \varepsilon\left(t^g_j + \frac{1}{1-\beta} \sum_{\omega\in\Omega} \rho_\omega\, \zeta^g_{j,\omega}\right) \tag{100}$$

subject to

$$\zeta^g_{j,\omega} \geq -u_j(y_j, Y_j, p_0, \boldsymbol{p}; \omega) - r^g_{j\omega} - t^g_j, \quad \forall\, \omega \in \Omega \tag{101}$$

$$y_j + y^+_{j\omega} - y^-_{j\omega} \in [\underline{y}_{j\omega},\ \overline{y}_{j\omega}], \quad \forall\, \omega \in \Omega \tag{102}$$

$$y_j, Y_j, \zeta^g_{j,\omega} \geq 0. \tag{103}$$

8.5 Risk-Aware Stochastic Competitive Equilibrium (R-SCER)

We now formalise the notion of competitive equilibrium for the risk-aware market augmented with insurance trading.

Definition 3 (Risk-Aware SCER with Insurance). *Given a stochastic market $\mathcal{M} = (m, n, \Omega, \mathcal{P})$ and common risk parameters (ε, β), a market outcome $(\boldsymbol{x}^*, \boldsymbol{y}^*, \boldsymbol{X}^*, Y^*, r^{c*}, r^{g*}, p^*_0, \boldsymbol{p}^*, q^*)$ is a* risk-aware stochastic competitive equilibrium with recourse and insurance *(R-SCER) if the following conditions hold:*

- **(F)** ***Feasibility.*** *The physical allocations satisfy the constraints of the risk-aware planner problem* (60)–(65).
- **(UM)** ***Utility Maximisation.*** *At prices $(p^*_0, \boldsymbol{p}^*, q^*)$, each consumer i maximises their risk-adjusted utility: $(x^*_i, \boldsymbol{X}^*_i, r^{c*}_i)$ solves* (96)–(99). *Similarly, each producer j maximises their risk-adjusted utility: (y^*_j, Y^*_j, r^{g*}_j) solves* (100)–(103).
- **(WL)** ***Walras' Law.*** *The market value of aggregate excess demand is zero in each market:*

$$p^*_0 \cdot z_0 = 0, \tag{104}$$

$$p^*_\omega \cdot z_\omega = 0, \qquad \forall\, \omega \in \Omega, \tag{105}$$

$$q^*_\omega \cdot \left(\sum_{i\in[n]} r^c_{i\omega^*} + \sum_{j\in[m]} r^g_{j\omega^*}\right) = 0, \qquad \forall\, \omega \in \Omega, \tag{106}$$

where z_0 and z_ω are the aggregate excess demands defined as in Definition 2.

Remark 5. *Condition* (106) *is the natural Walras' law for the insurance market: if there is excess supply of insurance in scenario ω, the corresponding price q^*_ω is zero. A strictly positive insurance price implies that the insurance market clears exactly in that state.*

Remark 6. *The assumption of common risk parameters* (ε, β) *is substantive. As discussed in Section 8.2.2, the alignment of CVaR multipliers requires all agents to agree on the risk-aversion weight ε and the confidence level β. The insurance market aligns the tail-weighting* across agents *(the γ-vectors), but it cannot reconcile fundamentally different risk preferences. We discuss the extension of the mechanism which relaxes to different risk-aversion β in a later section.*

8.5.1 Existence of R-SCER..

Theorem 4 (Existence of R-SCER). *Let* $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*, t^*, u^*)$ *be a solution to the risk-aware planner problem* (66) *with optimal dual variables* $(\lambda^*, \nu^*_\omega, \gamma^*_\omega, \eta^*_\omega)$ *and the associated bound/non-negativity multipliers. Define prices:*

$$p^*_0 := \lambda^*, \qquad p^*_\omega := \frac{\nu^*_\omega}{\tilde{\rho}^*_\omega}, \qquad q^*_\omega := \tilde{\rho}^*_\omega, \qquad \forall\, \omega \in \Omega, \tag{107}$$

where $\tilde{\rho}^*_\omega = (1-\varepsilon)\,\rho_\omega + \gamma^*_\omega$.

Then there exist insurance transfers ${r^c_{i\omega}}^*$ *and* ${r^g_{j\omega}}^*$ *such that* $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*, r^{c*}, r^{g*}, p^*_0, \boldsymbol{p}^*, q^*)$ *is an R-SCER.*

Proof. The proof establishes each R-SCER condition in turn.

Proof of (F): Feasibility. This holds directly, since $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*)$ is a feasible solution to the planner problem.

Proof of (UM): Utility Maximisation. We show that the planner's KKT conditions coincide with those of each insured agent at the prices (107).

Insurance prices align CVaR multipliers. Consider the Lagrangian of consumer i's insured problem (96)–(99). Let $\gamma^c_{i,\omega} \geq 0$ denote the multiplier on the CVaR shortfall constraint (97). Stationarity with respect to $r^c_{i\omega}$ gives:

$$\partial_{r^c_{i\omega}} \mathcal{L}^c_i = (1-\varepsilon)\rho_\omega - q_\omega + \gamma^c_{i,\omega} = 0 \implies q_\omega = (1-\varepsilon)\rho_\omega + \gamma^c_{i,\omega}. \tag{108}$$

The right-hand side is precisely the agent's CVaR-tilted weight $\tilde{\rho}^c_{i,\omega}$, so equilibrium insurance prices coincide with each agent's tilted probability mass: $q_\omega = \tilde{\rho}^c_{i,\omega}$. An analogous computation for producer j gives $q_\omega = \tilde{\rho}^g_{j,\omega}$.

Setting $q^*_\omega = \tilde{\rho}^*_\omega = (1-\varepsilon)\rho_\omega + \gamma^*_\omega$ ensures that every agent's tilted weight equals the planner's:

$$\tilde{\rho}^c_{i,\omega} = \tilde{\rho}^g_{j,\omega} = \tilde{\rho}^*_\omega, \qquad \forall\, i, j, \omega, \tag{109}$$

or equivalently, every agent's CVaR multiplier equals the planner's: $\gamma^c_{i,\omega} = \gamma^g_{j,\omega} = \gamma^*_\omega$, assuming that the agents share knowledge of the same probability distribution and ϵ.

Real-time variables. Comparing the agent's condition (89) (with $\tilde{\rho}^c_{i,\omega} = \tilde{\rho}^*_\omega$):

$$\tilde{\rho}^*_\omega(v_{i\omega} - p^*_\omega - \alpha^+_i) + (\mu^c_{i,\omega} - \underline{\mu}^c_{i,\omega}) - \sigma^{c+}_{i,\omega} = 0$$

with the planner's condition (72):

$$\tilde{\rho}^*_\omega(v_{i\omega} - \alpha^+_i) - \nu^*_\omega + (\mu^c_{i,\omega} - \underline{\mu}^c_{i,\omega}) - \sigma^{c+}_{i,\omega} = 0,$$

these coincide if and only if $\tilde{\rho}^*_\omega\, p^*_\omega = \nu^*_\omega$, which holds by (107). The conditions for $x^-_{i\omega}$, $y^+_{j\omega}$, and $y^-_{j\omega}$ match by the same argument.

Day-ahead variables. The agent's condition (88) (with $\tilde{\rho}^c_{i,\omega} = \tilde{\rho}^*_\omega$) reads:

$$\sum_\omega \tilde{\rho}^*_\omega v_{i\omega} - \Big(\sum_\omega \tilde{\rho}^*_\omega\Big) p_0^* + \sum_\omega (\mu^{\mathrm{c}}_{i,\omega} - \underline{\mu}^{\mathrm{c}}_{i,\omega}) - \sigma^{\mathrm{c}}_i = 0.$$

Since $\sum_\omega \tilde{\rho}^*_\omega = (1-\varepsilon) + \varepsilon = 1$, this simplifies to:

$$\sum_\omega \tilde{\rho}^*_\omega v_{i\omega} - p_0^* + \sum_\omega (\mu^{\mathrm{c}}_{i,\omega} - \underline{\mu}^{\mathrm{c}}_{i,\omega}) - \sigma^{\mathrm{c}}_i = 0.$$

Comparing with the planner's condition (71), these coincide if and only if $\lambda^* = p_0^*$, which holds by (107). The producer's day-ahead condition matches analogously.

Since each agent's problem is convex with the same feasible region that the planner imposes on that agent's variables, the KKT conditions are necessary and sufficient. Therefore, at prices $(p_0^*, \boldsymbol{p}^*, q^*)$, every agent maximises their risk-adjusted utility at the planner's allocation (for any choice of insurance transfers consistent with complementary slackness). We illustrate a construction below that provides one such choice.

Remark 7. *We construct explicit insurance transfers satisfying* (95). *Define each agent's loss at the planner's solution:*

$$L^c_{i\omega} := -u_i(x_i^*, \boldsymbol{X}_i^*, p_0^*, \boldsymbol{p}^*; \omega), \qquad L^g_{j\omega} := -u_j(y_j^*, Y_j^*, p_0^*, \boldsymbol{p}^*; \omega).$$

By Walras' law (established below), the price-dependent terms cancel when summing over all agents:

$$\sum_{i\in[n]} L^c_{i\omega} + \sum_{j\in[m]} L^g_{j\omega} = \ell^*(\omega) := -W2(\boldsymbol{x}^*, \boldsymbol{y}^*; \omega), \qquad \forall\, \omega \in \Omega. \tag{110}$$

Choose constants δ^c_i and δ^g_j (independent of ω) satisfying $\sum_i \delta^c_i + \sum_j \delta^g_j = 0$, and define:

$$r^{c\ *}_{i\omega} := -L^c_{i\omega} + \tfrac{1}{n+m}\,\ell^*(\omega) + \delta^c_i, \tag{111}$$

$$r^{g\ *}_{j\omega} := -L^g_{j\omega} + \tfrac{1}{n+m}\,\ell^*(\omega) + \delta^g_j. \tag{112}$$

Insurance market clearing (95) *holds by* (110)*:*

$$\sum_i r^{c\ *}_{i\omega} + \sum_j r^{g\ *}_{j\omega} = -\ell^*(\omega) + \ell^*(\omega) + 0 = 0.$$

Moreover, each agent's insured loss becomes:

$$L^c_{i\omega} + r^{c\ *}_{i\omega} = \tfrac{1}{n+m}\,\ell^*(\omega) + \delta^c_i, \tag{113}$$

which is an affine increasing function of the aggregate planner loss $\ell^*(\omega)$ with the same positive coefficient for all agents. Consequently, all agents share the same set of tail scenarios (those in which $\ell^*(\omega)$ exceeds its $(1-\beta)$-quantile), ensuring that the complementary slackness conditions on the CVaR constraints coincide across all agents and match the planner's (77).

Proof of (WL): Walras' Law. The day-ahead and real-time Walras' law conditions (104)–(105) follow from complementary slackness on the planner's balance constraints (60)–(61), exactly as in the risk-neutral case (Theorem 1):

$$\lambda^*\Big(\sum_j y_j^* - \sum_i x_i^*\Big) = 0 \implies p_0^* \cdot z_0 = 0,$$

since $p_0^* = \lambda^*/(1+\varepsilon)$, and similarly for each ω. The insurance Walras' law (106) holds by construction: the transfers (111)–(112) satisfy exact clearing (95), so the left-hand side of (106) is $q^*_\omega \cdot 0 = 0$. □

Remark 8 (Price interpretation). *The price formulas* (107) *admit a natural economic interpretation. The day-ahead price* $p_0^* = \lambda^*$ *is the planner's marginal value of day-ahead energy. The real-time price* $p_\omega^* = \nu_\omega^* / \tilde{\rho}_\omega^*$ *is the marginal value of real-time energy in scenario* ω*, normalised by the CVaR-tilted weight rather than the raw probability; this reduces to* ν_ω^*/ρ_ω *when* $\varepsilon = 0$*. The insurance price* $q_\omega^* = \tilde{\rho}_\omega^* = (1-\varepsilon)\rho_\omega + \gamma_\omega^*$ *is the standard* Arrow–Debreu state price*: the cost today of one unit of consumption tomorrow contingent on scenario* ω*. It decomposes into a fair-pricing component* $(1-\varepsilon)\rho_\omega$ *and a tail-risk premium* γ_ω^* *that is strictly positive only on scenarios in the CVaR tail. Because* $\sum_\omega q_\omega^* = 1$*, the insurance prices form a probability measure on* Ω *– the equilibrium worst-case (or risk-neutral pricing) measure under which the planner's allocation can be evaluated as a risk-neutral optimum.*

Remark 9 (Insurance transfers). *The construction* (111)–(112) *admits a clear interpretation: each agent's loss is pooled, and a pro-rata share of the aggregate loss* $\ell^*(\omega)$ *is redistributed. The constants* δ_i^c, δ_j^g *represent budget-neutral lump-sum transfers that can be chosen freely (subject to summing to zero) to satisfy any requirements without affecting the equilibrium conditions. This construction is explicit, in contrast to the pure existence arguments in [6] and [11].*

8.6 Welfare Theorems for Risk-Aware SCER

Lemma 4 (Common worst-case measure at any R-SCER). *Let* $(\boldsymbol{x}^*, \boldsymbol{y}^*, \boldsymbol{X}^*, \boldsymbol{Y}^*, r^{c*}, r^{g*}, p_0^*, \boldsymbol{p}^*, q^*)$ *be an R-SCER, and let* $\gamma_{i,\omega}^c, \gamma_{j,\omega}^g$ *be the multipliers on the agents' CVaR shortfall constraints* (97), (101)*. Define*

$$\hat{Q}^*(\omega) \;:=\; q_\omega^*, \qquad \forall\, \omega \in \Omega.$$

Then $\hat{Q}^*$ *is a probability measure on* Ω *that lies in the risk set* $\mathcal{Q}_{RA}$ *of* ρ_{RA}*, and:*

(1) $\hat{Q}^*$ *is a worst-case measure for every agent's risk-adjusted utility: for each consumer* i *and each producer* j*,*

$$\rho_{RA}\big[-u_i(\cdot\,;\omega) - r_{i\omega}^{c\;*}\big] = \mathop{\mathbb{E}}_{\hat{Q}^*}\big[-u_i(\cdot\,;\omega) - r_{i\omega}^{c\;*}\big], \tag{114}$$

$$\rho_{RA}\big[-u_j(\cdot\,;\omega) - r_{j\omega}^{g\;*}\big] = \mathop{\mathbb{E}}_{\hat{Q}^*}\big[-u_j(\cdot\,;\omega) - r_{j\omega}^{g\;*}\big]. \tag{115}$$

(2) $\hat{Q}^*$ *is a worst-case measure for the aggregate planner loss:*

$$\rho_{RA}\big[\ell^*(\omega)\big] \;=\; \mathop{\mathbb{E}}_{\hat{Q}^*}\big[\ell^*(\omega)\big], \tag{116}$$

where $\ell^*(\omega) := -W2(\boldsymbol{x}^*, \boldsymbol{y}^*; \omega)$*.*

Proof. By insurance stationarity (108), the equilibrium first-order condition with respect to $r_{i\omega}^c$ gives $q_\omega^* = (1-\varepsilon)\rho_\omega + \gamma_{i,\omega}^c$, and analogously $q_\omega^* = (1-\varepsilon)\rho_\omega + \gamma_{j,\omega}^g$ for each producer. Hence every agent's CVaR multiplier is the same:

$$\gamma_{i,\omega}^c = \gamma_{j,\omega}^g = q_\omega^* - (1-\varepsilon)\rho_\omega =: \gamma_\omega^\dagger, \qquad \forall\, i, j, \omega.$$

The agents' threshold stationarity gives $\sum_\omega \gamma_{i,\omega}^c = \varepsilon$, so $\sum_\omega \gamma_\omega^\dagger = \varepsilon$, and consequently

$$\sum_\omega \hat{Q}^*(\omega) \;=\; \sum_\omega q_\omega^* \;=\; \sum_\omega \big[(1-\varepsilon)\rho_\omega + \gamma_\omega^\dagger\big] \;=\; (1-\varepsilon) + \varepsilon \;=\; 1,$$

so $\hat{Q}^*$ is a probability measure. The agents' slack stationarity gives $\gamma_\omega^\dagger \le \varepsilon\rho_\omega/(1-\beta)$. Defining $Q^*(\omega) := \gamma_\omega^\dagger/\varepsilon$, we have $\sum_\omega Q^*(\omega) = 1$ and $Q^*(\omega) \le \rho_\omega/(1-\beta)$, so $Q^* \in \mathcal{Q}_\beta$ and $\hat{Q}^* = (1-\varepsilon)P + \varepsilon Q^* \in \mathcal{Q}_{\mathrm{RA}}$.

To establish (114), observe that the agent's CVaR shortfall multipliers $\gamma^c_{i,\omega} = \gamma^\dagger_\omega = \varepsilon\, Q^*(\omega)$. This identifies Q^* as a worst-case measure in the dual representation for the CVaR applied to consumer i's insured loss $-u_i - r^c_{i\omega}{}^*$ (Lemma 3). Combining with the trivial identification of P as the worst-case measure for $\mathbb{E}[\cdot]$, Lemma 2 yields $\hat{Q}^*$ as a worst-case measure for ρ_{RA} applied to the agent's insured loss. The producer case (115) is identical. For the planner case (116), the same argument applies with the planner's CVaR multipliers γ^*_ω in place of $\gamma^\dagger_\omega$; that these coincide ($\gamma^*_\omega = \gamma^\dagger_\omega$) is precisely what the price identification $q^*_\omega = \tilde{\rho}^*_\omega$ delivers, and is verified directly through the planner's KKT in the proof of Theorem 5 below. □

The lemma states that at equilibrium, all market participants and the planner share a common worst-case measure $\hat{Q}^*$, and that this measure is given directly by the equilibrium insurance prices: $\hat{Q}^*(\omega) = q^*_\omega$. The Arrow–Debreu prices for state-contingent claims are not just any pricing system but are specifically the worst-case probability measure under which the entire risk-aware market behaves as risk-neutral. Under this common measure, every coherent risk measure in sight collapses to a linear expectation, restoring the additivity that makes the risk-neutral aggregation argument work.

Having established the existence of R-SCER via Theorem 4, we now state the welfare theorems for the risk-aware setting.

Theorem 5 (First Welfare Theorem for R-SCER). *If* $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*, r^{c*}, r^{g*}, p^*_0, \boldsymbol{p}^*, q^*)$ *is an R-SCER, then* $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*)$ *is a risk-adjusted welfare-maximising allocation, that is, it solves* (66).

Proof. The argument parallels the proof of Theorem 2 for the risk-neutral case, with the dual representation of ρ_{RA} playing the role that linearity of expectation played in the risk-neutral setting.

Since the market outcome is an R-SCER, condition **(UM)** gives that each consumer i is utility-maximising at (x^*_i, X^*_i, r^{c*}_i) and each producer j is utility-maximising at (y^*_j, Y^*_j, r^{g*}_j), given prices $(p^*_0, \boldsymbol{p}^*, q^*)$. Let $\mathcal{A}$ denote the programme obtained by summing the agents' risk-adjusted utility-maximisation objectives and aggregating their individual constraints. Since each agent's problem is solved independently at fixed prices, the equilibrium allocation $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*, r^{c*}, r^{g*})$ is an optimal solution to $\mathcal{A}$.

We show that $\mathcal{A}$ and the planner's problem (66) have the same optimal allocations, by verifying that they share the same objective and feasible set.

Objective. Each consumer i's objective at the equilibrium can be rewritten using ρ_{RA}:

$$(1-\varepsilon)\sum_\omega \rho_\omega\,(u_i + r^c_{i\omega}) - \sum_\omega q^*_\omega r^c_{i\omega} - \varepsilon\,\mathrm{CVaR}_\beta[-u_i - r^c_{i\omega}] \;=\; -\rho_{\mathrm{RA}}\big[-u_i - r^c_{i\omega}\big] - \sum_\omega q^*_\omega r^c_{i\omega}.$$

By Lemma 4, the worst-case measure $\hat{Q}^*$ collapses the risk measure to a linear expectation: $\rho_{\mathrm{RA}}[-u_i - r^c_{i\omega}] = \mathbb{E}_{\hat{Q}^*}[-u_i - r^c_{i\omega}]$. Therefore consumer i's objective equals

$$\mathop{\mathbb{E}}_{\hat{Q}^*}[u_i + r^c_{i\omega}] - \sum_\omega q^*_\omega r^c_{i\omega} \;=\; \mathop{\mathbb{E}}_{\hat{Q}^*}[u_i] + \sum_\omega \hat{Q}^*(\omega)\, r^c_{i\omega} - \sum_\omega q^*_\omega r^c_{i\omega} \;=\; \mathop{\mathbb{E}}_{\hat{Q}^*}[u_i],$$

where the final equality uses the identification $\hat{Q}^*(\omega) = q^*_\omega$ from Lemma 4: the expected payout of the insurance position under the worst-case measure exactly cancels the premium paid for it. The producer case is identical, with u_j in place of u_i.

Summing over all agents:

$$\sum_{i\in[n]} \mathop{\mathbb{E}}_{\hat{Q}^*}[u_i] + \sum_{j\in[m]} \mathop{\mathbb{E}}_{\hat{Q}^*}[u_j] \;=\; \mathop{\mathbb{E}}_{\hat{Q}^*}\Big[\sum_i u_i + \sum_j u_j\Big].$$

By the day-ahead and real-time balance constraints from **(F)**, the price-dependent terms in the quasi-linear utilities cancel when summed across consumers and producers, leaving $\sum_i u_i + \sum_j u_j = \mathrm{W2}(\boldsymbol{x}^*, \boldsymbol{y}^*; \omega) = -\ell^*(\omega)$. Hence the aggregated objective of $\mathcal{A}$ becomes

$$\mathop{\mathbb{E}}_{\hat{Q}^*}\big[-\ell^*(\omega)\big] \;=\; -\mathop{\mathbb{E}}_{\hat{Q}^*}\big[\ell^*(\omega)\big] \;=\; -\rho_{\mathrm{RA}}\big[\ell^*(\omega)\big],$$

where the last equality uses (116). Up to the sign convention (the planner minimises $\rho_{\mathrm{RA}}[\ell^*]$ while $\mathcal{A}$ maximises its negative), this is precisely the planner's objective in (66).

Constraints. The bound constraints (63), (64) and non-negativity (65) appear identically in each agent's individual problem and in the planner's problem. The market-clearing constraints (60), (61) and the insurance-clearing constraint (95) hold by condition **(F)**, and the planner's CVaR shortfall and threshold constraints (62) are satisfied with primal variables $t^* := \sum_k t_k^{c*}$ and $u_\omega{}^* := [\ell^*(\omega) - t^*]^+$.

Since $\mathcal{A}$ and the planner's problem share the same objective and feasible set, any optimal solution to $\mathcal{A}$ is also optimal for the planner. Therefore $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*)$ solves (66). □

Theorem 6 (Second Welfare Theorem for R-SCER). *If $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*)$ is a feasible welfare-maximising allocation for the risk-aware planner problem* (66)*, then there exist prices $(p_0^*, \boldsymbol{p}^*, q^*)$ and insurance transfers (r^{c*}, r^{g*}) such that the market outcome $(\boldsymbol{x}^*, \boldsymbol{y}^*, X^*, Y^*, r^{c*}, r^{g*}, p_0^*, \boldsymbol{p}^*, q^*)$ is an R-SCER.*

Proof. This is a direct consequence of Theorem 4. Equilibrium prices are constructed from the dual multipliers of the risk-aware planner problem via the identification (107), and the insurance transfers are constructed via (111)–(112). Theorem 4 establishes that all three R-SCER conditions are satisfied. □

Remark 10 (Risk-neutral and pure-CVaR limits). *When $\varepsilon = 0$, the CVaR penalty vanishes: $\gamma_\omega^* = 0$ for all ω, $\tilde{\rho}_\omega^* = \rho_\omega$, and $q_\omega^* = \rho_\omega$. The price identifications reduce to $p_0^* = \lambda^*$ and $p_\omega^* = \nu_\omega^*/\rho_\omega$, recovering the risk-neutral formulas* (43) *for energy. The insurance market is priced at $q_\omega^* = \rho_\omega$. So, the premium $\sum_\omega q_\omega^* r_{i\omega}^c$ exactly equals the expected payout $\sum_\omega \rho_\omega r_{i\omega}^c$. Hence, in this case, the agents are indifferent, and the insurance market is redundant. The R-SCER reduces to the SCER of Definition 2 on the physical allocation, and the risk-neutral welfare theorems (Theorems 2–3) are special cases of Theorems 5–6.*

At the opposite extreme $\varepsilon = 1$, the planner optimises pure CVaR. The expected-utility component drops from the agent's objective, and $q_\omega^* = \gamma_\omega^*$ since $(1-\varepsilon)\rho_\omega = 0$. All probability mass in $\hat{Q}^*$ comes from the CVaR tail, with $\hat{Q}^*$ supported on the worst $(1-\beta)$-fraction of scenarios. The Arrow–Debreu prices then concentrate entirely on tail scenarios, and trades in non-tail scenarios are free.

8.7 Existence of R-SCER under heterogeneous risk-preferences

The results of Sections 8.5.1–8.6 assume that all agents share common risk parameters (ε, β). In practice, market participants may differ in their attitudes toward tail risk. A renewable generator exposed to weather uncertainty may adopt a different risk posture than a dispatchable thermal plant. In this section, we characterise precisely which forms of heterogeneity are compatible with competitive equilibrium and which are not.

8.7.1 Heterogeneous β and common ϵ. We now show that heterogeneous confidence levels β_i and β_j can be supported with a competitive equilibrium, provided the social planner adopts the appropriate confidence level.

Suppose each consumer i uses confidence level $\beta_i \in [0, 1)$ and each producer j uses $\beta_j \in [0, 1)$, while all agents (and the planner) share a common $\varepsilon \in [0, 1]$. Define

$$\beta_{\min} := \min\big\{ \min_{i\in[n]} \beta_i,\ \min_{j\in[m]} \beta_j \big\}, \tag{117}$$

and partition the set of agents into

$$S_{\min} := \{k : \beta_k = \beta_{\min}\}, \qquad S_{>} := \{k : \beta_k > \beta_{\min}\}.$$

Insurance stationarity (108) continues to give $q_\omega = (1-\varepsilon)\rho_\omega + \gamma^c_{k,\omega}$, so the agent's CVaR multiplier is $\gamma^c_{k,\omega} = q_\omega - (1-\varepsilon)\rho_\omega$. Summing over ω and using the agent's threshold stationarity $\sum_\omega \gamma^c_{k,\omega} = \varepsilon$ gives $\sum_\omega q_\omega = 1$ – the standard probability-measure constraint, common to every agent regardless of β_k. The slack stationarity, on the other hand, depends on β_k:

$$\gamma^c_{k,\omega} + \eta^c_{k,\omega} = \frac{\varepsilon\,\rho_\omega}{1-\beta_k}, \qquad \eta^c_{k,\omega} \geq 0, \tag{118}$$

giving the upper bound $\gamma^c_{k,\omega} \leq \varepsilon\,\rho_\omega/(1-\beta_k)$, equivalently

$$q_\omega \leq (1-\varepsilon)\,\rho_\omega + \frac{\varepsilon\,\rho_\omega}{1-\beta_k}. \tag{119}$$

The right-hand side is decreasing in β_k, so the tightest constraint is imposed by the agent with smallest β. The binding constraint reads

$$q_\omega \leq (1-\varepsilon)\,\rho_\omega + \frac{\varepsilon\,\rho_\omega}{1-\beta_{\min}}, \qquad \forall\,\omega \in \Omega, \tag{120}$$

which is precisely the constraint that arises from a planner using confidence level $\beta_{\min}$.

A critical structural consequence of (118) is that for any agent $k \in S_>$ (with $\beta_k > \beta_{\min}$), the slack dual ${\eta^c}_{k,\omega} > 0$ for every scenario — including tail scenarios — because $\gamma^c_{k,\omega} \leq \varepsilon\rho_\omega/(1-\beta_{\min}) < \varepsilon\rho_\omega/(1-\beta_k)$. Complementary slackness then forces ${\zeta^c_{k,\omega}}^* = 0$ for all ω. When the equilibrium tail multiplier $\gamma^c_{k,\omega} > 0$ as well, the CVaR shortfall constraint must bind with zero slack:

$$L^c_{k\omega} - r^*_{k\omega} = t^{c*}_k, \qquad \forall\,k \in S_>,\ \forall\,\omega \in \Omega \text{ with } \gamma^c_{k,\omega} > 0. \tag{121}$$

Agents in $S_>$ require a *constant* insured loss across the entire tail. This rules out the equal-share comonotone construction (111)–(112), in which insured losses vary with $\ell^*(\omega)$, and necessitates a different transfer construction.

Theorem 7 (Existence under heterogeneous β). *Given common $\varepsilon \in [0,1]$ and heterogeneous confidence levels $\{\beta_i\}_{i\in[n]}$, $\{\beta_j\}_{j\in[m]}$, let $(\boldsymbol{x}^*, \boldsymbol{y}^*, \boldsymbol{X}^*, \boldsymbol{Y}^*, t^*, u^*)$ solve the risk-aware planner problem* (66) *with confidence level $\beta_{\min}$ as defined in* (117)*. Let $(\lambda^*, \nu^*, \gamma^*, \eta^*)$ be the associated optimal dual variables. Define prices via* (107)*: $p^*_0 = \lambda^*$, $p^*_\omega = \nu^*_\omega/\tilde{\rho}^*_\omega$, $q^*_\omega = \tilde{\rho}^*_\omega = (1-\varepsilon)\rho_\omega + \gamma^*_\omega$.*

Then there exist insurance transfers ${r^c_{i\omega}}^*$ and ${r^g_{j\omega}}^*$ such that the market outcome is an R-SCER.

Proof. The price identification $q^*_\omega = \tilde{\rho}^*_\omega$ aligns each agent's tilted weight with the planner's, and β_k enters only through the slack dual $\eta^c_{k,\omega}$, which does not couple to physical variables. Dual feasibility $\eta^c_{k,\omega} \geq 0$ for each agent k follows from $\gamma^*_\omega \leq \varepsilon\rho_\omega/(1-\beta_{\min}) \leq \varepsilon\rho_\omega/(1-\beta_k)$.

Construction of insurance transfers. We construct one possible transfer that respect the structural requirement (121). Define each agent's per-scenario loss $L_{k\omega}$ at the equilibrium prices and the aggregate loss $\ell^*(\omega)$ as before, and set $t^* := \sum_k t^{c*}_k$. For each agent k, define:

$$r^*_{k\omega} = L_{k\omega} - t^{c*}_k - w_k\big(\ell^*(\omega) - t^*\big), \tag{122}$$

where the *risk-bearing weights* $w_k \geq 0$ satisfy $\sum_k w_k = 1$ and are chosen as:

$$w_k = \begin{cases} 0, & \text{if } k \in S_>, \\ 1/|S_{\min}|, & \text{if } k \in S_{\min}. \end{cases} \tag{123}$$

Market clearing. We verify (95):

$$\sum_k r^*_{k\omega} = \sum_k L_{k\omega} - \sum_k t_k^{c*} - \Big(\sum_k w_k\Big)\big(\ell^*(\omega) - t^*\big) = \ell^*(\omega) - t^* - \big(\ell^*(\omega) - t^*\big) = 0.$$

Insured losses. Each agent's insured loss equals

$$L_{k\omega} - r^*_{k\omega} = t_k^{c*} + w_k\big(\ell^*(\omega) - t^*\big). \tag{124}$$

For agents $k \in S_>$: $w_k = 0$, so the insured loss equals the constant t_k^{c*}, satisfying (121). For agents $k \in S_{\min}$: the insured loss is an affine increasing function of $\ell^*(\omega)$ with positive coefficient, so these agents share a common tail set with the planner.

CVaR complementary slackness. For agents $k \in S_>$: $\eta^c_{k,\omega} > 0$ for all ω, hence $\zeta^{c\;*}_{k,\omega} = 0$, and the shortfall constraint reads $0 = L_{k\omega} - r^*_{k\omega} - t_k^{c*} = 0$, holding with equality. Both $\gamma^c_{k,\omega} \cdot 0 = 0$ and $\eta^c_{k,\omega} \cdot 0 = 0$ are satisfied. For agents $k \in S_{\min}$: the argument is identical to the homogeneous case (Theorem 4), since these agents share the planner's confidence level $\beta_{\min}$. □

9 Experiments and Results

9.1 Synthetic Experiments

We first evaluate the risk-aware formulation in a controlled synthetic market with ten generators and one inelastic demand. Six generators are modeled as low-cost unreliable resources, representing renewable supply, while four generators are modeled as higher-cost reliable resources. Motivated by Huisman et al. [8], who show that variable renewables can generate heavy-tailed market outcomes, we model renewable uncertainty through common supply shock states: normal, mild shortfall, and severe shortfall. These shocks are applied across unreliable generators, inducing correlated low-renewable scenarios. The resulting market setup is summarized in Table 1.

Table 1. Synthetic experiment setup.

Component	Specification
Number of generators	10
Unreliable generators	6
Reliable generators	4
Scenario probabilities	Dirichlet near-uniform, $\kappa = 1500$
Unreliable generator costs	Uniform on $[8, 14]$/MWh
Reliable generator costs	Uniform on $[35, 55]$/MWh
Renewable shock states	normal, mild shortfall, severe shortfall
Cost to curtail demand	\$1000/MWh

Figure 1 shows that this construction produces an adverse low-welfare scenarios resulting from renewable shortfalls. We then compare deterministic clearing, risk-neutral stochastic clearing, and the CVaR-augmented formulation. The risk-neutral stochastic formulation achieves the highest expected social surplus, reflecting its objective of optimizing average social surplus across scenarios. In contrast, the CVaR-augmented formulation improves mean welfare in the worst 5% of scenarios, indicating that risk aversion trades off some average efficiency for better protection against adverse tail outcomes, as shown in Figure 2.

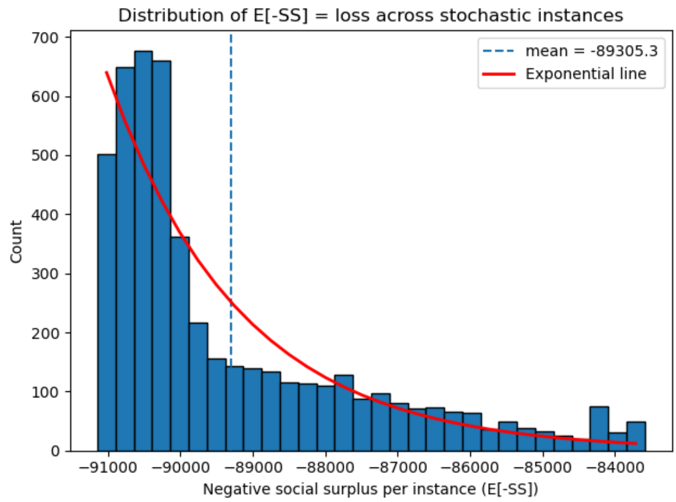

Fig. 1. Heavy-tailed welfare distribution induced by correlated renewable shortfall states in the synthetic experiment.

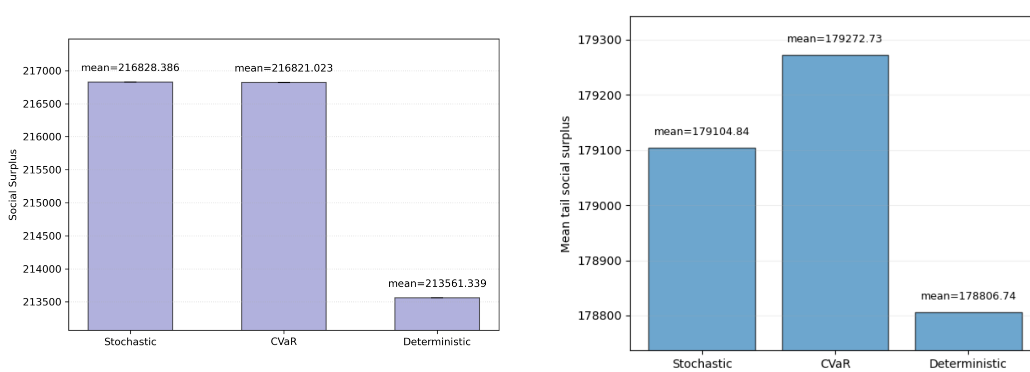

(a) Expected social surplus

(b) Mean social surplus in the worst 5% of scenarios

Fig. 2. Synthetic experiment results. The CVaR-augmented formulation is evaluated against deterministic and risk-neutral stochastic clearing using average welfare and tail welfare.

9.2 Empirical evaluation on IM3 GO WEST

We next evaluate the same clearing rules on scenarios constructed from the IM3 GO WEST dataset [1], which provides time-indexed renewable generation and load profiles for a Western Interconnection test system. We use solar and wind profiles as unreliable generators and thermal units as reliable generators. To test robustness across system sizes, we vary the number of renewable generators while holding the number of thermal generators fixed.

Table 2 shows that the CVaR-augmented formulation achieves the highest tail welfare across all tested configurations. The deterministic benchmark performs worst because its day-ahead schedule is optimized for a single forecasted system state, whereas the stochastic and CVaR formulations choose day-ahead commitments while anticipating real-time recourse across all scenarios. The

Table 2. Mean social surplus in the worst 5% of IM3 GO WEST scenarios.

Configuration	Stochastic	CVaR	Deterministic
3 wind, 3 solar, 4 thermal	2.580637×10^7	2.581649×10^7	2.578110×10^7
9 wind, 9 solar, 4 thermal	2.585322×10^7	2.588694×10^7	2.582500×10^7
20 wind, 20 solar, 4 thermal	2.642656×10^7	2.647268×10^7	2.627681×10^7

gap between CVaR and risk-neutral stochastic clearing also increases in the larger renewable configuration, suggesting that the value of risk-aware clearing becomes more pronounced as exposure to renewable uncertainty grows.

10 Conclusions

In this work, we design a centralised mechanism for the electricity market that a risk-aware social planner can adopt, which delivers improved tail-welfare guarantees relative to its risk-neutral counterpart. To sustain a stochastic competitive equilibrium in a decentralised market populated by risk-aware agents, we show that it is necessary to introduce Arrow-Debreu securities, an "insurance market" with trading instruments indexed to every scenario, thereby completing the market. Our analysis demonstrates that the day-ahead, real-time, and insurance prices can all be obtained as dual variables of the risk-aware social-welfare optimisation problem solved by the planner, yielding a unified pricing framework that simultaneously reflects expected energy value, temporal flexibility across the two stages, and the cost of bearing risk in the tail of the scenario distribution. Further, we experimentally show how market-clearing using this risk-aware mechanism can yield better welfare in the tail events, thereby promising a more robust mechanism. Our experiment validates this using both synthetic experiments and data collected from a Western Interconnection test system.

10.1 Future Work

Our complete-markets construction relies on a full set of Arrow–Debreu securities, an idealisation that is difficult to implement in real electricity markets. The most natural extensions concern *market incompleteness* is characterising the welfare loss when only a small number of risk instruments are available, identifying minimal sets of instruments that approximately decentralise the planner, and resolving the multiplicity of equilibria identified by Gérard et al. [6] when markets are incomplete.

References

[1] K. Z. Akdemir, K. Oikonomou, J. Kern, and N. Voisin. 2023. *IM3 GO WEST Parameter Search Dataset.* https://doi.org/10.57931/1923267

[2] Marcus Ang, Jie Sun, and Qiang Yao. 2018. On the dual representation of coherent risk measures. *Annals of Operations Research* 262, 1 (2018), 29–46.

[3] Kenneth J Arrow. 1964. The role of securities in the optimal allocation of risk-bearing. *The review of economic studies* 31, 2 (1964), 91–96.

[4] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. 1999. Coherent Measures of Risk. *Mathematical Finance* 9, 3 (1999), 203–228.

[5] Yury Dvorkin. 2019. A chance-constrained stochastic electricity market. *IEEE Transactions on Power Systems* 35, 4 (2019), 2993–3003.

[6] Henri Gérard, Vincent Leclère, and Andy Philpott. 2018. On risk averse competitive equilibrium. *Operations Research Letters* 46, 1 (2018), 19–26.

[7] Ye Guo, Cong Chen, and Lang Tong. 2021. Pricing multi-interval dispatch under uncertainty part I: Dispatch-following incentives. *IEEE Transactions on Power Systems* 36, 5 (2021), 3865–3877.

[8] Ronald Huisman, Evangelos Kyritsis, and Christiaan Stet. 2022. Fat Tails due to Variable Renewables and Insufficient Flexibility: Evidence from Germany. *The Energy Journal* 43, 5 (2022), 231–247.

[9] Avinash N Madavan, Nathan Dahlin, Subhonmesh Bose, and Lang Tong. 2025. Risk-Sensitive Security-Constrained Economic Dispatch: Pricing and Algorithm Design. *IEEE Transactions on Power Systems* (2025).

[10] Robert Mieth, Matt Roveto, and Yury Dvorkin. 2020. Risk trading in a chance-constrained stochastic electricity market. *IEEE Control Systems Letters* 5, 1 (2020), 199–204.

[11] Andy Philpott, Michael Ferris, and Roger Wets. 2016. Equilibrium, uncertainty and risk in hydro-thermal electricity systems. *Mathematical Programming* 157, 2 (2016), 483–513.

[12] Geoffrey Pritchard, Golbon Zakeri, and Andrew Philpott. 2010. A single-settlement, energy-only electric power market for unpredictable and intermittent participants. *Operations research* 58, 4-part-2 (2010), 1210–1219.

[13] R. Tyrrell Rockafellar and Stanislav Uryasev. 2000. Optimization of conditional value-at-risk. *Journal of Risk* 2, 3 (2000), 21–41.

[14] Alexander Shapiro and Yorai Wardi. 1994. Nondifferentiability of the steady-state function in discrete event dynamic systems. *IEEE Trans. Autom. Control.* 39 (1994), 1707–1711. https://api.semanticscholar.org/CorpusID:123616717

[15] Tomás Tapia, Zhirui Liang, Charalambos Konstantinou, and Yury Dvorkin. 2024. Electricity market-clearing with extreme events. *IEEE Transactions on Energy Markets, Policy and Regulation* 3, 2 (2024), 194–208.

[16] Lucien Werner, Nicolas Christianson, Alessandro Zocca, Adam Wierman, and Steven Low. 2023. Pricing uncertainty in stochastic multi-stage electricity markets. In *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 1580–1587.

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

[18] Victor M Zavala, Kibaek Kim, Mihai Anitescu, and John Birge. 2017. A stochastic electricity market clearing formulation with consistent pricing properties. *Operations Research* 65, 3 (2017), 557–576.