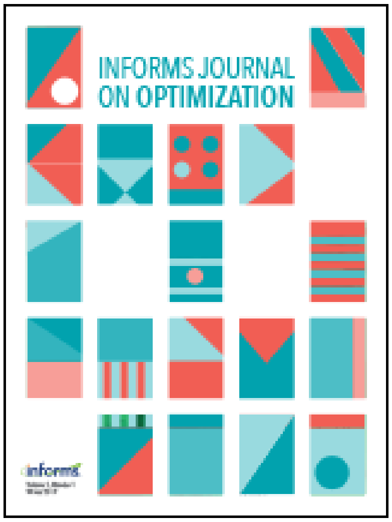

INFORMS Journal on Optimization

Pricing Wind: A Revenue Adequate, Cost Recovering Uniform Price Auction for Electricity Markets with Intermittent Generation

Golbon Zakeri, Geoffrey Pritchard, Mette Bjorndal, Endre Bjorndal

informs
http://pubsonline.informs.org/journal/ijoo/

INFORMS JOURNAL ON OPTIMIZATION
Vol. 1, No. 1, Winter 2019, pp. 35–48
ISSN 2575-1484 (print), ISSN 2575-1492 (online)

Pricing Wind: A Revenue Adequate, Cost Recovering Uniform Price Auction for Electricity Markets with Intermittent Generation

Golbon Zakeri,[a] Geoffrey Pritchard,[b] Mette Bjorndal,[c] Endre Bjorndal[c]

[a] Department of Engineering Science, University of Auckland, Auckland 1010, New Zealand; [b] Department of Statistics, University of Auckland, Auckland 1010, New Zealand; [c] Department of Business and Management Science, Norwegian School of Economics, 5045 Bergen, Norway

Contact: g.zakeri@auckland.ac.nz, http://orcid.org/0000-0002-6888-8451 (GZ); g.pritchard@auckland.ac.nz, http://orcid.org/0000-0002-5660-4503 (GP); mette.bjorndal@nhh.no (MB); endre.bjorndal@nhh.no (EB)

Received: September 11, 2017
Revised: February 6, 2018; May 11, 2018
Accepted: May 20, 2018
Published Online in Articles in Advance: October 23, 2018

https://doi.org/10.1287/ijoo.2018.0002

Abstract. With greater penetration of renewable generation, the uncertainty faced in electricity markets has increased substantially. Conventionally, generators are assigned a predispatch quantity in advance of real time based on estimates of uncertain quantities. Expensive real-time adjustments then need to be made to ensure demand is met as uncertainty takes on a realization. We propose a new stochastic-programming market-clearing mechanism to optimize predispatch quantities given the uncertainties' probability distribution and the costs of real-time deviation. This model differs from similar mechanisms previously proposed in that predispatch quantities are not subject to any network or other physical constraints, nor do they have a dedicated financial settlement. We establish revenue adequacy in each scenario (as opposed to "in expectation"), welfare enhancement, and expected (but not with probability one) cost recovery (including deviation costs) for this market-clearing mechanism. We also establish that this market-clearing mechanism is social-welfare optimizing.

Funding: This work was supported by the Energy Education Trust of New Zealand.

Keywords: electricity markets • uniform pricing • locational marginal pricing • stochastic programming • wind power • regulation

1. Introduction

The increasing penetration of intermittent renewable (wind and solar) power generation has led to much discussion of how to integrate these uncertain sources into electricity markets. The essential challenge is to devise a single operational and trading framework that can accommodate traditional thermal-generation plants (large, inflexible units with slow ramp rates and lead times measured in hours), intermittent renewable plants (with unpredictable power fluctuations on time scales measured in minutes), and hydroelectric or other flexible plants (which help to integrate the other two generation types but have limited capacity and possibly high cost).

From a mathematical perspective, the problem may be naturally conceived as one of stochastic optimization. There is a conceptual first stage, in which inflexible plants must be scheduled before the demands that will be placed on it are fully known, and a second stage in which the preparations made in the first stage must be adapted to the immediate needs of a realized situation. The first stage is often identified with a day-ahead market (although the actual time scale may be something other than 24 hours), the second with a real-time or balancing market. Recently, several authors (e.g., Bouffard et al. 2005, Wong and Fuller 2007, Morales et al. 2009, Rud 2009, Pritchard et al. 2010, Morales et al. 2014, and Zavala et al. 2017, Bjorndal et al. 2018) have presented reformulations of the traditional optimal power flow problem as two-stage stochastic programs of this sort. The models in the existing literature, however, do not entail the set of desirable properties that the stochastic auction designed in this paper entails.

The starting point of this paper is the observation that the main function of the first stage is to produce a primal, rather than a dual, solution. Even though day-ahead dispatch is most commonly implemented as a separate market in its own right, with its own prices and settlement, its real purpose is not to trade energy or set prices, but to produce a physical dispatch that anticipates and hedges against the exigencies of real time in the best possible way.

Day-ahead markets do also enable their participants to manage risk (at least for those risks relating to the last 24 hours before real time) and to forecast (second-stage) prices, but from an optimization viewpoint, these are secondary functions that can be—and in practice often are—achieved in other ways. Indeed, deregulated electric power sectors everywhere have found it necessary to invent still other markets for trading risk derivatives and other forecasting mechanisms—both clear signs that day-ahead markets are a rather imperfect tool for these jobs.

Having cast a captious eye over the first-stage prices, we may then notice that the constraints with which they are associated are also somewhat questionable. In existing literature, the first-stage primal solution is designed to comply with physical system constraints (imposed by transmission line capacities, security, etc.) even though it is understood that these constraints really apply only to an as-yet-unknown modification of this solution and that the unmodified solution will never be physically implemented. (Inasmuch as the initial dispatch will be modified later, it could be thought of as "nonphysical" although this does not, of course, mean that it is unimportant.)

In this paper, we develop a problem formulation that eliminates both first-stage physical constraints and the prices associated with them, leaving the primal solution as the only interpreted output of the first stage and relegating all financial transactions to the second stage. Perhaps surprisingly, this formulation turns out to have several desirable properties.

In general, these kinds of stochastic-programming (SP) models exhibit revenue adequacy in expectation (see Wong and Fuller 2007, Morales et al. 2009, and Pritchard et al. 2010). That is, the total expected payments made by consumers will equal or exceed the total expected compensation paid to generators. However, revenue adequacy need not occur in every scenario; an example is given in Pritchard et al. (2010). The model proposed in the present paper is free from this defect.

Another desirable property is cost recovery: the compensation paid for each dispatched energy tranche to the firm that offered it should equal or exceed the cost implied by the way it was offered. If this cost is assumed (as, e.g., in Morales et al. 2009) to reflect the supplier's actual costs, the implication is that the system will not create financial losses for suppliers. Our model presents a uniform pricing model that is proved to have cost recovery in expectation, including the cost of any deviations.

We demonstrate here that the decisions made by an independent system operator, who clears the market based on maximizing expected total welfare, align with those made by individual agents as long as those agents are truthful (i.e., a competitive equilibrium is reached) as well as risk neutral.

One can also ask the usual question asked of a stochastic program; that is, do the above properties remain valid when the sample incorporated in the stochastic program is not the full underlying distribution? Revenue adequacy is maintained for our model even if we encounter a scenario that we had not included in the sample. Furthermore, we establish an asymptotic result that demonstrates expected cost recovery as long as the chosen sample is "rich enough" in being reflective of the true distribution.

2. The New SP Model

We assume a power system modeled by DC load flow approximation, consisting of a collection of nodes connected by lines. The lines are assumed lossless so that, if F is a vector of flows on the lines, the net power $\tau_n(F)$ imported into node n by F is a linear function of F. The capabilities of the transmission system are represented by the requirement $F \in U$; we assume U to be a bounded convex polyhedral set with $0 \in U$.

The supply side of the market will comprise a finite collection of tranches with tranche j consisting of a quantity g_j of energy offered at a node $n(j)$. Note that a single-generation plant may (and in practice often will) offer multiple tranches, which together make up a stepwise supply function. Let $T(n)$ denote the set of tranches (from all suppliers) offered at node n.

2.1. The Real-Time Problem

Although the problem formulation we have in mind has two stages, we begin by considering the second stage alone as much of the analysis does not depend on any details of the first stage. Presentation of the full two-stage stochastic program is deferred to Section 2.3.

We represent real-time dispatch by the following optimal power flow problem.

$$\begin{array}{lll} [RT(x,d)]\text{: } \min_{X,F} & \sum_j (c_j X_j + r_j^+(X_j - x_j)_+ + r_j^-(X_j - x_j)_-) & \\ \text{s.t.} & \tau_n(F) + \sum_{j \in T(n)} X_j = d_n & \forall n \quad [\pi_n] \\ & 0 \le X_j \le g_j & \forall j \\ & F \in U. & \end{array} \tag{1}$$

Here X_j represents the quantity of energy to accept (dispatch) from tranche j (for each j), and F is the vector of associated line flows. The notations t_+ and t_-, for a scalar quantity t, denote $\max(t, 0)$ and $\max(-t, 0)$, respectively. The parameter $d = (d_n)$ gives the demand for energy at each node. We initially regard the demands as fixed parameters; later (Section 2.3), they are random.

The parameter $x = (x_j)$ reflects a nonstandard offer structure: each offered tranche j has an associated per-unit cost (or ask price) c_j, with additional costs applying when the dispatch exceeds (r_j^+) or falls short of (r_j^-) the level x_j.

Here x_j represents a previously established setpoint or predispatch, and the objective terms including it represent the "ramping" costs of making real-time adjustments away from this level. (Ramping constraints, of the form $x_j - a_j \le X_j \le x_j + b_j$ for some a_j and b_j, could also have been included without difficulty, but for the sake of simplicity, they do not appear in the present paper.) The reader may also observe that the assumed form of the adjustment costs closely resembles the expression for conditional value at risk, suggesting that the objective of $RT(x,d)$ might also be thought of as a measure of suppliers' risk; in this regard, see Philpott et al. (2016). In a market setting, both the c_j and the r_j^+, r_j^- are to be regarded as offer parameters nominated by suppliers. We assume $0 \le x_j \le g_j$, $r_j^+ \ge 0$, and $r_j^- \ge 0$, but it is possible that $c_j < 0$.

Note that $RT(x,d)$ can be written as a linear program. It is a variant of the standard dispatch problem considered, for example, in Philpott and Pritchard (2004). In the present paper, we limit ourselves to the LP form to simplify the mathematical analysis, but we note that a conceptually simple generalization to a convex program would allow for such things as quadratic transmission line losses or a more general form for the real-time adjustment costs. The latter might be quite desirable as real power-plant operating costs can be much more complex than the simple approximation assumed here.

Note that, if $r_j^+ = r_j^- = 0$, then $RT(x,d)$ reduces to a standard dispatch problem with constant marginal cost for each tranche offered. The interpretation of this case is that all supply is infinitely flexible, and the cost of deploying "reserve" is nothing but the actual production cost. In this case, the real-time dispatch problem is the only one required, and any day-ahead activity would be redundant.

We assume in this paper that $RT(x,d)$ is always feasible; this can be achieved by (for example) allowing unlimited load-shedding—which is effectively a high-priced offer of supply—at any node. Although feasibility achieved in this manner hardly addresses the concerns of system operators (or their customers), it does serve to dispose of the mathematical difficulties that would be created if we had to deal with infeasible problems. To the same end, we have assumed no minimum take from any offer (i.e., the lower bound on X_j is zero) although such a bound would not otherwise be difficult to incorporate.

It should be noted that the real-time adjustment rates r_j^+ and r_j^- apply only to adjustments within a given partially dispatched tranche. Larger adjustments may bring other tranches into (or out of) play, increasing the overall cost of adjustment.

The introduction of the adjustment costs means that there is no longer a simple "merit order" in which the offered tranches are dispatched. This in itself is nothing new: transmission network constraints can cause expensive tranches to be dispatched ahead of cheaper ones even in the conventional DC optimal power flow (DCOPF). However, we now have no particular merit order even among tranches representing supply from a single-generation unit, that is, a single machine. The reader will appreciate that this creates an implementation problem: if a single unit divides its capacity into multiple tranches for offering purposes, the cost parameters r_j^+, r_j^-, and c_j will have been determined for each tranche under the (possibly mistaken) assumption that certain other ("lower-cost") tranches have already been fully dispatched. But the solution of Equation (1) may dispatch the tranches in some other way, which may not even be physically realizable. In the present paper, we skirt this issue by assuming that when a unit divides its offer in this way, the parameter r_j^+ has the same value for all tranches offered by that unit; similarly, r_j^-. The tranches are then uniquely ordered by their c_j values (which without loss of generality are distinct) and are cleared by Equation (1) in that order.

Now suppose that (X^*, F^*) is a primal optimal solution of $RT(x,d)$ and that $\pi = (\pi_n)$ is the vector of optimal dual prices for the energy-balance constraints (i.e., energy prices at the nodes). We assume that energy is traded at the prices π_n with no side payments. That is, the consumers at node n collectively pay $\pi_n d_n$ for their demand d_n, and the provider of tranche j receives $\pi_{n(j)} X_j^*$ for the energy it supplies.

According to the Lagrangian duality theorem (see, e.g., Whittle 1971), the Lagrangian

$$\begin{aligned} L_\pi(X,F) = {} & \textstyle\sum_j (c_j X_j + r_j^+ (X_j - x_j)_+ + r_j^- (X_j - x_j)_-) \\ & + \textstyle\sum_n \pi_n \left(d_n - \tau_n(F) - \sum_{j \in T(n)} X_j\right) \end{aligned} \tag{2}$$

is minimized, subject to the remaining constraints $0 \le X_j \le g_j\ \forall j$ and $F \in U$ at (X^*, F^*). This leads to several useful observations.

Proposition 1 (Revenue Adequacy). *The revenue raised from consumers is sufficient to pay the suppliers. That is,*

$$\sum_n \pi_n d_n \ge \sum_j \pi_{n(j)} X_j^*.$$

Proof. Considering the terms containing F in Equation (2), we see that $\sum_n \pi_n \tau_n(F)$ is maximized over $F \in U$ at F^*. By comparison with the feasible point $0 \in U$ (all zero flows), this gives $\sum_n \pi_n \tau_n(F^*) \geq 0$. But, because (X^*, F^*) satisfies the energy-balance constraints in $RT(x,d)$, $\tau_n(F^*) = d_n - \sum_{j \in T(n)} X_j^*$. The result follows. ■

Note that this result holds for all values of the parameters x and d. In particular, it is not necessary to choose the set points x in any particular way; that is, revenue adequacy in each scenario is guaranteed as we simply solve an ordinary OPF in real time.

Proposition 2 (Pricing and Dispatch of Individual Offers). *The dispatch of tranche j is related to the energy price at its local node in the following way:*

If $\pi_{n(j)} < c_j - r_j^-$,	*then* $X_j^* = 0$.
If $\pi_{n(j)} = c_j - r_j^-$,	*then* $0 \leq X_j^* \leq x_j$.
If $c_j - r_j^- < \pi_{n(j)} < c_j + r_j^+$,	*then* $X_j^* = x_j$.
If $\pi_{n(j)} = c_j + r_j^+$,	*then* $x_j \leq X_j^* \leq g_j$.
If $\pi_{n(j)} > c_j + r_j^+$,	*then* $X_j^* = g_j$.

Proof. Considering the terms containing X_j in Equation (2), we see that for each j,

$$(c_j - \pi_{n(j)})X_j + r_j^+(X_j - x_j)_+ + r_j^-(X_j - x_j)_-$$

is minimized over $0 \leq X_j \leq g_j$ at X_j^*. The result follows. ■

Corollary 1 (Supplier Margins). *The supplier's margin on tranche j (i.e., the difference between the market revenue received for it and the cost terms corresponding to it in the objective of Equation (1)) is*

$$m_j(x_j, \pi_{n(j)}) = \begin{cases} -r_j^- x_j, & \text{if } \pi_{n(j)} \leq c_j - r_j^- \\ (\pi_{n(j)} - c_j)x_j, & \text{if } c_j - r_j^- \leq \pi_{n(j)} \leq c_j + r_j^+ \\ (\pi_{n(j)} - c_j - r_j^+)g_j + r_j^+ x_j, & \text{if } \pi_{n(j)} \geq c_j + r_j^+. \end{cases} \tag{3}$$

In particular, the margin is positive if $\pi_{n(j)} > c_j$ *and negative if* $\pi_{n(j)} < c_j$.

Proof. The margin made by the supplier of tranche j will be

$$\pi_{n(j)}X_j^* - c_j X_j^* - r_j^+(X_j^* - x_j)_+ - r_j^-(X_j^* - x_j)_-.$$

In light of Proposition 2, this margin can be expressed as in Equation (3). ■

Remark 1. If the cost terms for tranche j in the objective of Equation (1) are assumed to reflect suppliers' actual costs, then the supplier's margin is its gross profit.

In the standard DCOPF market formulation (see, e.g., Philpott and Pritchard 2004), supplier margins are always nonnegative. This is intuitively clear because, when a tranche is dispatched, it either sets the price at its location (so that its supplier receives a zero margin over the offered price) or it is infra-marginal and receives a positive markup. This is not the case in the present paper.

The possibility of a negative margin (i.e., of the market accepting a supplier's offer in a way that fails to cover the supplier's costs) is a new and undesirable feature of our model. It is, however, a necessary property of any single-settlement, energy-only market of this kind. Consider, for example, the situation in which a tranche is dispatched at the first stage ($x_j > 0$) but not at the second stage ($X_j^* = 0$). The supplier of such a tranche cannot receive any revenue (no matter how the prices are computed) but may incur setup costs to achieve the indicated x_j.

Note also from Equation (3) that the supplier's margin $m_j(x_j, \pi_{n(j)})$ is uniquely determined (provided the local nodal price $\pi_{n(j)}$ is uniquely determined) even in degenerate cases in which X_j^* is not uniquely determined. Such a case may occur, for example, if two identical tranches are offered at the same node.

2.2. The Dual Real-Time Problem

Now consider the dual $DRT(x,d)$ of the linear program $RT(x,d)$. When we express $RT(x,d)$ as a linear program, the parameters x and d are constraint right-hand sides; thus, they appear in $DRT(x,d)$ only as objective coefficients. That is, the feasible set S of $DRT(x,d)$ does not depend on x or d. From this observation, we can derive a result that is useful later in the paper and may also be of independent interest.

Lemma 1. *There exists a constant C, depending only on the feasible set S of $DRT(x,d)$, such that for any x, the set*

$$\{d : RT(x,d) \text{ has unique energy price duals satisfying } |\pi_n| \le C\}$$

has full Lebesgue measure.

Remark 2. This result offers reassurance that the prices generated by the proposed mechanism are both well determined and bounded unless the demands experienced are exceptional in the sense of belonging to a measure-zero subset of the demand space.

Proof. Because S is a (possibly unbounded) polyhedral set with finitely many vertices, we can choose C so that $\|v\|_\infty \le C$ for all vertices v of S. The optimal set of any problem $DRT(x,d)$ will be a face f (possibly just a single vertex) of S; if the energy price variables (π_n) are uniquely determined (i.e., constant on f), then they will satisfy $|\pi_n| \le C$. It, therefore, suffices to show that, for a fixed x, the set

$$\Delta_x = \{d : RT(x,d) \text{ has nonunique energy price duals } (\pi_n)\}$$

has Lebesgue measure zero.

For each face f of S on which the (π_n) variables are not constant, choose points v_f, w_f differing in at least one π_n component. When $DRT(x,d)$ has such a face f as its optimal set, its objective vector is orthogonal to $v_f - w_f$. Because this objective vector includes the demands d_n as coefficients corresponding to the variables π_n, we have $a_f^T d + b_f = 0$ for some nonzero vector a_f and scalar b_f. So

$$\Delta_x \subseteq \bigcup_f \{d : a_f^T d + b_f = 0\}.$$

Because S has only finitely many faces, the latter set is a finite union of hyperplanes with Lebesgue measure zero. ∎

2.3. The Stochastic Programming Problem

We turn now to the problem of selecting optimal set points x prior to solution of the real-time problem (1). Let $J(x,d)$ denote the optimal value of the problem $RT(x,d)$. Consider the first-stage problem

$$[FP(\mu)]: \quad \min_x E\left[J(x,D)\right], \tag{4}$$

where D is a random vector of demands with distribution μ. Note that this is an unconstrained optimization: we place no a priori constraints on the decision variable x.

Although the motivation for the present paper was the presence of random generation, only the demands can be random in Equation (4). This is because the uncertain wind- and solar-power sources that we have in mind usually have low or zero marginal costs of production so that they behave more like random (negative) demands than traditionally offered generation. The form of Equation (4), thus, accommodates them in a natural way: D should be regarded as a residual demand, net of some uncertain generation that has been treated as demand. The main disadvantage of such a viewpoint is that it ignores the possibility of curtailing the uncertain generation, something that might become desirable in situations in which the system is oversupplied, and prices fall to very low or negative levels. With this in mind, we remark that it is possible to model intermittent generation explicitly in our formulation with the offer quantities g_j in Equation (1) becoming random variables. Although we have not pursued this possibility in our main results, it is illustrated in Example 3.1.

We can bring Equations (1) and (4) together as the two-stage stochastic programming problem

$$\begin{array}{llll} [SP]: \min_{x,X,F} & E\left[\sum_j \left(c_j X_j + r_j^+ (X_j - x_j)_+ + r_j^- (X_j - x_j)_-\right)\right] & & \\ \text{s.t.} & \tau_n(F) + \sum_{j \in T(n)} X_j = D_n & \forall n \text{ w.p.1 } [\pi_n] & \\ & 0 \le X_j \le g_j & \forall j \text{ w.p.1} & \\ & F \in U & \text{w.p.1.} & \end{array} \tag{5}$$

Here the demands D_n, the second-stage decision variables X_j and F, and the second-stage dual variables π_n are all random variables; that is, they have values depending on an outcome ω selected from a sample space Ω. To simplify notation, we follow the usual convention of suppressing the dependence on ω for random quantities, thus, for example, F rather than $F(\omega)$. Constraints at the second stage are to be satisfied with probability one (w.p.1).

Note again that the variables x_j are unconstrained; the first-stage solution is not explicitly required to meet demand or satisfy other physical constraints. This differs from previous treatments (e.g., Pritchard et al. 2010) and has implications for prices. However, we have $0 \le x_j \le g_j$ (as assumed in Section 2.1) as a consequence of Theorem 1.

Let (x^*, X^*, F^*) be optimal for Equation (5); then x^* is optimal for Equation (4) (see Dantzig and Madansky 1961 or Bouffard et al. 2005). As with the real-time problem, Lagrangian duality (for this case, see Shapiro et al. 2009) can be used to establish properties of the solution. The Lagrangian of Equation (5)

$$L_\pi(x, X, F) = E\left[\sum\nolimits_j (c_j X_j + r_j^+ (X_j - x_j)_+ + r_j^- (X_j - x_j)_-)\right] \\ + E\left[\sum\nolimits_n \pi_n \left(D_n - \tau_n(F) - \sum\nolimits_{j \in T(n)} X_j\right)\right] \tag{6}$$

is minimized, subject to the remaining constraints $0 \le X_j \le g_j \ \forall j$ w.p.1 and $F \in U$ w.p.1 at (x^*, X^*, F^*). This leads to some further useful observations.

Theorem 1 (Optimal Set Points Are Quantiles of Real-Time Dispatch). *The optimal set point x_j^* for tranche j is a $\frac{r_j^+}{r_j^+ + r_j^-}$ quantile of the probability distribution of X_j^*.*

Remark 3. See also the similar result (theorem 4) in Zavala et al. (2017). It should be noted that X_j^* is not an exogenous random variable; rather, both it and x_j^* depend on the parameters of Equation (5). Theorem 1 should, thus, be regarded as describing the structure of solutions to Equation (5).

Proof. Considering the terms containing x_j in Equation (6), we see that

$$E\left[r_j^+ (X_j^* - x_j)_+ + r_j^- (X_j^* - x_j)_-\right]$$

is minimized over the unconstrained variable x_j at x_j^*. If $r_j^+ = r_j^- = 0$, this is a trivial result, but otherwise, the problem is a variant of the well-known newsvendor stochastic optimization problem. The result follows as in, for example, Koenker (2005) and Shapiro et al. (2009). ■

Theorem 2 (Expected Supplier Margins). *The expected supplier margin on each tranche is nonnegative.*

Proof. Considering the terms containing x_j or X_j in Equation (6), we see that

$$E\left[(c_j - \pi_{n(j)}) X_j + r_j^+ (X_j - x_j)_+ + r_j^- (X_j - x_j)_-\right]$$

is minimized, subject to $0 \le X_j \le g_j$ w.p.1 (and x_j unconstrained) at (x_j^*, X_j^*). It follows by comparison with the feasible solution $x_j = X_j = 0$ that

$$E\left[(c_j - \pi_{n(j)}) X_j^* + r_j^+ (X_j^* - x_j^*)_+ + r_j^- (X_j^* - x_j^*)_-\right] \le 0.$$

That is, the *expected* margin made by the supplier of tranche j (see Corollary 1) is nonnegative. ■

Remark 4. This result offers some solace to suppliers: although this market model may deliver them negative margins on occasion, their margins are at least nonnegative in expectation.

2.4. A Single-Settlement Market

In this subsection, we make some remarks about how the model in this paper might translate into the operation of an electricity market.

Our model has two stages of optimization, but only the second stage has prices. If the two stages are thought of as “day-ahead” and “real-time,” it is only the real-time stage that constitutes a market in the usual sense with goods (energy) being traded and financial settlement made. There is no day-ahead market of the kind that exists in various real-world deregulated power systems; all power is effectively traded in real time.

It is worth considering how the usual functions of a day-ahead market—risk management and price discovery—might be accomplished in a deregulated system that does not include one. The best existing model for this may be the New Zealand Electricity Market (NZEM), which has functioned effectively without a day-ahead market since 1996. The risk-management function is accomplished, in the New Zealand context, via separate markets for energy hedges (contracts for differences); it does not appear to be necessary for these to be deeply integrated with the primary market. Price discovery occurs via the publication of predispatch schedules in the period leading up to "gate closure" (the offer submission deadline) two hours in advance of real time.

One important advantage of the day-ahead market is its ability to synthesize collective forecasts from the trading activity of its participants. This paper does not consider how the demand probability distribution required in Equation (5) is to be obtained, but if it is devised by the system operator alone, then it may be of lower quality than a "crowdsourced" forecast with many contributors. A possible solution might be a separate market for short-term (one day) futures and other financial derivatives.

Returning to the model of the present paper, we observe that its first stage has primarily a physical, rather than an economic, function. Its purpose is to produce a primal solution only: the optimal set points x. Although there is no day-ahead price to inform the market, there can be something arguably better: an ensemble forecast of possible real-time prices computable once the set points are known.

The possibility of negative margins suggests a need for uplift payments to suppliers to ensure that their costs are recovered. The most straightforward approach would be a simple uplift payment in each market-trading period equal to the negative part of the supplier's margin on each tranche. Alternatively, the result of Theorem 2 suggests a way to make the uplift payments smaller: aggregate the margin over multiple market-trading periods and make an uplift payment equal to the negative part of the total. As the number of trading periods aggregated over increases, the required uplift payments should tend to zero.

It should be noted that the latter approach is not entirely consistent with the ideas of Stoft (2002) regarding appropriate incentives for longer-term investment. However further discussion of this point is beyond the scope of the present paper.

2.5. Competitive Equilibrium

The Lagrangian (6) offers a way to regard the optimum as the solution to a competitive game. It may be expressed as

$$\begin{aligned} L_\pi(x, X, F) = &-\sum\nolimits_j E[(\pi_{n(j)} - c_j)X_j - r_j^+(X_j - x_j)_+ - r_j^-(X_j - x_j)_-] \\ &- E\Big[\sum\nolimits_n \pi_n \tau_n(F)\Big] + E\Big[\sum\nolimits_n \pi_n D_n\Big]. \end{aligned} \tag{7}$$

Consider a game in which each energy tranche j is offered by an agent A_j desiring to maximize the supplier margin on that tranche while an additional agent A_{trans} controls the transmission system and desires to maximize the transmission constraint rental $\sum_n \pi_n \tau_n(F)$. If the game is perfectly competitive and all agents are risk-neutral (so that expected values will serve as their objectives) then Equation (7) shows that the optimum (x^*, X^*, F^*) is an equilibrium for the game.

In economic language, the competitive solution is social-welfare maximizing. However, this assertion presupposes that the system operator's view of the future (including the probability distributions and risk preferences implied by the problem formulation) is identical to that of the market participants generally. The reader may wonder if a similar result could be proved when the agents in the market have different views and risk preferences. In this case, if the market for trading risk is complete (i.e., there are enough instruments, such as hedge contracts, available to agents), then the results obtained in Ralph and Smeers (2015) can be applied. In this case, we can demonstrate that the risk-adjusted system position will coincide with the risk-averse, competitive equilibrium.

3. Illustrative Examples

In this section, we lay out two simple examples to illustrate the properties of our model.

3.1. Single-Node Example

Consider a single-node system in which two firms and one intermittent supplier offer in energy and a deterministic demand of 100 MW. The intermittent supplier (firm I1) faces two equally likely production scenarios, one low and one high. The cost of generation and ramp up and down costs as well as capacities of generation are provided in Table 1.

Table 1. Offer Costs (in \$) and Quantities (in MW) for All Firms

Firm	Cost of generation	Ramp up	Ramp down	Capacity
F1	10	1	1	50
F2	20	5	0.001	80
I1	0.01	1	0.001	20 in scenario 1 50 in scenario 2

The dispatch problem can then be formulated as

$$
\begin{aligned}
\text{[Ex1]}\quad \min_{X,x} \sum_{\omega\in\{1,2\}} \frac{1}{2}\Big(& 10X_1^\omega + 20X_2^\omega + 0.01X_3^\omega \\
& +1(X_1^\omega - x_1)_+ + 1(X_1^\omega - x_1)_- \\
& +5(X_2^\omega - x_2)_+ + 0.001(X_2^\omega - x_2)_- \\
& +1(X_3^\omega - x_3)_+ + 0.001(X_3^\omega - x_3)_-\Big) \\
\text{s.t.}\quad & X_1^\omega + X_2^\omega + X_3^\omega = 100 && \omega\in\{1,2\} \\
& 0 \le X_1^\omega \le 50 && \omega\in\{1,2\} \\
& 0 \le X_2^\omega \le 80 && \omega\in\{1,2\} \\
& 0 \le X_3^1 \le 20 \\
& 0 \le X_3^2 \le 50 \\
& x_i \ge 0 && i\in\{1,2,3\}.
\end{aligned}
$$

The solution to this problem is outlined in Table 2.

We note that solving the expected value problem by tailoring the dispatch to the expected wind generation will yield advisory positions of 50, 15, and 35 for F1, F2, and I1, respectively. When faced with scenario 1, that is, wind availability of only 20 MW, the system would have to resort to dispatching generator 2 and incur \$75.00. The value of stochastic solution in this problem is then approximately 5%.

For this problem, the optimal value is \$800.38. The stochastic programming duals for the demand balance constraints in scenarios 1 and 2 are \$10.0005 and \$9.9995, respectively. As in Dantzig and Madansky (1961), these duals (already available from the stochastic program), scaled by the corresponding scenario's probability, will produce the real-time prices that are used for financial transactions in each scenario. These are 20.001 and 19.999, respectively. Table 3 outlines the revenues, costs, and profits of the generators in this simple example.

This example demonstrates the mechanics of the proposed market-clearing and settlement schemes. It is clear here that we maintain revenue adequacy in each scenario and that generators recover cost in expectation. We also point out that if a "physical" first-stage constraint, such as proposed in Pritchard et al. (2010), is added to the model, welfare decreases. Here, if we ensure that that first-stage advisory position must meet demand, that is, $\sum_i x_i = 100$, then the optimal dispatch cost increases to \$815.37.

3.2. Two-Node Example

The example depicted in Figure 1 is taken from Pritchard et al. (2010). It has two inelastic demand scenarios laid out in Table 4 (this can be thought of as demand net of renewable generation), which, although they agree as to the total load, differ markedly in the location of the load. The thermal generation offer is completely inflexible (requiring $X_i(\omega_1) = X_i(\omega_2) = x_i$), and the hydros are completely flexible with indicated deviation costs (2.00 dollars for either upward or downward deviation per MW). The line is lossless.

The optimal advisory position for this problem is given by $x^* = (0, 3, 5)$, that is, three units from the thermal and all five units of hydro 2. The optimal real-time dispatches for this example are given by $X^*(\omega_1) = (0, 3, 5)$ for scenario 1 and $X^*(\omega_2) = (1, 3, 4)$ for the second scenario. Furthermore, the nodal prices are given by $\pi_A^{\omega_1} = \pi_B^{\omega_1} = 15\frac{1}{3}$ for scenario 1 and $\pi_A^{\omega_2} = 27$ and $\pi_B^{\omega_2} = 8$ for scenario 2. For this example, the total payment in scenario 1 is $\$122\frac{2}{3}$, which is precisely what is collected from the consumer. In scenario 2, the payment to generators adds up to \$140, and the collected revenues from demand is \$197. The difference between the collected revenues and payments to the generators is entirely a result of congestion rent in this scenario. Clearly, the settlement mechanism is revenue adequate in each scenario. We note that the treatment of this example in Pritchard et al. (2010) finds that the

Table 2. Advisory and Real-Time Generation (in MW) for the Firms

Firm	Advisory position	Real-time generation
F1	50	50 in scenario 1 50 in scenario 2
F2	30	30 in scenario 1 0 in scenario 2
I1	50	20 in scenario 1 50 in scenario 2

Table 3. Computation of Payments (in $) to the Generators

Firm and scenario	Revenue	Cost	Profit	Expected profit
F1 scenario 1	1,000.05	500	500.05	500
F1 scenario 2	999.95	500	499.95	
F2 scenario 1	600.03	600	0.03	0
F2 scenario 2	0	0.03	−0.03	
I1 scenario 1	400.02	0.23	399.79	699.62
I1 scenario 2	999.95	0.5	999.45	

collected revenues (under the settlement scheme proposed in that paper) in the first scenario are insufficient to cover the participant payments.

4. Continuity Properties

The probability distribution of the demands D in Equation (4) is inevitably a modeling approximation to the real world. In practice, it is likely to be a discrete distribution, consisting of a finite collection of scenarios with attached probabilities. In this section, we consider the possibility of the actual demand vector being drawn from a "true" probability distribution different from the one used in Equation (4) to determine set points x. In particular, this affects the result of Theorem 2: expected supplier margins are nonnegative *if* the set points have been chosen optimally for the demand distribution that will actually be experienced.

Theorem 3. *Let $\mu_1, \mu_2, \ldots$ be probability distributions of demand, and suppose that the μ_k converge in distribution to a limiting distribution μ with the supports of μ and the μ_k all lying within a bounded set. Let x^{k*} be a minimizer of $FP(\mu_n)$. Then any limit point x^* of $\{x^{k*}\}$ is a minimizer of $FP(\mu)$.*

Remark 5. Our intended interpretation of this result casts μ in the role of the true distribution of demand and the μ_k as modeling approximations thereof. The result suggests that the minimizers of $FP(\mu_k)$ may be regarded as approximations to a minimizer of $FP(\mu)$.

Proof. Apply theorem 12 of chapter 9 in Birge and Louveaux (1997). The uniform bound on the supports of μ_k, μ and the fact that $J(\cdot, D)$ is convex and finite-valued with probability one ensure that the required hypotheses are satisfied. ∎

We turn now to the supplier margins considered in Corollary 1 and Theorem 2. Let $M_j(x, d)$ be the margin made by the supplier of tranche j in problem $RT(x, d)$; recall that this margin may be expressed as $m_j(x_j, \pi_{n(j)})$, where m_j is the function given in Equation (3). Note that m_j is a continuous function of two variables.

Theorem 4. *Let the demand vector D be governed by a probability distribution μ, which has bounded support and is absolutely continuous with respect to Lebesgue measure. Then $E[M_j(x, D)]$ is continuous in x.*

Remark 6. Combining Theorems 3 and 4 offers solace to suppliers even when the modeled distribution of demand differs from the true one. The expected supplier margin—which, according to Theorem 2, is nonnegative for the stochastic program with demand distribution modeled exactly—varies continuously as the modeled demand distribution (and associated optimal set points) are varied away from this ideal.

Proof. Let the sequence (x^k) converge to x. Let the demand vector d be such that $RT(x, d)$ has uniquely determined prices (π_n). (According to Lemma 1, the random vector D will take on such a value with probability one.) As noted

Figure 1. A Two-Node System with Random Loads

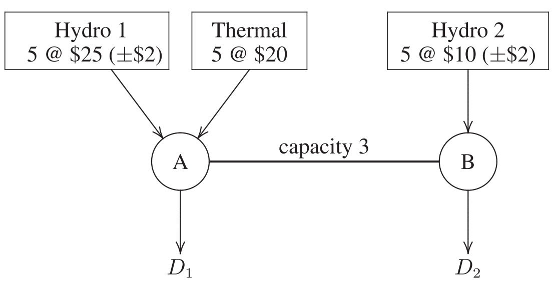

Table 4. Scenarios for the Two-Node Problem

Scenario	Probability	D_1	D_2
ω_1	0.6	2	6
ω_2	0.4	7	1

in Section 2.2, the feasible set S of $DRT(x, d)$ does not depend on either x or d. It follows from this that, for sufficiently large k, $DRT(x^k, d)$ will have its optimum on the same face of S as $DRT(x, d)$; in particular, the energy prices (π_n) will be the same for $RT(x^k, d)$ as they are for $RT(x, d)$. We then have

$$m_j(x_j^k, \pi_{n(j)}) \to m_j(x_j, \pi_{n(j)})$$

by continuity of m_j. That is, $M_j(x^k, d) \to M_j(x, d)$. This gives $M_j(x^k, D) \to M_j(x, D)$ with probability one. The result follows by the dominated convergence theorem with Lemma 1 providing the necessary boundedness for the duals. ■

The result of Theorem 4 does not hold if the demand distribution is discrete rather than continuous. It may even fail if the components d_n of demand are discretely distributed at some nodes n and continuously distributed at other nodes; see the counterexample to this effect in Sen (1993).

5. Numerical Results

As mentioned in the previous section, the probability distribution of the demands (or demands net of renewable generation) D in Equation (4) is inevitably an approximation to the real world. Although Theorem 4 offers some reassurance regarding the effects of the approximation errors, it is also worthwhile to have realistic computational evidence to demonstrate these results. To this end, we undertook a comprehensive study on the full New Zealand system, for which the data and the system parameters are readily available. We outline the experiments and report on the results.

5.1. The New Zealand System

The New Zealand electricity market, established in 1996, is one of the most mature in the world. The country consists of two islands, the north and south islands, each with a high-voltage AC power grid; the islands are interconnected by a high-voltage DC link. The NZEM is a single-settlement, energy-only market with no separate day-ahead market. For each half-hour period, generators may offer up to five (price, quantity) tranches per power plant; some demand-side bids, in the same format, are also present. Offers for each period may not be changed after a point two hours prior to the beginning of the time period; this point is referred to as gate closure.

The optimal power flow over the NZEM is a side-constrained network optimization problem that is solved for each half-hour trade period to determine dispatched quantities and locational marginal prices of electricity (see Alvey et al. 1998). Although the software (known as SPD for "scheduling, pricing and dispatch") that solves this problem is proprietary, the New Zealand Electricity Authority has created an invaluable research tool by providing a free replica of SPD along with all the input data necessary to replicate the final pricing run of each half hour over the past decade (see New Zealand Electricity Authority 2017). The replica that is used to build our stochastic programming model on is called vectorized scheduling pricing and dispatch (vSPD).

The analysis in this section considers uncertainty because of wind generation only. For the period under consideration, almost all of the electricity production from wind comes from the wind farms located in the central and southern (i.e., Wellington) regions of the north island as depicted in Figure 2 (see wind generation near Te Apiti and Mill Creek). The central north island region has about 300 MW of nameplate wind-generation capacity; the Wellington region has 200 MW. Together they contribute 4%–5% of the electricity production in New Zealand.

5.2. Design of the Experiments

To assess the policies delivered by the stochastic program and associated pricing scheme presented in this paper, we need to solve a number of two-stage stochastic programs derived from vSPD in conjunction with wind-production scenarios. As outlined in the modeling sections, we form a two-stage stochastic program in which the first stage set points are decided upon prior to gate closure and generation decisions are made in real time to meet the demand given that there is generation from wind in real time, and this varies by scenario. The second-stage

Figure 2. (Color online) Power Plants in the North Island of New Zealand

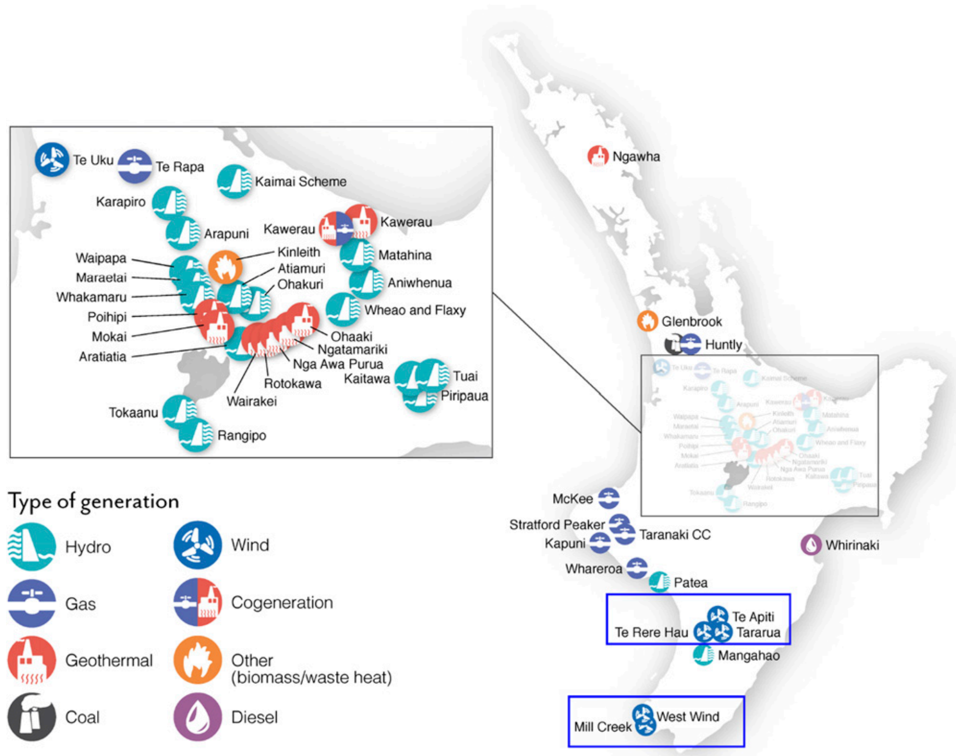

problems are, therefore, replications of vSPD for various scenarios of wind generation. In particular, the sets of constraints laid out in Equation (5) are defined by the vSPD constraints, so they capture meeting (net) demand in every scenario, respecting the tranche capacities, and include line capacity constraints and loop constraints as well as complying with piecewise linear losses (even though we had dispensed with this concept in the theory section of the paper).

To produce wind scenarios, we utilize the quantile regression method of Pritchard (2011) and Khazaei et al. (2014). This method uses separate quantile regressions to predict various quantiles of the near-future changes in wind power by nonlinear (cubic spline) functions of the present wind power; these predictions are then used as scenarios. In our experiments, we generate scenarios for the variation in wind power over the next two hours from

Figure 3. (Color online) A 12-Scenario Forecast Model for Changes in Wind Power from North Island (NI) Wind Farms over the Next Two Hours

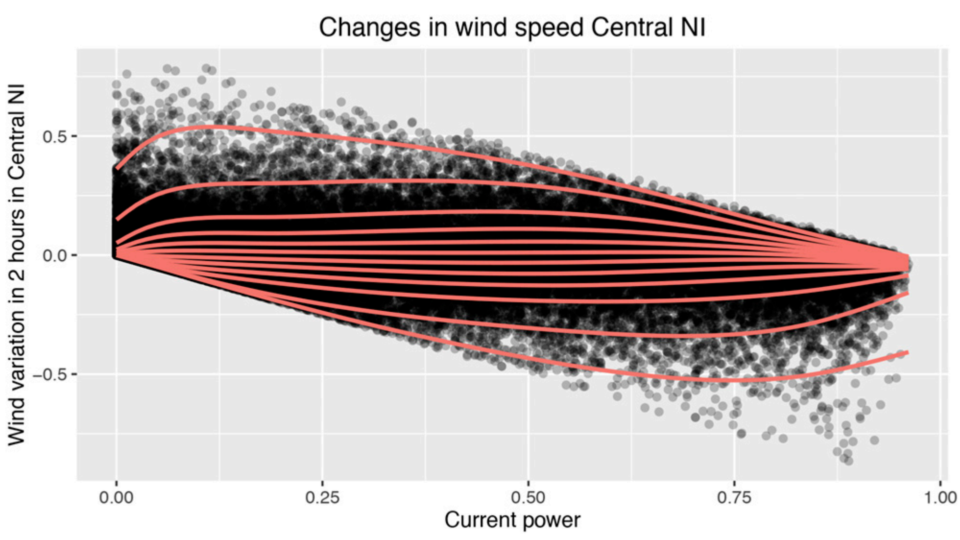

Table 5. Nonzero Shortfall Summary Statistics for January 1 to December 31, 2014

Statistic	$\kappa = 10$	$\kappa = 100$
Maximum shortfall	\$489.00	\$489.00
Third quarter shortfall	\$33.33	\$46.00
Median shortfall	\$13.800	\$31.50
Mean shortfall	\$31.80	\$48.65
First quarter shortfall	\$3.90	\$13.21
Minimum shortfall	\$0.058	\$0.20
Total shortfall	\$10,670.5	\$15,347.6

the 1st, 5th, 10th,..., 95th and 99th quantiles. We also treat the wind power available from the two regions as conditionally independent given their present power outputs; that is, the two-hour changes are independent. Correlation measurements and plots suggest this is a reasonable assumption. Our regression parameters are derived from an underlying data set of half-hour measurements for the period from January 1, 2011, to December 31, 2013.

Equipped with these scenarios and the full vSPD dispatch model, we form a two-stage stochastic program of the form (5) for the NZEM. We run our experiments for all periods of the calendar years 2014 and 2015 (35,040 half-hour trade periods); note that wind realizations during this time are out of sample with respect to the original data set. As a proxy for the costs of deviation from a preselected set point (r_j^+ and r_j^-), we use the reciprocals of the maximum possible ramp rates of each power plant, scaled by $\kappa = 10$ and $\kappa = 100$ (thereby rendering two different sets of experiments).

We use GAMS in conjunction with CPLEX (as the original vSPD model is writen in GAMS, this was natural for the modification to the two-stage stochastic programs) and implemented our code on a PC workstation. The stochastic programs solve in a short time; on average, it takes them about 30–40 seconds to solve with the largest solve time of under two minutes for the data examined. This methodology is thereby demonstrated to be implementable for a real system in real time. Once each stochastic program is solved, we obtain a first-stage (advisory) predispatch policy. We then obtain the real-time dispatch and prices by solving the real-time problem with the historically realized wind data for the period. We report on the shortfall considering every offer stack (so all generators and all nodes). More than 96% of the periods examined in 2014 and more than 97% of periods examined in 2015 demonstrated full cost recovery for all generators. Table 5 contains the (nonzero) shortfall from cost recovery if market settlement occurred in every half-hour period. The figures reported in Table 5 are based on the approximately 17,520 periods pertaining to the year 2014 but only take into account periods in which a shortfall actually occurred, that is, approximately 700 periods. Similarly, Table 6 contains the analogous information for the calendar year 2015. In the last row in each table, we report the total amount of shortfall incurred through each year.

Once we extend the settlement by aggregating payments over a 24-hour period, the total annual (nonzero) shortfall reduces dramatically. For 2014, this total annual shortfall is \$575 and \$988 for $\kappa = 10$ and $\kappa = 100$, respectively (this is as compared with the total shortfall figures reported in the last lines of Tables 5 and 6). In 2015, the total annual shortfall becomes \$140 and \$1,267 for $\kappa = 10$ and $\kappa = 100$, respectively.

The welfare enhancement, measured by the difference between the cost of generation resulting from the implementation of the SP policy, compared with historical generation, is positive but small, an approximate 0.1% improvement. This is because, even with the larger multiplier of $\kappa = 100$, the cost of deviation is very small compared with the cost of generation. Stochastic programming methodology is perhaps better suited to systems with less flexibility and higher penetration of renewables. In fact, our results confined to dry winters, when hydro

Table 6. Nonzero Shortfall Summary Statistics for January 1 to December 31, 2015

Statistic	$\kappa = 10$	$\kappa = 100$
Maximum shortfall	\$370.00	\$470.00
Third quarter shortfall	\$13.21	\$54.24
Median shortfall	\$5.40	\$34.50
Mean shortfall	\$10.88	\$50.29
First quarter shortfall	\$3.15	\$24.61
Minimum shortfall	\$0.26	\$0.53
Total shortfall	\$2,508.7	\$9,152.7

resources are scarce and we need expensive thermal ramps, report a fivefold to tenfold increase in the value of stochastic solution. However the full-scale NZEM data were readily available and valuable in gauging the implementability of our methodology, which led us to run our large-scale experiments to prove that the proposed stochastic program is implementable in real-world systems. Further details of our experiments and their study in the context of risk is the subject of a new report (see Cory-Wright and Zakeri 2018).

6. Conclusions and Future Directions

In this paper, we introduced an improvement on the recent stochastic programming–based market-clearing models for electricity markets with large penetration of renewables. We have established revenue adequacy in each scenario, expected cost recovery (for each agent), and incentive compatibility for our proposed model. We have also established asymptotic results that guarantee preservation of these properties for the model so long as the selected sample is rich enough. Our results establish a strict improvement over the stochastic-programming models and pricing mechanisms proposed in the literature thus far.

We have presented numerical results that are obtained from implementing our method on out-of-sample data from the NZEM. These results are the implementation of an SP methodology for individual periods; effectively a one step look ahead. Although this is clearly an improvement over the current deterministic model, it can be enhanced by considering multiple time periods in a multistage stochastic model. In such an extension, however, the set points introduced in this paper become constrained by the need to remain in proximity to previous operating points. Development of an effective stochastic multistage model is left as future work.

Other further developments of the model could include stochasticity in the parameters g_j (random tranche sizes—an explicit offer mechanism for intermittent generators) and/or U (random transmission capabilities, allowing for line outages). The first of these involves bringing randomness to additional constraint right-hand sides, so in principle, it is only a mild extension of the formulation in the present paper. The second is more difficult to envisage as it makes the constraint coefficient matrix a random variable.

A more ambitious extension would be to include integer variables or other nonconvexity in the first stage. Although it would certainly be desirable to model unit commitment, the considerable theoretical challenges put this idea well beyond the scope of the present paper. We note, however, that in one respect our approach seems well suited to such a development: the disregard for first-stage prices makes it unnecessary to consider how a nonconvex formulation would affect them.

Acknowledgments

The authors thank Professor Shmuel Oren for excellent questions that provoked this paper. The authors also thank Professor Alexander Shapiro for insightful comments that helped to establish the asymptotic results, Ryan Cory-Wright for computational work, and the paper's anonymous referees for various observations and additions.

References

Alvey T, Goodwin D, Ma X, Streiffert D, Sun D (1998) A security-constrained bid-clearing system for the New Zealand wholesale electricity market. *IEEE Trans. Power Systems* 13(2):340–346.

Birge J, Louveaux F (1997) *Introduction to Stochastic Programming* (Springer, New York).

Bjorndal E, Bjorndal M, Midthun K, Zakeri G (2018) Congestion management in a stochastic dispatch model for electricity markets. *Eur. J. Oper. Res.* Forthcoming.

Bouffard F, Galiana F, Conejo A (2005) Market clearing with stochastic security. *IEEE Trans. Power Systems* 20(4):1818–1826.

Cory-Wright R, Zakeri G (2018) On efficiency, savings, wealth transfers and risk aversion in electricity markets with uncertain supply. *Eur. J. Oper. Res.* Forthcoming.

Dantzig G, Madansky A (1961) On the solution of two stage linear programs under uncertainty. Neyman J, ed. *Proc. 4th Berkeley Symp. Math. Statist. Probab.*, Vol. 1, (University of California Press, Berkeley) 165–176.

Khazaei J, Zakeri G, Pritchard G (2014) The effects of stochastic market clearing on the cost of wind integration: A case of New Zealand electricity market. *Energy Systems* 5(4):657–675.

Koenker R (2005) *Quantile Regression* (Cambridge University Press, Cambridge, UK).

Morales JM, Conejo AJ, Perez-Ruiz J (2009) Economic valuation of reserves in power systems with high penetration of wind power. *IEEE Trans. Power Systems* 24(2):900–910.

Morales JM, Zugnoa M, Pineda S, Pinson P (2014) Electricity market clearing with improved scheduling of stochastic production. *Eur. J. Oper. Res.* 235(3):765–774.

New Zealand Electricity Authority (2017) vSPD. Accessed February 17, 2017, http://www.emi.ea.govt.nz/Tools/vSPD.

Philpott A, Pritchard G (2004) Financial transmission rights in convex pool markets. *Oper. Res. Lett.* 32(2):109–113.

Philpott A, Ferris M, Wets R (2016) Equilibrium, uncertainty, and risk in hydro-thermal electricity systems. *Math. Programming* 157(2): 483–513.

Pritchard G (2011) Short-term variations in wind power: Some quantile-type models for probabilistic forecasting. *Wind Energy* 14(2): 255–269.

Pritchard G, Zakeri G, Philpott A (2010) A single settlement energy only electric market for unpredictable and intermittent participants. *Oper. Res.* 58(4, pt 2):1210–1219.

Ralph D, Smeers Y (2015) Risk trading and endogenous probabilities in investment equilibria. *SIAM J. Optim.* 25(4):2589–2611.

Rud L (2009) Essays on electricity markets. Technical Report, SNF Report No. 10/09 (Institute for Research in Economics and Business Administration, Bergen, Norway).

Sen S (1993) Subgradient decomposition and differentiability of the recourse function of a two-stage stochastic linear program. *Oper. Res. Lett.* 13(3):143–148.

Shapiro A, Dentcheva D, Ruszczynski A (2009) *Lectures on Stochastic Programming: Modeling and Theory* (SIAM, Philadelphia).

Stoft S (2002) *Power System Economics: Designing Markets for Electricity* (Wiley–IEEE Press, Hoboken, NJ).

Whittle P (1971) *Optimization Under Constraints* (Wiley, Hoboken, NJ).

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

Zavala VM, Kim K, Anitescu M, Birge J (2017) A stochastic electricity market clearing formulation with consistent pricing properties. *Oper. Res.* 65(3):557–576.