

Market Design Considerations for Scarcity Pricing: A Stochastic Equilibrium Framework

Anthony Papavasiliou and Yves Smeers** and Gauthier de Maere d'Aertrycke****

ABSTRACT

Scarcity pricing is a mechanism for improving the valuation of reserve capacity in real-time electricity markets. The goal of scarcity pricing is to mitigate the missing money problem and enhance investment in flexible resources. The implementation of scarcity pricing is underway in a number of U.S. markets, including Texas and PJM. The implementation is also currently under consideration in Belgium. As the mechanism was originally conceived in the context of a U.S.-style two-settlement system, its implementation in a European setting poses a number of interesting market design dilemmas which can affect the back-propagation of scarcity prices to forward day-ahead markets for energy and reserve capacity. We propose a modeling framework for analyzing these market design choices based on stochastic equilibrium, and use this modeling framework in order to represent and analyze a wide range of market design proposals. We report results on a case study of the Belgian electricity market.

Keywords: Scarcity pricing, operating reserve demand curve, stochastic equilibrium

1. INTRODUCTION

Scarcity pricing is the principle of pricing electricity at a value above the marginal cost of the marginal unit during conditions of high system stress, according to the incremental value that flexible capacity offers to the system in terms of keeping loss of load probability in check. Concretely, scarcity pricing is implemented by including an adder to the real-time price on top of the marginal cost of the marginal unit, and by rewarding that same adder to standby reserve capacity. The effect of this mechanism is that (i) it rewards flexible resources for being available, even if not activated, and (ii) it rewards flexible resources for reacting to system imbalances when the system is short on flexible capacity. Through economic arbitrage between generation and reserve capacities in real time and day ahead, scarcity pricing creates the potential of giving rise to a long-term investment signal for building flexible capacity or mobilizing demand response that can deliver security to the system.

Numerical analyses of the Belgian market (Papavasiliou and Smeers, 2017; Papavasiliou et al., 2018) have demonstrated the potential of scarcity pricing to restore the financial viability of flexible technologies in Belgium, and also to create a strong investment signal for mobilizing demand response. In response to these encouraging indicators about the potential of scarcity pricing to attract flexibility in the Belgian market, the present paper discusses concrete market design measures that would enable scarcity pricing to function effectively in the context of the Belgian market design.

We propose a range of increasingly disruptive measures for the evolution of Belgian market design that would enable scarcity pricing to deliver its intended benefits. The resulting electricity market design is analyzed using a stochastic equilibrium framework. Stochastic equilibrium allows

*Corresponding author. Center for Operations Research and Econometrics, UCLouvain. E-mail: anthony.papavasiliou@uclouvain.be.

**Center for Operations Research and Econometrics, UCLouvain.

***Center for Economic Expertise, Modelling, and Studies, Tractebel/ENGIE.

us to quantify the impact of scarcity pricing on real-time energy and reserve prices, as well as the back-propagation of these prices to the day-ahead market through arbitrage. On the basis of our analysis, our concrete recommendation to the Belgian regulator is to proceed with the introduction of a real-time market for reserve capacity in Belgium.

2. CONTEXT

2.1 Principles of Scarcity Pricing

Scarcity pricing has been proposed as an approach for the precise valuation of reserve services (Hogan, 2005). The principle of scarcity pricing is to add a correction to the real-time price which rewards generators that can respond rapidly when balancing the system. The theoretical justification of the approach is that it adjusts the real-time price of energy and reserve capacity such that the resulting dispatch of profit-maximizing generators would reproduce the optimal dispatch that would be obtained if the contribution of reserve capacity towards reducing the loss of load probability would be accounted for (Hogan, 2013).

The mechanism has an equivalent, intuitive interpretation in terms of an operating reserve demand curve (ORDC). The rationale for an ORDC interpretation can be developed as follows. Consider the marginal value of reserve capacity when there is very little capacity left. For example, in Texas, when the system reserve capacity drops below 2000 MW, the system operator follows a series of emergency actions, up to and including involuntarily curtailment of demand, in order to prevent cascading outages (ERCOT, 2019). Effectively, the marginal value of reserve capacity under these conditions is equal to the value of lost load, which in Texas is set administratively to 9000 \$/MWh (ERCOT, 2019). When abundant reserve capacity is available in the system (e.g. above 5000 MW in Texas (ERCOT, 2019)), the marginal value of capacity is equal to zero. For intermediate values, the marginal value of reserve depends on the loss of load probability¹, because a marginal increment in reserve capacity has a marginal effect on system welfare which is proportional to the loss of load probability². This reasoning gives rise to the introduction of a demand function for operating reserve capacity that the system operator submits to a multi-product auction that simultaneously clears energy and reserve in the market. The effect of this demand function is that, under conditions of scarcity in reserve capacity, it lifts the energy price by a scarcity adder, which also applies to reserve capacity. A simplified formula for this adder when there exists a single type of reserve is given by the following expression:

$$(VOLL - \lambda) \cdot LOLP(R). \tag{1}$$

The notation here is as follows: *VOLL* corresponds to the value of lost load, λ is a proxy of the marginal cost of the marginal unit, *R* is the amount of remaining reserve capacity, and *LOLP* is the loss of load probability. Note that, as the system becomes tight (*R* decreases), adding this term to the energy price tends to push the energy price to *VOLL*. The distinction with a pure energy-only market is that this occurs in a smooth and more predictable fashion, since it is not rare that the system reaches a level where *LOLP*(*R*) is non-zero, even if no load shedding occurs. When abundant capacity is available (*R* is very large), the adder dissipates and has no effect on the energy price.

In a two-settlement system, the scarcity adder directly impacts in real time those resources that can rapidly be dispatched upward: they receive the scarcity adder in addition to the marginal cost of the marginal unit. But this scarcity signal is not meant to only apply to real-time operations. Arbitrage between day-ahead and real-time markets back-propagates the scarcity signal to day-ahead markets, and thus creates a favorable environment for all resources that can offer reserve capacity.

¹The loss of load probability *LOLP*(*R*) is the probability of load shedding, given that the system is carrying *R* MW of reserve capacity. It can be expressed as $LOLP(R) = \mathbb{P}[Imb \geq R]$, i.e. as the probability that system imbalance *Imb* exceeds *R* MW of reserve capacity.

²A mathematically rigorous proof of this intuition is provided by Hogan (2013), see also Papavasiliou and Smeers (2017).

Such resources are inherently required in systems with significant shares of renewable power supply. With that being said, a notable difference between scarcity pricing and capacity mechanisms is the built-in “pay for performance” attribute of the scarcity pricing mechanism. Indeed, under scarcity pricing, the stress of the system is signaled by the real-time price which is enhanced by a scarcity adder, therefore it is in the best interest of resources to perform exactly when the system is most stressed (otherwise they pay for their shortfall in real time, or forgo profit opportunities). In a capacity mechanism, this performance attribute needs to be closely specified in the mechanism (by defining an ad-hoc de-rating of capacities depending on their characteristics, or penalties for unavailability during stress events) and requires ex-post monitoring of those performances.

2.2 Implementation of Scarcity Pricing

Scarcity pricing in the form of operating reserve demand curves has already been implemented in ERCOT (2015), and has recently been introduced in PJM (Hogan and Pope, 2019), see also FERC number EL19-58-000 and ER19-1486-000 of May 2020. In Europe, the Belgian transmission system operator computes and publishes scarcity prices one day after real-time market clearing, although these prices are not currently being used for settlement purposes. Furthermore, the Belgian transmission system operator has performed an ex-post simulation of how scarcity prices would have affected the Belgian market in 2017 (ELIA, 2018), based on telemetry data. In September 2020 the Belgian transmission system operator launched a public consultation on its findings regarding the design of a scarcity pricing mechanism for implementation in Belgium ELIA (2020), in response to the proposal set forth in the present paper.

Two legal documents that have recently been published by the European Commission and the European Parliament indicate a favorable view towards scarcity pricing. *Scarcity pricing* is referred to as *shortage pricing* in these documents. The articles in question are article 44, paragraph 3 of the Electricity Balancing Guideline³ (European Commission, 2017) and Article 20, paragraph 3 of the Clean Energy Package⁴ (European Union, 2019).

For a broader discussion of energy-only markets supported by scarcity pricing and capacity markets, as well as the ongoing debate in the European context, the reader is referred to paragraphs 1.1 and 1.2 of Papavasiliou and Smeers (2017).

2.3 U.S. Two-Settlement Systems and the European Market Design

Given these appealing features of scarcity pricing and the increasing experience that is accumulating through the adoption of the mechanism, the Belgian regulator issued a request in 2015 for investigating the potential impact of scarcity pricing on the financial viability of combined cycle gas turbines. The analysis (Papavasiliou and Smeers, 2017; Papavasiliou et al., 2018) concluded that scarcity pricing can have a tangible effect on the financial viability of CCGT units, in the sense that its introduction can shift these units from not being able to recover their long-run investment costs to being able to recover their long-run investment costs. The Belgian regulator subsequently requested from the

³“Each TSO may develop a proposal for an additional settlement mechanism separate from the imbalance settlement, to settle the procurement costs of balancing capacity pursuant to Chapter 5 of this Title, administrative costs and other costs related to balancing. The additional settlement mechanism shall apply to balance responsible parties. This should be preferably achieved with the introduction of a **shortage pricing function**. If TSOs choose another mechanism, they should justify this in the proposal. Such a proposal shall be subject to approval by the relevant regulatory authority.”

⁴“Member States with identified resource adequacy concerns shall develop and publish an implementation plan with a timeline for adopting measures to eliminate any identified regulatory distortions or market failures as a part of the State aid process. When addressing resource adequacy concerns, the Member States shall in particular take into account the principles set out in Article 3 and shall consider:

...
(c) introducing a **shortage pricing function** for balancing energy as referred to in Article 44(3) of Regulation 2017/2195;
...”

Belgian system operator to conduct a parallel run on how the scarcity adder would have evolved given the available reserve capacity that was recorded by ELIA telemetry in 2017 (ELIA, 2018). In addition, the Belgian regulator requested a concrete market design proposal for the specific changes that would be required in the design of the Belgian day-ahead and real-time market in order to permit scarcity prices to take effect, and back-propagate through financial arbitrage⁵ to the forward day-ahead market. This is the focus of the present publication. Although the original investigation of scarcity pricing in 2015 (Papavasiliou and Smeers, 2017; Papavasiliou et al., 2018) was motivated by the financial viability of CCGT units, the present publication is further concerned with the implications of the mechanisms for demand-side resources that can offer reserve services to the system.

The motivation behind the request of the Belgian regulator is the fact that the European electricity market design deviates to a certain extent from the U.S. two-settlement system for which scarcity pricing was conceived. We discuss some of the market design differences, and how they affect the behavior of scarcity pricing, in turn.

Real-time market for reserve capacity. Reserve capacity is a service that is traded in real time in U.S. markets. This is achieved through security constrained economic dispatch (SCED), and reserve capacity is consequently priced according to the effect of reserve requirements. In the context of scarcity pricing, an operating reserve demand curve can either be introduced explicitly in SCED, or implicitly by approximating the value of the real-time reserve capacity price through equation 1, as an ex post calculation. By contrast, most European markets, including Belgium, are limited to auctioning reserve capacity in the week-ahead or day-ahead. Bilateral trading of reserve capacity is subsequently allowed, and takes place to a limited extent as real time approaches. Come real time, the so-called balancing service providers that are required to make this reserve available are only required to make their pre-committed reserve capacity available at the beginning of a balancing interval. Any excess reserve capacity or shortage in reserve capacity that appears in real time is not explicitly priced or traded. The only thing that is priced in real time is balancing *energy*, these payments are limited to those resources that are activated for balancing, and the payment is reflective of fuel cost, as opposed to the marginal value of the capacity to system security. Thus, there is no mechanism for rewarding reserve capacity for simply being available (as opposed to being activated), and there is no signal that reflects the level of stress in the system, as measured by the amount of reserve capacity that is available in real time.

Co-optimization of energy and reserves in day-ahead markets. European day-ahead electricity markets trade energy separately from reserve capacity in a day-ahead energy exchange. This energy exchange covers the majority of the continental European region, including Belgium, and will soon be expanded to include a number of additional member states (e.g. Greece). Depending on the specific member state, reserve capacity is traded separately, either before the clearing of the energy exchange (e.g. Belgium, Germany) or after the clearing of the energy exchange (e.g. Spain, Italy). The Belgian setup of clearing reserve capacity before the energy exchange is reminiscent of the separate auctioning of reserve and energy, as implemented for example in the original California market design. Since scarcity pricing relies on the formation of reserve prices through the arbitrage between energy and reserve capacity, the choice of separating the auctioning of energy from the auctioning of reserve capacity introduces certain inaccuracies in the valuation of reserve capacity, because it relies on the precise anticipation of reserve capacity opportunity cost, which in itself relies on the precise anticipation of energy prices. By contrast, this opportunity cost is evaluated endogenously in a co-optimization of energy and reserves through a multi-product auction, which is the current practice

⁵Financial arbitrage refers to the practice of exploiting deviations between day-ahead and real-time prices in order to generate speculative profits. In intuitive terms, financial arbitrage works as follows: if the day-ahead price for energy is low, compared to the real-time price of energy, then traders will tend to take positions in the day ahead, and close them in real time, thereby exerting an upward pressure on day-ahead prices, and vice versa, until day-ahead prices converge to real-time prices.

in U.S. day-ahead power pools. Additional incentive effects such as price reversals of reserve products (Oren, 2001) do not concern us here, and we limit our investigation to the implications of this market design choice on ORDC adders, assuming that agents bid truthfully.

Virtual trading. Scarcity pricing is a real-time mechanism. The back-propagation of scarcity prices to a signal that can reward long-run investments in reserve capacity relies on the arbitrage between energy and reserve capacity, and on the arbitrage between forward prices and real-time prices. The former is discussed in the previous paragraph. The latter relies on the specific arrangements that are set up in markets in order to allow for market participants to align forward prices with real-time prices. Such arrangements include aligning as much as possible day-ahead and real-time auctions, allowing opportunity cost bidding in the day-ahead market, or enabling virtual trading. Virtual trading of energy is allowed in various U.S. markets (Hogan, 2016; Parsons et al., 2015), whereas it is not allowed for reserve capacity. In European markets, virtual trading of energy is not permitted explicitly. Nevertheless, it may take place implicitly, to a limited extent, since market participants bid portfolios in the day-ahead market. Since part of these portfolio bids entails forecasts of renewable supply or demand that is served by the entities that submit these bids, these forecasts constitute private information of the bidder. Thus, although not allowed for explicitly, virtual trading may be possible to a certain extent also in the European market design. Since scarcity pricing affects real-time energy and reserve prices, and since virtual trading permits the back-propagation of these prices to the day-ahead time frame, virtual trading has a role in the function of a scarcity pricing mechanism.


Other aspects. The real-time pricing of energy in European electricity markets is inconsistent on the buying and on the selling side. At the time of this writing, suppliers of balancing energy are paid a pay-as-bid price, whereas ‘consumers’ of balancing energy (i.e. resources that deviate from their day-ahead positions) pay a uniform imbalance price. We conduct our analysis under the assumption that these distortions are not present, and that real-time energy is priced as a single product.

As a legacy attribute of system operation, certain European system operators prefer to assign as much of the balancing responsibility to market participants as possible. Real-time operations are thus treated as a service for supporting the day-ahead market, which is viewed by the Belgian system operator as the spot market. This is in stark contrast to U.S. market design, where the real-time market is the spot market, whereas the day-ahead market functions as a forward market. A tangible consequence of the European viewpoint is that so-called balancing responsible parties which trade energy in the day-ahead market are expected to deviate as little as possible from their day-ahead positions, even if this is beneficial towards balancing the system. We attempt to model this aspect of real-time market design in our analysis.

The complex treatment of non-convex operating costs is out of scope for the present analysis. U.S. day-ahead auctions separate commitment and dispatch from pricing, and resort to uplifts for covering any inconsistencies that may emerge between the day-ahead price and the cleared quantities. By contrast, the European day-ahead market clearing platform determines prices and dispatch quantities within a single market clearing model. The day-ahead auction relies on so-called paradoxically rejected bids in order to address inconsistencies between cleared quantities and prices, according to which those bids that would suffer economic losses are not accepted, even if they could enhance welfare by being accepted. In our present analysis, we ignore issues that relate to non-convexities by assuming a simplified relaxation of unit commitment which still captures the economical significance of unit commitment as an irrevocable day-ahead decision that affects the real-time profitability of a conventional unit.

European day-ahead markets trade transmission capacity according to a zonal design. The representation of zonal markets is out of scope for the present analysis, since its precise modeling (Aravena and Papavasiliou, 2017) would obscure the main messages of our analysis. Consequently, we ignore transmission constraints in the present paper.

Figure 1: The chain of designs that have been considered in our analysis.



2.4 Research Goal and Outline

Taking stock of the aforementioned differences between European and U.S. market design, we propose a family of models that approximate a possible evolution of Belgian market design towards a U.S.-style two-settlement system. In analyzing this evolution, we propose concrete changes to European real-time and day-ahead market design that are increasingly disruptive, and discuss the implications of these changes on the effectiveness of scarcity pricing. The effectiveness of scarcity pricing is assessed on the basis of the profitability of flexible resources, which include both fast-moving CCGT units as well as demand-side resources that can offer reserve services to the system. These are the resources that are most significantly affected by the introduction of scarcity pricing.

The evolution that we investigate in our analysis is depicted graphically in figure 1. We briefly summarize the chain of evolutions as follows:

(i) We embark on our analysis under the assumption that suppliers of balancing energy face the same price as consumers of balancing energy.

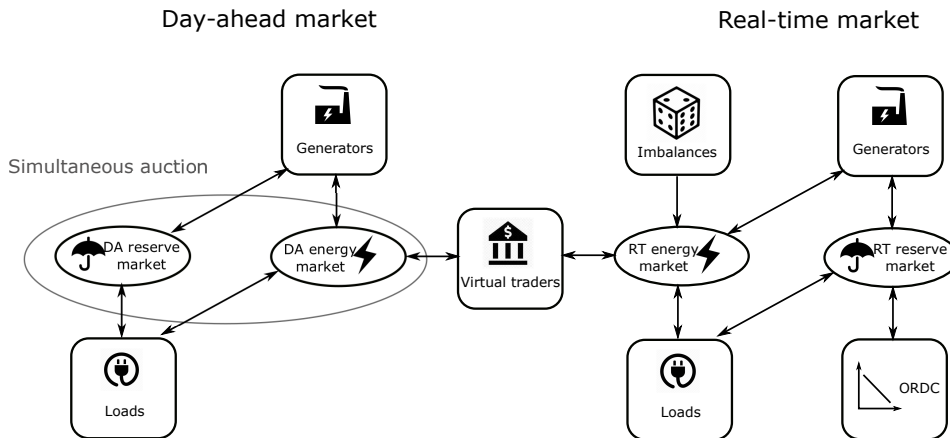
(ii) The least disruptive change in market design that we consider is the introduction of a market for real-time reserve capacity. Concretely, this implies that the real-time price of energy is enhanced with a scarcity adder, and also that a scarcity adder is applied to the amount of reserve capacity that is available *after* imbalances have been cleared within a given imbalance interval. This scarcity adder is computed on the basis of the excess reserve capacity that remains available to the system after activation for balancing, which is a signal of scarcity in the system.

(iii) We then consider the implications of lifting the penalties that agents incur for changing their real-time positions relative to their day-ahead trades.

(iv) The next step that we consider is the introduction of virtual trading in day-ahead energy markets, which would imply a significant departure from the current European market design, since it would signify the role of the day-ahead market as a forward market.

(v) The final step that we consider is the simultaneous trading of energy and reserve capacity in the European day-ahead market. This step presents a number of institutional and technological challenges, that we discuss in further detail after presenting our results.

The main methodological contribution of the present work is the grounding of our analysis on a stochastic equilibrium modeling framework. This stochastic equilibrium framework (Ralph and Smeers, 2015), has been employed extensively in the context of capacity expansion equilibrium (Ehrenmann and Smeers, 2011; de Maere d’Aertrycke et al., 2017; Abada et al., 2017). The framework is adapted to the analysis of day-ahead and real-time markets in the present paper. We demonstrate how this framework is appropriate for explaining the arbitrage mechanisms that induce price formation in day-ahead markets, and how it can be used for quantifying various metrics of economic significance, including clearing prices in real-time and day-ahead markets, as well as the economic surplus of market participants. We illustrate the agility of the modeling framework by using it to represent the wide variety of market designs that are represented in figure 1, and demonstrate its practical usefulness in providing a basis for concrete policy recommendations.

Figure 2: A schematic illustration of the US-style two-settlement system.

The use of stochastic equilibrium as a means of economic modeling is in stark contrast to proposals of stochastic market clearing that have been discussed in the literature, see Bouffard et al. (2005); Pritchard et al. (2010); Zavala et al. (2017); Morales et al. (2014). Whereas stochastic market clearing is a proposal for pricing electricity which presents serious implementation challenges (e.g. related to the definition of scenarios, effects on incentives, computational challenges), stochastic equilibrium (which has an equivalent stochastic programming formulation in a risk-neutral setting) is a powerful tool for policy analysis.

Concretely, the remainder of the paper is developed as follows. In section 3 we present the stochastic equilibrium formulation of a U.S.-style two-settlement system. In section 4 we discuss the formulation of the European market model. In section 5 we present an application of stochastic equilibrium in a case study of the Belgian market and discuss our findings. In section 6 we summarize our concrete recommendations for the introduction of scarcity pricing in the Belgian electricity market, and propose directions of future research.

3. STOCHASTIC EQUILIBRIUM MODELING OF A US-STYLE TWO-SETTLEMENT SYSTEM

The core model of a U.S.-style two-settlement system that we base our analysis on is presented in figure 2. Market agents are represented in squares in this figure, and they correspond to four categories: generators, loads, virtual traders, and the system operator. The products that are traded are indicated in ellipses and they correspond to energy and reserve capacity. Multiple types of reserve products may exist, depending on their corresponding quality (measured in terms of response time). Indeed, we model a secondary and tertiary reserve product in our numerical case study, but collapse the full set of reserve products into a single type of reserve at this stage for the sake of simplifying the exposition. The arrows in the figure indicate whether a given agent is active in a given market.

Note that the sequential nature of two-settlement systems is accounted for by representing the sequence of day-ahead and real-time market clearing through a scenario tree, and establishing a separate market clearing condition for the day-ahead and real-time markets. Concretely, uncertain real-time conditions are indexed by $\omega \in \Omega$, where Ω corresponds to a set of real-time scenarios such as the realization of load and renewable supply forecast errors, as well as contingencies. The link between the two markets is the settlement of forward market positions at real-time prices, and the attitude of market agents who are active in day-ahead markets towards risky real-time outcomes; we capture this effect through coherent risk measures⁶. Note that virtual traders in figure 2 are only active

⁶Coherent risk measures are risk measures which satisfy the properties of convexity (the marginal cost of risk is increasing),

in the energy market, and not the reserve market.

In the present and following section, we will highlight the most significant aspects of the models that we develop. The exact formulation of the models is presented in an electronic supplement⁷.

3.1 Real-Time Market Equilibrium

In real time, generators trade energy and reserve. Assuming price-taking behavior, generators aim to maximize their real-time profits, which are expressed as the difference between their revenues and the fuel cost that they incur in real time. Thus, a generator g in a given real-time scenario $\omega \in \Omega$ maximizes the following:

$$\max_{p_{g,\omega}^{RT} \geq 0, r_{g,\omega}^{RT} \geq 0} \lambda_{\omega}^{RT} \cdot p_{g,\omega}^{RT} + \lambda_{\omega}^{R,RT} \cdot r_{g,\omega}^{RT} - C_g \cdot p_{g,\omega}^{RT} \quad (2)$$

where the variables $p_{g,\omega}^{RT}$ and $r_{g,\omega}^{RT}$ represent the real-time supply of power and reserve respectively, λ_{ω}^{RT} and $\lambda_{\omega}^{R,RT}$ correspond to the real-time price of energy and reserve respectively, and C_g is the marginal cost of the generator.

Note that the model presented above already departs from European market design: the European market does not trade capacity in real time. Consequently, there does not exist a real-time price for reserve in European markets, and the second term in the objective function is entirely absent.

Note, also, the precise meaning that we attribute to the reserve capacity variable $r_{g,\omega}^{RT}$ throughout the paper: $r_{g,\omega}^{RT}$ is the amount of reserve capacity that remains available *after* reserve has been activated in order to clear an imbalance within a given real-time operating interval. This is in contrast to the amount of reserve capacity that is available *before* clearing an imbalance. As we will argue more extensively in section 5, this distinction is crucial in markets where real-time settlement intervals extend over multiple minutes (imbalance intervals are 15 minutes in a large number of European markets, following the E.U. Balancing Guideline⁸, whereas they have been reduced to 5 minutes in a number of U.S. markets).

Generators must respect two major types of operating constraints in real time, that affect the availability of reserve capacity: capacity, and ramp limits⁹. Constraints on generation capacity are expressed as follows:

$$(\alpha_{g,\omega}^{G,RT,+}) : \quad p_{g,\omega}^{RT} + r_{g,\omega}^{RT} \leq P_{g,\omega}^{RT,+} \cdot y_g, \quad (3)$$

where $P_{g,\omega}^{RT,+}$ is the available capacity of a generator in real time, and where y_g is a unit commitment variable which indicates whether a unit has been committed or not in the day-ahead time frame. As such, y_g is *not* decided in real time, but is instead fixed according to the solution of the day-ahead equilibrium.

Ramp constraints are expressed as follows:

$$(\beta_{g,\omega}^{G,RT}) : \quad r_{g,\omega}^{RT} \leq R_g, \quad (4)$$

where R_g is the ramping capability of a generator within the time horizon that is afforded by the

monotonicity (if one lottery dominates another for every outcome, then it is associated with a lower risk), translation invariance (a sure payoff tomorrow is as good as having the payoff today), and positive homogeneity (doubling the cost doubles the risk). The term has been coined by Artzner et al. (1999).

⁷https://perso.uclouvain.be/anthony.papavasiliou/public_html/Supplement.pdf

⁸See, for example, article 53(1) of the Electricity Balancing Guideline (European Commission, 2017).

⁹We follow Hogan (2013) in referring to ramp limits as the constraints that relate to how quickly a resource can be activated in order to counter an imbalance *within* a given imbalance interval. We are not interested in modeling constraints on the schedule set-points *between* imbalance intervals. These inter-period ramp rate constraints would introduce intractable inter-temporal linkages, while detracting from the focus of the analysis on secondary and tertiary reserves, which are activated *within* imbalance intervals.

response time of the reserve product.

The dual multipliers of the corresponding constraints are indicated on the left of the constraints. Minimum production constraints can be expressed similarly, but we do not develop them here in order not to overburden the notation.

We can express the profit-maximizing behavior of loads in real time analogously. We do not develop this in detail here, and refer the reader instead to the online supplement. Instead, we explicitly discuss the role of the system operator in real time, as it relates directly to the scarcity pricing design. In real time, the system operator procures capacity according to an operating reserve demand curve. This demand curve is described by a set RL of price-quantity pairs $(V_l^R, D_l^R), l \in RL$, that describe a downward-sloping demand function¹⁰. Concretely, the system operator procures reserve in real time so as to maximize its real-time profit, as quantified by the real-time ORDC:

$$\max_{d_{l,\omega}^{R,RT} \geq 0} \sum_{l \in RL} V_l^R \cdot d_{l,\omega}^{R,RT} - \lambda_{\omega}^{R,RT} \cdot d_{l,\omega}^{R,RT}. \quad (5)$$

Here, $d_{l,\omega}^{R,RT}$ corresponds to the amount of real-time reserve capacity that is procured by the system operator.

The system operator limits its procurement according to its bid quantity:

$$(\alpha_{l,\omega}^{R,RT}) : \quad d_{l,\omega}^{R,RT} \leq D_l^{R,RT}, l \in RL. \quad (6)$$

The real-time market equilibrium is completed by the market clearing conditions of the energy and reserve markets:

$$\sum_{g \in G} p_{g,\omega}^{RT} = \sum_{l \in L} d_{l,\omega}^{RT} \quad (7)$$

$$\sum_{g \in G \cup L} r_{g,\omega}^{RT} = \sum_{l \in RL} d_{l,\omega}^{R,RT} \quad (8)$$

where G corresponds to the collection of generators, L corresponds to the collection of loads, and $d_{l,\omega}^{RT}$ is the real-time energy demand of loads.

The real-time market equilibrium for a given real-time scenario $\omega \in \Omega$ can thus be expressed as the collection of the profit maximization problems of generators, loads and the system operator (expressed equivalently through their KKT conditions) and the market clearing conditions.

As we discuss in section 2, scarcity pricing relies on two types of arbitrage. The first is the arbitrage between energy and reserve capacity, and it is already captured in the model that we have presented in the present section. The second is the arbitrage between day ahead and real time. We proceed next to discuss how this arbitrage is modeled through coherent risk measures. Coherent risk measures are not the only means of modeling financial arbitrage, but they are appropriate for our needs since they allow for a mathematical formulation of the market equilibrium.

3.2 Coherent Risk Measures

Similarly to the real-time market equilibrium, the day-ahead equilibrium involves the profit maximization conditions of individual agents and the market clearing of the forward day-ahead market, which is a purely financial market. The introduction of a forward market for energy and reserves implies that the description of agents needs to be completed by a definition of their attitude towards the uncertainty of real-time market clearing outcomes. Risk aversion can be modeled in different

¹⁰The procedure for deriving this demand function is explained in a number of references related to the Texas system (ERCOT, 2015), PJM (Hogan and Pope, 2019), as well as the Belgian system (ELIA, 2018; Papavasiliou et al., 2019). The derivation of the demand curve for the numerical results that are presented in this paper is presented in the electronic supplement.

ways. In this work, the choice is made to model the risk attitude of agents by coherent risk measures (Artzner et al., 1999; Shapiro et al., 2009), see also the definition in footnote 6. This machinery allows us to develop a quantitative description for the formation of day-ahead energy and reserve prices through the arbitrage of profit opportunities between day ahead and real time, which is an essential element in the effective functioning of scarcity pricing. Note that our approach also applies to the more general setting of convex risk measures (Follmer and Schied, 2002).

A risk measure \mathcal{R} is a function that maps uncertain outcomes, defined as random functions $Z(\omega)$, to the extended real line $\bar{\mathbb{R}} = \mathbb{R} \cup \{+\infty\} \cup \{-\infty\}$. The profit function of an agent i is taken as argument of the risk measure. This profit function Γ_i depends on decisions made by different agents and hence is specific for each agent. Agent decisions can depend on uncertainty ω . Hence the profit function will also depend on this uncertainty.

One can now define the risk measure \mathcal{R}_i as the minimum of an uncertainty expectation \mathbb{E}_{q_i} on the profit function Γ_i . This gives the risk measure \mathcal{R}_i as:

$$\mathcal{R}_i(\Gamma_i) = \min_{q_i \in \mathcal{M}_i} \mathbb{E}_{q_i}[\Gamma_i] = \min_{q_i \in \mathcal{M}_i} \sum_{\omega \in \Omega} q_{i,\omega} \cdot \Gamma_i, \quad (9)$$

with \mathcal{M}_i the space where the risk-averse probability distribution q_i belongs.

Risk measures which are coherent obey the properties that are summarized in footnote 6. The definition of coherent risk measures is presented in detail in Shapiro et al. (2009). Note from Shapiro et al. (2009) that the supergradient of a coherent risk measure, $\partial \mathcal{R}_i(\Gamma_i)$, can be expressed equivalently as the set of risk-neutral probabilities \tilde{q}_i which minimize the worst-case payoff of an agent over the uncertainty set \mathcal{M}_i . Mathematically, the following holds:

$$\partial \mathcal{R}_i(\Gamma_i) = \arg \min_{q_i \in \mathcal{M}_i} \mathbb{E}_{q_i}[\Gamma_i] \ni \tilde{q}_i. \quad (10)$$

The supergradient of risk measure \mathcal{R}_i with respect to a variable a can be obtained using the chain rule:

$$\frac{\partial \mathcal{R}_i(\Gamma_i)}{\partial a} = \sum_{\omega \in \Omega} \frac{\partial \mathcal{R}_i}{\partial \Gamma_i} \frac{\partial \Gamma_i}{\partial a} = \mathbb{E}_{\tilde{q}_i} \left[\frac{\partial \Gamma_i}{\partial a} \right]. \quad (11)$$

Equations (10) and (11) will be employed in the following sections in order to describe the back-propagation of real-time prices to the day-ahead market. We provide a more detailed discussion about coherent risk measures, the derivations of the present section, and their application in a specific example, in the electronic supplement.

3.3 Day-Ahead Market Equilibrium

In the context of our problem, the random payoff function Γ_i of a generator which is described in section 3.2 is the real-time profit that the generator can earn from supplying power and reserve in real time, net of the settlement of power and reserve that the generator has sold in the day-ahead market. Concretely, a generator $g \in G$ is described by the following profit maximization problem in the day-ahead market:

$$\begin{aligned} \max_{y_g \geq 0, r_g^{DA} \geq 0, p_g^{DA}} & \lambda^{DA} \cdot p_g^{DA} + \lambda_g^{R,DA} \cdot r_g^{DA} - K_g \cdot y_g \\ & + \mathcal{R}_g[\Pi_{g,\omega}^{RT}(y_g) - \lambda_\omega^{RT} \cdot p_g^{DA} - \lambda_\omega^{R,RT} \cdot r_g^{DA}] \end{aligned} \quad (12)$$

Here, p_g^{DA} and r_g^{DA} is the power and reserve sold by generator g in the day-ahead market, respectively. The day-ahead energy and reserve prices are denoted as λ^{DA} and $\lambda^{R,DA}$ respectively. We assume that generators act as price takers in the day-ahead market. The fixed cost of committing a

generator is denoted as K_g .

The real-time payoff of the generator, which is the argument of the risked profit \mathcal{R}_g , consists of the following terms: (i) the profit of the generator in the real-time market, (ii) the payment required by the generator for buying back its day-ahead position in the energy market at the real-time price of energy, and (iii) the payment required by the generator for buying back its day-ahead position at the real-time price of reserve. Note that $p_{g,\omega}^{RT}$ is the final position generator g in real time, not its adjustment in real time with respect to the day-ahead position.

In order to arrive to a model with economic interpretation, we use a convex relaxation of the commitment decision $y_g \in \{0, 1\}$ by replacing it with the requirement that y_g lie in the unit interval:

$$(\delta_g) : \quad y_g \leq 1 \quad (13)$$

Using a well-known result by Balas (1998), it is straightforward to extend this model in order to represent the convex hull of the feasible set of a unit that respects a wide range of operating constraints, including production minimum, ramp rates, and minimum up and down times. This is discussed in detail by Hua and Baldick (2016).

Note that the day-ahead profit maximization problem of a generator allows for *virtual trading* of energy, since the generator is not constrained by its physical generation capacity. In particular, a generator can sell a negative quantity of power, sell more energy than permitted by its technical maximum (or even sell without actually owning a physical unit), or sell power in the day-ahead market without necessarily committing a unit. By contrast, we do not allow for the virtual trading of reserve. Concretely, we impose limits on the day-ahead trading of reserve capacity by introducing the following constraint:

$$(\beta_g^{G,DA}) : \quad r_g^{DA} \leq R_g. \quad (14)$$

The real-time profit $\Pi_{g,\omega}^{RT}(y_g)$ of the generator is the result of the profit maximization problem described in section 3.1:

$$\Pi_{g,\omega}^{RT}(y_g) = \max_{p_{g,\omega}^{RT}, r_{g,\omega}^{RT} \geq 0} \{ \lambda_{\omega}^{RT} \cdot p_{g,\omega}^{RT} + \lambda^{R,RT} \cdot r_{g,\omega}^{RT} - C_g \cdot p_{g,\omega}^{RT}, \text{ s.t. (3), (4)} \} \quad (15)$$

Note that the profit $\Pi_{g,\omega}^{RT}(y_g)$ is a *function* of the day-ahead commitment decision y_g . Using standard convexity arguments, it can be shown that $\Pi_{g,\omega}^{RT}(y_g)$ is a concave function of y_g , and that its supergradient with respect to y_g is $p_{g,\omega}^{RT,+} \cdot \alpha_{g,\omega}^{G,RT,+}$, where $\alpha_{g,\omega}^{G,RT,+}$ is the dual multiplier of constraint (3). This is useful for expressing the day-ahead market equilibrium through the KKT conditions of the generator profit maximization.

The profit maximization of loads can be introduced analogously. Note that we can introduce an ORDC in the day-ahead market clearing by following exactly the same procedure as in section 3.1.

The day-ahead market equilibrium can then be completed by the introduction of day-ahead market clearing conditions for energy and reserve:

$$\sum_{g \in G} p_g^{DA} = \sum_{l \in L} d_l^{DA}, \quad (16)$$

$$\sum_{g \in G \cup L} r_g^{DA} = \sum_{l \in RL} d_l^{R,DA}. \quad (17)$$

The collection of (i) the KKT conditions of the day-ahead profit maximization of generators, loads, and the system operator, and (ii) the market-clearing conditions of the day-ahead energy and reserve market constitute the day-ahead market equilibrium.

3.4 Back-Propagation of Prices

Having introduced a scenario tree for representing the sequence of day-ahead market clearing followed by real-time clearing, and coherent risk measures for quantifying the attitude of agents towards risky real-time profits, we now have in place all the necessary machinery for describing the arbitrage forces that shape day-ahead energy prices. Let us consider the first-order optimality conditions of generators with respect to day-ahead energy production, p^{DA} . Using equation (10) and the chain rule of equation (11), we have:

$$(p_g^{DA}) : \quad \lambda^{DA} = \mathbb{E}_{q_g} [\lambda_\omega^{RT}], \quad (18)$$

where

$$q_g \in \arg \min_{q_g \in \mathcal{M}_g} \mathbb{E}_{q_g} [\Pi_{g,\omega}^{RT} - \lambda_\omega^{RT} \cdot p_g^{DA} - \lambda_\omega^{R,RT} \cdot r_g^{DA}] \quad (19)$$

is the risk-neutral measure of generator g at the market equilibrium.

Equation (18) expresses in a quantitative way the mechanism by which inter-temporal arbitrage serves to shape day-ahead prices as the expectation of real-time prices under the risk-neutral measure of an agent. Likewise, the first-order optimality conditions with respect to day-ahead reserve capacity, r_g^{DA} , yield a back-propagation formula for day-ahead reserve prices.

Different choices of risk sets \mathcal{M}_g represent different attitudes of agents towards risk. For example, the conditional value at risk (Rockafellar and Uryasev, 2002) employed by Ehrenmann and Smeers (2011) can be represented by defining

$$\mathcal{M}_g = \{q_g : q_{g,\omega} \leq \frac{P_\omega}{\alpha_g}, \sum_{\omega \in \Omega} q_{g,\omega} = 1, q_{g,\omega} \geq 0, \omega \in \Omega\}, \quad (20)$$

where P_ω is the probability of outcome ω and α_g is the percentile of the CVaR measure of agent g . As a concrete illustration, the formulation of the full stochastic equilibrium model using a CVaR risk measure is presented in the electronic supplement of the present publication.

As a special case, setting $\alpha_g = 1$ corresponds to a risk-neutral agent. This reduces the set \mathcal{M}_g to a single element $(P_\omega, \omega \in \Omega)$, and establishes the equivalence between stochastic equilibrium and stochastic programming in a risk-neutral setting. We exploit this equivalence in the case study of section 5.

In the following subsections, we digress slightly from the exposition of the U.S. and E.U. models. The brief digression serves two purposes. (i) In section 3.4.1 we demonstrate the agility of our proposed model in reproducing phenomena that are documented in an informal way in economic reports. (ii) Section 3.4.2 is directly relevant to the conclusions of our paper: it illustrates, by example, that a risk-neutral model may fail to account for the beneficial impacts of virtual trading in aligning day-ahead and average real-time prices.

3.4.1 Day-Ahead Risk Premium

In order to illustrate the agility of our modeling setup, we demonstrate its application on the formation of day-ahead risk premia. In its 2017 market monitoring report of the Texas market, Potomac Economics (2018) discusses risk premia in the ERCOT day-ahead market. The report focuses specifically in the origin of high risk premia in the summer months of 2017:

“Risk is lower for loads purchasing in the day-ahead market and higher for generators selling day-ahead. The higher risk for generators is associated with the potential of incurring a forced outage and having to buy back energy at real-time prices. This explains why the highest premiums occurred during the summer months in 2017 with the highest relative demand and highest prices.”

This argument indicates that generators are “less” eager to commit to day-ahead trades than loads, because generators face the risk of not having their capacity available in real time. Generators are therefore more selective in terms of the day-ahead price that they require. This exerts an upward pressure on day-ahead prices.

We attempt to capture this effect using our risk-averse formulation. Consider a system with three technologies: (i) base-load (nuclear) generation, with a capacity of 7100 MW and a marginal cost of 6.5 €/MWh; (ii) Shoulder (coal) generation, with a capacity of 2000 MW and a marginal cost of 25 €/MWh; and (iii) peak (gas) generation, with a capacity of 2200 MW and a marginal cost of 80 €/MWh. Suppose that coal has its full capacity available with a probability of 90%, and experiences an outage with a probability of 10%.

We consider a demand equal to 9000 MW. Thus, under normal operating conditions, coal is the marginal unit. If the coal unit fails, then gas covers the demand, and coal is exposed to a real-time price of 80 €/MWh. The average real-time price is 30.50 €/MWh. The day-ahead price which is obtained by the risk-averse equilibrium model is 30.80 €/MWh. We then lift demand to 10000 MW. Under normal operating conditions, gas is the marginal unit. If the coal unit fails, then we have load shedding, and coal is exposed to a real-time price of 1000 €/MWh. The average real-time price is 172.00 €/MWh. The day-ahead price that is obtained from the risk-averse equilibrium model is 176.80 €/MWh. Indeed, the model produces a risk premium, and this risk premium increases as demand increases. We are thus able to reproduce the effect that is described in the Potomac report.

3.4.2 Interplay between risk aversion and virtual trading

We now extend the previous example in order to gain some insight regarding the interplay between risk aversion and virtual trading. We use the same setup as in the previous paragraph, where the demand is fixed to 10000 MW.

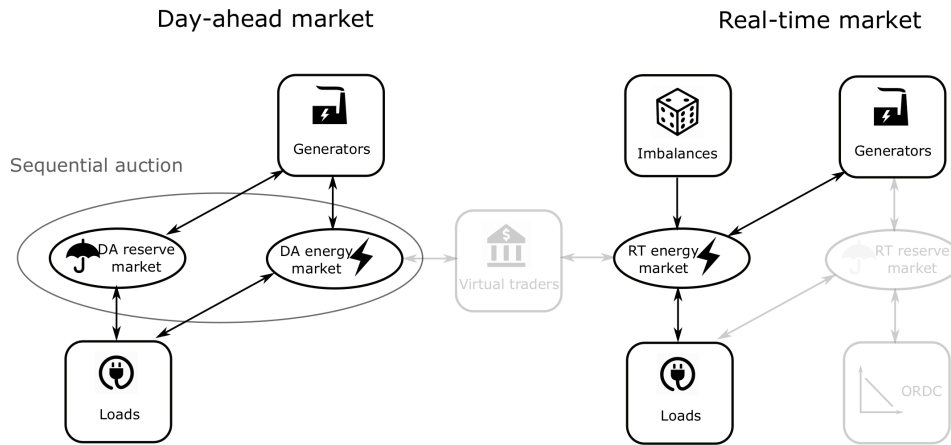
In the case where agents are risk averse (same settings as in the previous paragraph, except for the nuclear generator which is assumed to be risk neutral), the day-ahead energy price without virtual trading becomes 176.80 €/MWh. The day-ahead price that is computed from the risk-averse equilibrium model with virtual trading amounts to 172.00 €/MWh, and is driven by the risk-neutral nuclear generator which is able to engage in virtual trade. In the case where agents are risk-neutral, the day-ahead energy price resulting from the equilibrium model is equal to 172.00 €/MWh, regardless of whether or not we suppress virtual trading.

In the risk-neutral model, the results with and without virtual trading are identical. This may suggest that the role of virtual trading is irrelevant. Note, however, that virtual trading is useful for agents with different attitudes towards risk to trade that risk. In the risk-averse model, the presence of virtual trading makes the difference between price convergence under the statistical probability measure (when virtual trading is allowed) and positive risk premia (when virtual trading is not allowed). Therefore, a risk-neutral model is likely to miss the positive effects of virtual trading in terms of price convergence. We revisit this observation in section 6, when we discuss the conclusions of our case study.

4. STOCHASTIC EQUILIBRIUM MODELING OF THE EUROPEAN DESIGN

The European market design that we are interested in modeling is presented in figure 3. Compared to the U.S. two-settlement system model that is presented in figure 2, one may observe the following principal differences: (i) the European real-time market does not trade reserve capacity; (ii) reserves are cleared separately from energy in the European day-ahead market; and (iii) virtual trading in energy is not permitted in the European day-ahead market. We discuss these differences in turn, and present our approach towards modeling them.

Figure 3: A schematic illustration of the European day-ahead and real-time electricity market.



4.1 Real-Time Market Equilibrium

There exists a predominant view among European system operators that interprets real-time operations as a service, as opposed to a market. This point of view is reflected in the existence of balancing responsible parties (BRPs) which are responsible for balancing their financial (day-ahead or intraday) positions with their physical net injection or off-take. U.S. markets are, instead, based on the notion that individual resources should be adjusted, even if very close to real time, to a level that aids the system, even if this implies significant changes with respect to previously defined market positions.

The original motivation for the existence of BRPs is that, as long as BRPs take all the necessary actions to balance within their own perimeter of resources, there is little residual uncertainty left over to the system operator. This is a problematic point of view, since it ignores transmission constraints and the benefits of sharing cheap resources in real time, nevertheless it dictates European real-time operations for the time being. With the increasing integration of renewable resources, and the consequent need for market players to adjust their positions to the rapidly evolving real-time conditions of the system, it seems valuable to transition to a point of view whereby the day-ahead market is treated as a forward financial market, and real time is the spot market against which forward positions are settled financially.

There are two concrete side-effects of treating real-time operations as a set of balancing services, as opposed to the actual spot market for electricity. (i) There exists little if any consistency between the products traded in day ahead and in real time. Concretely, most European markets do not trade reserve capacity in real time. Instead, agents are expected to be able to deliver, in real time, the quantity of reserves that they have committed to in the day ahead. After this reserve capacity is delivered in real time, additional reserve capacity that is available to the system is not rewarded. (ii) Agents are expected to honor their spot market (day-ahead) positions, therefore any deviations from the spot market results are ‘discouraged’, and system operators expect that balancing responsible parties maintain power balance within their perimeter in real time.


The first effect is captured in our model by removing real-time reserve capacity payments from the objective function of market participants. The second effect is captured by penalizing deviations of agents from their day-ahead trading positions.

Concretely, we introduce slack variables that measure the upward ($s_{g,\omega}^{RT,+}$) and downward ($s_{g,\omega}^{RT,-}$) deviation of real-time power production of generators from their day-ahead positions:

$$p_{g,\omega}^{RT} - p_g^{DA} - s_{g,\omega}^{RT,+} \leq 0, \quad (21)$$

$$p_g^{DA} - p_{g,\omega}^{RT} - s_{g,\omega}^{RT,-} \leq 0. \quad (22)$$

Figure 4: The three-stage decision tree which is used for modeling the sequential clearing of reserve followed by energy in the day-ahead market.



We then penalize these slack variables by administrative penalties for upward (ϵ_g^+) and downward (ϵ_g^-) deviations. The profit maximization objective of generators in real time then reads as follows:

$$\max_{p_{g,\omega}^{RT} \geq 0, s_{g,\omega}^{RT,+} \geq 0, s_{g,\omega}^{RT,-} \geq 0} \lambda_{\omega}^{RT} \cdot p_{g,\omega}^{RT} - C_g \cdot p_{g,\omega}^{RT} - \epsilon_g^+ \cdot s_{g,\omega}^{RT,+} - \epsilon_g^- \cdot s_{g,\omega}^{RT,-}. \quad (23)$$

Note that the objective function no longer includes revenues from selling reserve capacity in real time, since there exists no real-time reserve market. The optimal objective function value of this profit maximization is denoted as $\Pi_{g,\omega}(y_g, p_g^{DA})$, and it is a *function* of the day-ahead unit commitment y_g and the day-ahead sales of energy p_g^{DA} .

The real-time profit maximization of loads needs to be adapted accordingly. Moreover, the system operator is absent in the real-time market, and there exist no market-clearing conditions for reserve capacity in real time. The real-time equilibrium is then written out as the collection of KKT conditions of the generator and load profit maximization problems, as well as the market clearing condition for energy. The precise collection of conditions is presented in the electronic supplement.

4.2 Day-Ahead Energy Exchange

The separation of energy and reserve markets has had a difficult history in U.S. market design (Oren, 2001). The difficulties that were present in the early California electricity market design were due to the migration of bidders to reserve markets of lower quality, due to the fact that the proposed market clearing rules were remunerating lower quality reserve products at a higher price. Eventually, the sequential clearing of reserves was abandoned in California, and the paradigm today in U.S. markets is to clear energy, reserve and transmission capacity simultaneously in multi-product day-ahead and real-time auctions.

By contrast, reserves are auctioned separately from energy in European day-ahead markets. For example, in Belgium reserves are auctioned before energy in the day-ahead market, while in Italy and Spain they are auctioned after the day-ahead energy market (Dominguez et al., 2019). Since capacity can be allocated to either energy or reserves, but not both, the challenge for market participants is to precisely anticipate the value of their capacity in one of these markets, in order to bid accurately in the other market. We propose a modeling framework for capturing this challenge, and its implication on the back-propagation of energy and reserve prices. This modeling setup is especially relevant for Belgium, where the system operator is considering a reorganization of reserve procurement so as to transition to dynamic sizing (De-Vos et al., 2019), in order to prevent the oversizing that takes place when one prepares for the worst-case scenario over a long stretch of time.

In order to model the increased uncertainty of sequential clearing, we modify the day-ahead market clearing model of section 3.3. Inspired by Dominguez et al. (2019), we specifically assume a three-stage scenario tree, as depicted in figure 4. The first stage of the tree is when agents decide for reserve capacity. Then the ‘type of day’ is revealed, which is the assessment of the TSO for what quantity of operating reserve will be required for the following day. Following the clearing of the reserve auction, the day-ahead energy market is cleared.

In order to fully describe the scenario tree, we further need to specify what is observed

within each node, and the transition probabilities: (i) The uncertainty which is observed in stage 2 is the state of the world which will influence the real-time imbalances of stage 3, and which from the point of view of the agents implies an observable price for day-ahead energy, λ_{ω}^{DA} . (ii) What is observed in the third stage is the realization of real-time renewable supply and real-time demand, which from the point of view of the agents implies a real-time price for energy and reserve. (iii) The first stage reveals to the agents the price of day-ahead reserve capacity, as a result of their competition. The set of second and third stage outcomes is denoted as Ω_2 and Ω_3 respectively.

The resulting nested equilibrium follows the approach of Philpott et al. (2013). At stage 2 and second-stage outcome $\omega \in \Omega_2$, a generator decides on its commitment $y_{g,\omega}$ and the amount of power that it wishes to sell in the day-ahead market. Since virtual trading in energy is not permitted, the generator must obey the following constraint:

$$(\alpha_{g,\omega}^{G,DA,+}) : \quad p_{g,\omega}^{DA} + r_g^{DA} \leq P_g^{DA,+} \cdot y_{g,\omega} \quad (24)$$

This constraint implies that a generator must start its unit up in the day-ahead energy market if it is to be cleared for selling energy in the day-ahead market, and may only sell up to the capacity that has not already been allocated for providing reserve.

When in node $\omega \in \Omega_2$, a generator discounts its risky real-time payoff according to a coherent risk measure $\mathcal{R}_{g,\omega}$. Note that the risk measure depends on the risk attitude of the agent g , but also on the outcome $\omega \in \Omega_2$. The dependence on $\omega \in \Omega_2$ is due to the fact that the conditional distribution of third-stage outcomes in figure 4 depends on $\omega \in \Omega_2$.

The generator profit maximization objective in the second stage then reads as follows:

$$\max_{y_{g,\omega} \geq 0, p_{g,\omega}^{DA} \geq 0} \lambda_{\omega}^{DA} \cdot p_{g,\omega}^{DA} - K_g \cdot y_{g,\omega} + \mathcal{R}_{g,\omega}(\Pi_{g,\omega}^{RT}(y_{g,\omega}, p_{g,\omega}^{DA}) - \lambda_{\omega'}^{RT} \cdot p_g^{DA}). \quad (25)$$

Note that the day-ahead reserve capacity, r_g^{DA} , is a *parameter* for this profit maximization. Using standard convex optimization arguments, we can observe that the optimal objective function of this profit maximization, which is denoted as $\Pi_{g,\omega}^{DA}(r_g^{DA})$, is a concave function of r_g^{DA} . This function is used in the day-ahead reserve equilibrium problem, which is presented in section 4.3.

The profit maximization problem of loads is developed analogously. The day-ahead energy exchange equilibrium is then written out as the collection of KKT conditions of the generator profit maximization problem, the KKT conditions of the load profit maximization problem, and the market clearing condition of the day-ahead energy market. The precise collection of conditions is developed in the electronic supplement.

4.3 Day-Ahead Reserve Auction

In the root node of the scenario tree of figure 4, a generator faces the risky returns of stage 2, $\Pi_{g,\omega}^{DA}(r_g^{DA})$, which are evaluated against a coherent risk measure \mathcal{R}_g . The objective of the generator in the day-ahead reserve auction can be written out as follows:

$$\max_{r_g^{DA} \geq 0} \lambda^{R,DA} \cdot r_g^{DA} + \mathcal{R}_g(\Pi_{g,\omega}^{DA}(r_g^{DA})). \quad (26)$$

The profit maximization of the generator is a convex optimization problem, which we can express equivalently using its KKT conditions. A similar process can be applied for loads. We can further introduce a day-ahead ORDC by including a transmission system operator as a buyer of reserve capacity. The full specification of the reserve equilibrium is completed by the market clearing condition on reserve capacity. The full collection of equilibrium conditions is provided in the electronic supplement.

5. RESULTS

In this section we compare the performance of the various designs that are described in figure 1. These designs are modeled on the basis of the US and EU models that we presented in sections 3, 4. We conduct our analysis under the assumption that all agents are risk-neutral, which allows us to cast our equilibrium problem as a stochastic program. This is required in order to be able to cope with the large scale of the realistic numerical simulation of the Belgian market. We revisit the implications of our assumption in the discussion in section 6.

Our analysis is based on the period of September 2015 - March 2016, for which we have access to detailed market data. This time interval was also the basis for the analysis presented by Papavasiliou et al. (2018), where the data is described in detail. The precise assumptions of the case study are presented in the electronic supplement.

5.1 Tested Designs

We will present our analysis as an evolution from the existing Belgian market design to an ideal US two-settlement system. We will report day-ahead prices for energy and reserves, and the profits of CCGT units as our key performance indicators. The profits of loads and additional key performance indicators are included and discussed in the electronic supplement.

Before presenting the results, we proceed to discuss the various market designs that we consider. Each model represents an evolution with respect to the existing market design. This evolution is presented in figure 1. We introduce one incremental change to the Belgian market design at a time, until we arrive to a US two-settlement system. We then analyze the impact of each change on market prices and generator profits. The path that we consider is not the only one from the existing Belgian design to a US two-settlement system, but is rather based on our estimate of increasingly disruptive changes. Thus, the early intermediate models are the easiest to implement, since they require the fewest changes to the existing Belgian design.

E.U. The E.U. model represents our first proxy of current Belgian market operations. The clearing of reserve precedes the clearing of energy in the day-ahead market. No virtual trading is allowed. There is no market for real-time reserve capacity. Deviations are penalized at 10% of the marginal cost of generators. The day-ahead demand for reserve capacity is set at the Belgian reserve requirements: 140 MW for secondary reserve (aFRR), and 350 MW for tertiary reserve¹¹ (mFRR). At this stage, we can also insert an energy adder, although this is not expected to have any effect, neither on system dispatch, nor on market clearing prices¹².

E.U.-Inelastic. The E.U.-Inelastic model represents our second proxy of the Belgian market. In this model, the following constraint is introduced in the real-time market:

$$s(\gamma_{g,\omega'}^{G,RT}) : \quad r_g^{DA} - r_{g,\omega'}^{RT} \leq 0, \omega' \in \Omega_3 \quad (27)$$

¹¹Note that this is the requirement for tertiary reserve from production (R3Prod). In particular, (i) we are ignoring the tertiary reserve requirements covered by demand (the co-called ICH product). This has no effect on the model, since the ICH requirements can easily be covered by scheduled demand, and demand is not eligible for covering R3Prod requirements. Also, (ii) we are ignoring R3 dynamic profile. The definition of R3 dynamic profile is idiosyncratic (https://www.elia.be/~media/files/Elia/Products-and-services/ProductSheets/S-Ondersteuning-net/S8_The-tertiary-reserve-Dynamic-Profile.pdf) and deviates from the simple definition, used in this paper, of a resource that is capable to respond within 15 minutes).

¹²Effectively, generators can adjust their bids in order to internalize the adder (for example, an adder of 1 €/MWh which generators can anticipate (even in expectation) will imply that generators will simply adjust their energy bid by 1 €/MWh down, and therefore the same exact outcome will prevail in terms of both dispatch as well as market clearing price as if the adder did not exist in the first place).

This constraint requires that the real-time reserve capacity be at least at the level of the day-ahead reserve capacity. Note that $r_{g,\omega}^{RT}$ corresponds to the excess headroom that is left over in a unit *after* an imbalance has been cleared within a given imbalance interval. Whether or not this constraint should be imposed depends on whether or not the system operator demands that units make their reserve capacity available, even *after* they have been activated for clearing an imbalance. The argument for removing this constraint is that a generator that has supported the system within an imbalance interval should not be held accountable for reserve capacity shortfall at the *end* of the imbalance interval. The argument for keeping the constraint is that the end of one imbalance interval signifies the beginning of a new imbalance interval, and therefore the system should be prepared, anew, to balance real-time uncertainty. The E.U. model corresponds to the first point of view (not including constraint (27)). The E.U.-Inelastic model corresponds to the second point of view (including constraint (27)). Together, these models envelope our proxy of current Belgian market operations.

RTReserve. This model emerges from the introduction of a real-time operating reserve capacity product in the Belgian electricity market. An identical demand function for operating reserve is introduced in the day-ahead market. The reader is referred to the electronic supplement of this article for a detailed description of how the ORDC is calibrated against historical data of the Belgian system.

RTReserve2. This model emerges from the removal of administrative penalties for real-time deviations from day-ahead positions. Concretely, we represent this effect by inserting constraints (21) and (22) to the RTReserve model. In terms of practical operations, the implementation of this measure could be facilitated by a transition to central dispatch in real time, whereby the system operator issues instructions to individual generators and posts real-time prices that are consistent with these instructions. This should be contrasted to the current paradigm, whereby balancing responsible parties cannot support the system needs (e.g. by increasing their net injection when the system is short) unless they speculate on the imbalance price value. This implies a risk for BRPs which would be alleviated in central dispatch, since in central dispatch the set point is determined by the system operator, and is consistent with the needs of the system (as expressed in the real-time price) as well as the profit maximizing behavior of individual resources.

V.T. This model is the evolution of RTReserve2 whereby we lift the physical constraints on the trading of reserve in the day-ahead market. Concretely, in constraint (3) we remove the production variable $p_{g,\omega}^{RT}$ from the left-hand side. Thus, we impose the requirement that any reserve that is contracted in the day-ahead reserve auction must be backed up by the commitment of generators in the day-ahead time frame. This corresponds to the day-ahead nomination of generators which is applied in the Belgian market, according to which reserve commitments must be backed up by physical capacity in the day-ahead time frame. On the other hand, energy is traded freely (i.e. without the backing of physical assets) in the day-ahead time frame, which corresponds to a departure from the current practice of the Belgian market, at least in principle.

U.S. This model is the evolution of V.T., whereby the energy and reserve markets are cleared simultaneously in the day ahead.

5.2 Prices

In what follows, we present average results of our models and historically observed prices. The model results are the average values over the full scenario tree that is described in figure 4, whereas the historical realizations can be thought of as sample realizations over this tree.

Table 1: Day-ahead energy price (€/MWh) for the models defined in section 5.1. Average real-time prices are identical to day-ahead prices.

Month	US	VT	RTReserve	RTReserve2	EU	EU-Inel.	Hist. DA	Hist. RT
1	41.00	41.00	41.00	41.12	35.61	56.39	52.50	39.51
2	31.17	31.17	31.17	31.31	30.68	32.63	55.41	61.04
3	46.88	46.88	46.88	47.05	30.60	66.23	43.12	36.57
4	37.44	37.38	37.36	37.49	28.77	51.5	35.94	33.31
5	41.25	41.25	41.25	41.32	27.17	63.36	32.61	29.48
6	21.74	21.71	21.73	21.87	19.61	26.42	25.39	21.80
7	21.14	21.12	21.13	21.23	20.74	21.44	27.13	25.11
Average	34.37	34.36	34.36	34.48	27.60	45.42	38.87	35.26

We present the energy prices for the different designs in table 1. The average real-time prices are identical to the average day-ahead prices (note, however, that the averages are different between the different market models), as indicated in the caption of the table. Note that the historical average day-ahead price is greater than the average historical real-time price, indicating a risk premium associated with day-ahead trading¹³. The E.U. and E.U.-Inelastic models, which are the closest proxies to the current Belgian design, are enveloping the historically observed day-ahead and real-time energy prices. Recall that the two models differ in terms of whether or not reserve capacity is required to be available after the activation of reserve within an imbalance interval. It is evident that this requirement may have a very significant impact on prices.

As we will explain later, the E.U.-Inelastic design results in an over-valuation of reserve capacity. On the other extreme, the E.U. model completely removes the ORDC from the real-time market. This results in an under-valuation of reserve. The advantage of an ORDC which is designed on the basis of loss of load probability is that the valuation of reserve self-adjusts to reasonable levels, where we will justify ‘reasonable’ later in the profit analysis of section 5.3. The loss of load probability in itself can be estimated using historical system data, and permits the appropriate calibration of the ORDC to a level which is consistent with system operator reliability standards.

The price reduction observed in the energy price of the E.U. model can be prevented by the introduction of a real-time market for reserve capacity. This restores the connection between real-time energy prices and the real-time value of reserve capacity. This explains the increase in energy prices which is observed in the RTReserve and RTReserve2 models. The transition to virtual trading has a minor effect on prices for the risk-neutral case. The transition to simultaneous also has a minor effect on prices.

The energy prices are largely linked to reserve prices, due to no-arbitrage conditions. Therefore, in order to understand the energy prices, we focus on understanding reserve prices, which we analyze below.

The price for fast reserve is shown in table 2 and for slow reserve in table 3. We present reserve prices for all models, as well as historical prices. The data source for the historical reserve price data is ELIA.

There are occasional differences between day-ahead and average real-time prices, but they are minor. For the E.U. and E.U.-Inelastic models, this price effectively reflects the marginal cost of activating resources in the day-ahead time frame, and is not necessarily reflective of the real-time value of reserve in keeping loss of load probability in check. For the RTReserve, RTReserve2, VT and US models, the price of reserve is driven by the demand side, and relates to the value of reserve in keeping a low loss of load probability.

¹³Note from the last two columns of the table that the historically observed differences between monthly averages of day-ahead energy prices and monthly averages of real-time energy prices in Belgium is relatively high compared to other markets (e.g. Texas), see also the discussion in section 2.1 of Papavasiliou et al. (2019).

Table 2: Day-ahead price of fast reserve (€/MWh) for the models defined in section 5.1.

Month	US	VT	RTReserve	RTReserve2	EU	EU-Inel.	Hist.
1	17.60	17.60	17.60	17.62	1.37	34.00	10.90
2	9.01	9.04	9.03	9.08	1.25	11.08	8.67
3	22.73	22.73	22.73	22.83	1.25	42.07	11.79
4	21.16	21.34	21.31	21.35	2.90	34.78	10.57
5	24.50	24.52	24.51	24.49	1.07	46.30	9.25
6	8.77	8.83	8.87	8.81	1.11	13.49	7.69
7	6.39	6.42	6.40	6.36	0.97	6.55	8.28
Average	15.74	15.78	15.78	15.79	1.42	26.90	9.59
Average (RT)	15.56	15.69	15.65	15.15	N/A	N/A	N/A

Table 3: Day-ahead price of slow reserve (€/MWh) for the models defined in section 5.1.

Month	US	VT	RTReserve	RTReserve2	EU	EU-Inel.	Hist.
1	10.82	10.83	10.86	11.00	1.37	34.00	4.66
2	6.32	6.34	6.34	6.43	1.25	11.08	4.66
3	13.73	13.75	13.75	14.01	1.25	42.07	4.66
4	14.34	14.21	14.44	14.52	2.90	34.78	4.66
5	16.77	16.76	16.79	16.83	1.07	46.30	4.71
6	6.90	6.87	6.91	6.93	1.11	13.49	6.08
7	5.02	5.00	5.02	5.00	0.97	6.55	7.46
Average	10.56	10.54	10.59	10.67	1.42	26.90	5.27
Average (RT)	10.70	10.54	10.52	10.17	N/A	N/A	N/A

5.3 Profits

In table 4 we present the profit results for the 8 CCGT units that were active in the market during the test period. Note that three other CCGT units were available on strategic reserve, and were therefore not actively producing power in the energy market. It is worth comparing these results to the capital investment cost of a typical CCGT unit. We assume an overnight cost of 676 \$/kW (EIA 2012 estimate), an exchange rate of 0.88 €/\$, annual discounting at a rate of return r , and T years of investment lifetime. We consider a range of r from 8 to 12%, and an investment lifetime of 25 to 30 years. Normalizing by the generator capacity (as opposed to the total yearly production), this gives CCGT investment costs ranging from 6.03 €/MWh to 8.66 €/MWh.

We have indicated the entries of table 4 according to how they compare to the running investment cost of a typical CCGT unit. Generators with a profit below 6.03 €/MWh are indicated in bold font, and correspond to units that are unable to recover fixed costs, even under optimistic assumptions about fixed costs. Generators that are in the range of 6.03 - 8.66 €/MWh are indicated in italic font, and correspond to units that earn a profit within the range of investment costs. These units are breaking even. Units indicated in normal font are earning a profit above, 8.66 €/MWh, and are therefore covering investment costs even under demanding investment requirements.

We observe that the two envelope models of the European design, EU and EU-Inelastic, cover a wide range of generator profits. Recall that EU-Inelastic corresponds to the EU design in which generators are forced to carry their day-ahead reserve capacity in real time, following activation, as indicated in constraint (27). Therefore, the extent to which generator capacity *after* activation must correspond to reserve capacity committed in the day-ahead market can shift a unit from making losses to earning excessive profits. Removing this requirement altogether, which is the case in the EU model, places 4 out of 8 units in a non-viable financial position. The introduction of a real-time market for reserve capacity (RTReserve and RTReserve2) restores 3 of these units to breaking even, and 1 of them to covering its investment costs comfortably. In a risk-neutral environment, this is the

Table 4: Generator profits (in €/MWh) for the models defined in section 5.1.

	US	VT	RTReserve	RTReserve2	EU	EU-Inel.
CCGT1	7.73	7.37	7.37	7.40	2.59	16.15
CCGT2	20.68	20.66	20.68	20.79	15.07	31.80
CCGT3	8.06	8.06	8.06	8.09	2.64	19.03
CCGT4	12.04	12.04	12.04	12.08	3.84	28.62
CCGT5	21.07	21.05	21.07	21.18	15.45	32.26
CCGT6	8.30	8.29	8.30	8.32	2.66	19.42
CCGT7	21.45	21.43	21.45	21.56	15.82	32.57
CCGT8	20.58	20.56	20.58	20.69	14.93	31.67

key market design change.

Interpreting the results. A major difficulty with the absence of a real-time market for reserve capacity is that it becomes difficult to value reserve capacity precisely. Using the no-arbitrage conditions of the stochastic equilibrium model, the back-propagation of the day-ahead price of reserve capacity when we are not forced to carry any reserve capacity after activation (E.U. model) can be expressed as follows:

$$\lambda^{R,DA} = \beta_g^{G,DA} + \mathbb{E}[\alpha_{g,\omega}^{G,DA}],$$

where $\beta_g^{G,DA}$ corresponds to flexibility scarcity in the day-ahead market (constraint (14)) and $\alpha_g^{G,DA}$ corresponds to capacity scarcity in the day-ahead market (constraint (24)). This signal is too *weak* to signal scarcity in the system.

On the other hand, when we are forced to carry the full amount of reserve after activation (EU-Inelastic model), the scarcity signal is too *strong*:

$$\lambda^{R,DA} = \beta_g^{G,DA} + \mathbb{E}[\alpha_{g,\omega}^{G,DA}] + \mathbb{E}[\gamma_{g,\omega'}^{G,RT}],$$

where $\gamma_g^{G,RT}$ corresponds to the requirement of carrying the day-ahead reserve after activation (constraint (27)).

The real-time ORDC automates this calculation in a self-correcting fashion, and arbitrage propagates this price to the day-ahead market, thereby signaling investment in reserve capacity in case of tight system conditions:

$$\lambda^{R,DA} = \beta_g^{G,DA} + \mathbb{E}[\alpha_{g,\omega}^{G,DA}] + \mathbb{E}[\lambda_{\omega'}^{R,RT}].$$

5.4 Additional Results.

Loads face an increase in energy prices due to scarcity prices, however they can offset this effect by offering reserve services to the system. In section 4 of the electronic supplement we quantify the level of engagement that allows loads to offset the increase in energy prices due to scarcity pricing with increased income from offering reserve services.

Section 5 of the electronic supplement provides a breakdown of the average daily welfare between market stakeholder profits, and reports the daily variance of these profits. In addition, we analyze the amount of fast and slow reserve capacity that is committed before and after activation for clearing imbalances, and we discuss the dependence between real-time energy prices and stress in the system, as measured by the amount of leftover reserve capacity *after* activation. We comment on these results in the electronic supplement.

6. CONCLUSIONS AND PERSPECTIVES

Reserve prices drive the profitability of flexible resources. Therefore, getting real-time reserve prices right is a fundamental aspect of sound market design in an environment of large-scale renewable energy integration, where flexible resources are needed for supporting system security. Flexible resources include both fast-moving thermal units, as well as demand-side resources. In the future, such flexibility could be sourced from distributed resources in medium and low-voltage grids, as well as resources from other sectors. A framework for the integration of distributed flexibility in wholesale electricity markets, and its interplay with operating reserve demand curves and scarcity pricing, is described in the work of Caramanis et al. (2016).

Our analysis underscores the importance of establishing a well-functioning real-time market as a necessary condition for rewarding reserve services adequately in a regime of power system operations that requires significant levels of flexibility. The most important measure in this direction is to introduce a market for real-time reserve capacity. Concretely, we propose the introduction of a scarcity adder for reserve capacity which is payable to standby real-time reserve capacity, and which also uplifts the Belgian imbalance price. The calculation of this adder requires the so-called Available Reserve Capacity (ARC), which is measured in real time by the system operator, and which is used for computing the loss of load probability that is required for computing the scarcity price. The Belgian system operator is already using ARC for computing scarcity prices ELIA (2018). These adders are posted on-line in the day following operations by the Belgian system operator, since October 2019.

The implementation of a real-time market for reserve capacity in Belgium, as described in the previous paragraph, is the least ground-breaking measure along the chain of evolutions that are indicated in figure 1. Nevertheless, as indicated in the analysis of section 5, it delivers the greatest benefits in a market with risk-neutral agents.

Coincidentally, our analysis indicates that, in a risk neutral setting, the more disruptive measures of introducing virtual trading and co-optimizing the trading of reserve and energy in the day-ahead time frame also have a lesser influence on prices and the profitability of flexible resources, as indicated in tables 1 and 4. A transition to explicit virtual trading would require a radical overturn of the European view of real-time markets as a service, which remains endemic in European policy debates, and is therefore likely to be a time-consuming effort. A future transition to co-optimization of energy and reserves is also likely to be challenging, as it raises computational challenges related to the uniform pricing approach that is adopted in European day-ahead markets based on paradoxically rejected bids, and requires the engagement of various stakeholders, including transmission system operators and market operators. Our conclusion, therefore, is that the most realistic and effective first step for the implementation of scarcity pricing in Belgium is the introduction of a real-time market for reserve capacity.

We note that our observations about the importance of virtual trading and simultaneous clearing of energy and reserve are only valid in a risk-neutral setting. As we have indicated in the example of section 3.4, the impact of virtual trading is likely to have a material impact on price convergence in the absence of risk-neutral agents¹⁴. In conducting our large-scale numerical simulation of the Belgian market, we are inevitably limited to a risk-neutral model in order to be able to cast the problem equivalently as a stochastic program. In future work, we are interested in extending this large-scale numerical simulation beyond the risk-neutral setting. We are currently exploring decomposition methods for achieving this goal. This extension presents an exciting prospect for future research.

¹⁴Note that, in this example, we have shown that virtual trading may alter the market equilibrium, even in a system where agents share homogeneous beliefs about uncertainty. The generalization of our work to a setting where agents hold heterogeneous beliefs about uncertainty is left for future research.

ACKNOWLEDGEMENTS

The author would like to thank William Hogan and Alain Marien for valuable discussions. This research has been supported financially by the Belgian Regulatory Commission for Electricity and Gas, the Francqui Foundation Research Professorship 2018-2021, and the EPOC 2030-2050 project (Belgian energy transition funds).

REFERENCES

- Abada, I., de Maere d'Aertrycke, G., Smeers, Y., 2017. On the multiplicity of solutions in generation capacity investment models with incomplete markets: a risk-averse stochastic equilibrium approach. *Mathematical Programming B* 165 (1), 5–69.
- Aravena, I., Papavasiliou, A., 2017. Renewable energy integration in zonal markets. *IEEE Transactions on Power Systems* 32 (2), 1334–1349.
- Artzner, P., Delbaen, F., Eber, J.-M., Heath, D., 1999. Coherent measures of risk. *Mathematical finance* 9 (3), 203–228.
- Balas, E., 1998. Disjunctive programming: Properties of the convex hull of feasible points. *Discrete Applied Mathematics* 89 (1-3), 3–44.
- Bouffard, F., Galiana, F. D., Conejo, A. J., November 2005. Market-clearing with stochastic security. *IEEE Transactions on Power Systems* 20 (4), 1827–1835.
- Caramanis, M. C., Ntakou, E., Hogan, W., Chakraborty, A., Schoene, J., April 2016. Co-optimization of power and reserves in dynamic T&D power markets with nondispatchable renewable generation and distributed energy resources. *Proceedings of the IEEE* 104 (4), 807–836.
- de Maere d'Aertrycke, G., Ehrenmann, A., Smeers, Y., 2017. Investment with incomplete markets for risk: The need for long-term contracts. *Energy Policy* 105, 571–583.
- De-Vos, K., Stevens, N., Devolder, O., Papavasiliou, A., Hebb, B., Matthys-Donnadieu, J., 2019. Dynamic dimensioning approach for operating reserves: Proof of concept in Belgium. *Energy Policy* 124, 272–285.
- Dominguez, R., Oggioni, G., Smeers, Y., 2019. Reserve procurement and flexibility services in power systems with high renewable capacity: Effects of integration on different market designs. *Electrical Power and Energy Systems* 113, 1014 – 1034.
- Ehrenmann, A., Smeers, Y., 2011. *Stochastic Equilibrium Models for Generation Capacity Expansion*. Vol. 163 of *Stochastic Optimization Methods in Finance and Energy*, International Series in Operations Research and Management Science, Part 2. Springer, pp. 273–310.
- ELIA, 2018. Study report on scarcity pricing in the context of the 2018 discretionary incentives.
- ELIA, 2020. Preliminary report on ELIAs findings regarding the design of a scarcity pricing mechanism for implementation in Belgium. Tech. rep., Belgian transmission system operator.
- ERCOT, 2015. ERCOT market training: Purpose of ORDC, methodology for implementing ORDC, settlement impacts for ORDC.
URL http://www.ercot.com/content/wcm/training_courses/107/ordc_workshop.pdf
- ERCOT, 2019. ERCOT Market Education, Basic Training Program, Module 6: Real-Time Operations.
- European Commission, 2017. Commission regulation (EU) 2017/2195 of 23 november 2017 establishing a guideline on electricity balancing. Tech. rep.
URL <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32017R2195&from=EN>
- European Union, 2019. Regulation (EU) 2019/943 of 5 june 2019 of the European parliament and council on the internal market for electricity. Tech. rep.
URL <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32017R2195&from=EN>
- Follmer, H., Schied, A., 2002. Convex measures of risk and trading constraints. *Finance and stochastics* 6 (4), 429–447.
- Hogan, W., September 2005. On an ‘energy only’ electricity market design for resource adequacy. Tech. rep., Center for Business and Government, JFK School of Government, Harvard University.
- Hogan, W., 2013. Electricity scarcity pricing through operating reserves. *Economics of Energy and Environmental Policy* 2 (2), 65–86.
- Hogan, W., 2016. Virtual bidding and electricity market design. *The Electricity Journal* 29 (5), 33–47.
- Hogan, W. W., Pope, S. L., 2019. PJM reserve markets: Operating reserve demand curve enhancements. Tech. rep., Harvard University.
URL https://sites.hks.harvard.edu/fs/whogan/Hogan_Pope_PJM_Report_032119.pdf

- Hua, B., Baldick, R., 2016. A convex primal formulation for convex hull pricing. *IEEE Transactions on Power Systems*.
- Morales, J. M., Zugno, M., Pineda, S., Pinson, P., 2014. Electricity market clearing with improved scheduling of stochastic production. *European Journal of Operational Research* 235 (3), 765–774.
- Oren, S. S., 2001. Design of ancillary services markets. In: 34th Hawaii International Conference on System Sciences.
- Papavasiliou, A., Smeers, Y., 2017. Remuneration of flexibility using operating reserve demand curves: A case study of Belgium. *The Energy Journal*, 105–135.
- Papavasiliou, A., Smeers, Y., Bertrand, G., 2018. An extended analysis on the remuneration of capacity under scarcity conditions. *Economics of Energy and Environmental Policy* 7 (2).
- Papavasiliou, A., Smeers, Y., de Maere d’Aertrycke, G., 2019. Study on the general design of a mechanism for the remuneration of reserves in scarcity situations. Tech. rep., UCLouvain.
URL <https://www.creg.be/sites/default/files/assets/Publications/Notes/Z1986Annex.pdf>
- Parsons, J., Colbert, C., Larrieu, J., Martin, T., Mastrangelo, E., February 2015. Financial arbitrage and efficient dispatch in wholesale electricity markets. Tech. rep., MIT Center for Energy and Environmental Policy Research.
- Philpott, A., Ferris, M., Wets, R., 2013. Equilibrium, uncertainty, and risk in hydrothermal electricity systems. *Mathematical Programming*.
- Potomac Economics, 2018. 2017 state of the market report for the ERCOT electricity markets.
- Pritchard, G., Zakeri, G., Philpott, A., 2010. A single-settlement, energy-only electric power market for unpredictable and intermittent participants. *Operations Research* 58 (4), 1210–1219.
- Ralph, D., Smeers, Y., 2015. Risk trading and endogenous probabilities in investment equilibria. *SIAM Journal on Optimization* 25 (4), 2589–2611.
- Rockafellar, R. T., Uryasev, S., 2002. Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance* 26, 1443–1471.
- Shapiro, A., Dentcheva, D., Ruszczyński, A., 2009. *Lectures on stochastic programming: modeling and theory*. Society for Industrial and Applied Mathematics.
- Zavala, V. M., Kim, K., Anitescu, M., Birge, J., 2017. A stochastic electricity market clearing formulation with consistent pricing properties. *Operations Research* 65 (3), 557–576.