Energy Security Improvements

ISO Discussion Paper

April 2019 – Version 1

1. Introduction and Summary

The electric power system in New England is undergoing a major transition. The owners of traditional power plants – nuclear, coal, and oil-fired – are permanently shuttering many of these stations due to economic and environmental pressures. The majority of the region's electricity, both currently and for the foreseeable future, is likely to come from newer, more efficient natural-gas fired generation and an array of renewable energy technologies, such as solar- and wind-powered generation.

This evolution comports well with the New England states' goals for a cleaner, greener regional power grid, yet it also presents new challenges. Both renewable and natural gas-based generation technologies rely on the "just-in-time" delivery of their energy sources. Solar- and wind-based power inherently vary with the weather. Less obviously, and of greater concern presently, is the just-in-time delivery of natural gas across interstate pipelines to the region's generating stations. During cold winter conditions, these pipelines rapidly reach capacity and are unable to fuel many of New England's power plants.

ISO New England (ISO) is concerned, given the power system's evolving resource mix and the region's constrained fuel delivery infrastructure, that there may be insufficient energy available to the New England power system to satisfy electricity demand during cold winter conditions. While there has been no loss of load attributable to insufficient energy supplies to date, we expect that industry trends will increase this risk over time unless proactive solutions are developed.

Giving heightened priority to this issue, in 2018, the Federal Energy Regulatory Commission (FERC) directed the ISO to submit "Tariff revisions reflecting improvements to its market design to better address regional fuel security concerns."[1] That directive arose amidst a contentious regulatory

[1] *ISO New England Inc.*, 164 FERC ¶ 61,003 at PP 2, 5 (2018).

ISO New England Inc.
One Sullivan Road
Holyoke, MA 01040-2841

process involving shorter-term, non-market actions to bolster the region's fuel supplies by delaying the retirement of the large Mystic Generating Station near Boston, Massachusetts. Expressing a clear preference for a different path forward, the FERC reaffirmed its "support for market solutions as the most efficient means to provide reliable electric service to New England consumers at just and reasonable rates," and expectations for the ISO "to develop longer-term market solutions."[2]

To that end, this paper shares the ISO's current perspective on underlying problems, root causes, and longer-term market solutions. While the power system's growing reliance on technologies with just-in-time energy sources poses new challenges, we believe these challenges have sensible solutions. Further, we readily agree with the FERC's affirmation that these challenges are most appropriately addressed through market mechanisms. As the technologies comprising New England's power grid continue to rapidly evolve, harnessing the forces of competition will provide the most cost-effective long-term solutions.

In practice, reducing the risks that arise in a power system increasingly reliant on just-in-time energy sources requires additional sources of energy supply (or reductions in demand) when gas pipelines are most constrained, renewable resources experience adverse weather, or both. At present, supplemental arrangements for fuel can enable existing fossil-fired generating stations to perform reliably during such conditions. Examples include arrangements by natural gas-fired generators to procure and maintain liquefied natural gas (LNG) inventories at existing LNG facilities in the Northeast (for use when the interstate pipelines are constrained during winter), and making advance transportation arrangements to enable backup fuel oil supplies to be replenished during winter at the region's dual-fuel (oil and gas) and oil-based power plants. Over the longer term, a broader array of capital investments may ultimately prove cost-effective alternatives, such as local "satellite" LNG storage facilities near generation stations, greater price-sensitive demand participation in the wholesale markets, and innovative electricity storage technologies (like grid-scale batteries) that can smooth out the intermittency of renewable energy resources.

Ultimately, the competitive power sector's willingness to undertake any of these reliability-enhancing-but-costly endeavors depends on their expected return on the investment. To date, the region's competitive power sector has made limited progress providing them. In recent winters, few natural-gas fired generators made advance arrangements for LNG inventories in New England; and by some measures, the generation fleet's fuel oil inventories for winter power generation are steadily declining with time due to both economic factors and emissions restrictions.

1.1 Problems and Causes

To facilitate productive discussions of the ISO's concerns and potential solutions, this paper begins with a deeper examination of the underlying problems and their root causes. Our focus is whether the ISO-administered wholesale electricity markets – which were not originally designed for the challenges just-in-time generation technologies have wrought – provide adequate financial

---

[2] *Id.* at PP 53, 54.

incentives for resource owners to make additional investments in supply arrangements that would be cost-effective and benefit the power system at times of heightened risk.

Our central conclusion is that, in many situations, the answer is no. Even when such supplemental supply arrangements would be cost-effective from society's standpoint as a means to reduce reliability risks, the current suite of market products do not provide sufficient financial incentives for market participants to undertake them. The root cause is logical enough. Making these discrete investments, if they meaningfully reduce the risk of electricity supply shortages (and therefore the risk of high prices), entails up-front costs to the generator – yet reduce the energy market price the generator receives. The value that society places on making the supplemental supply arrangement is based on the high price it *avoids* with the investment. However, the value the generator places on the same arrangement is based on the lower price it *receives* in the energy market with the investment. This value difference, in turn, results in a divergence between the social and private benefit of the investment – a situation we call a *misaligned incentives* problem.

In effect, given how New England's power system has evolved, generation owners now face an economic "catch-22": If a generator does not make, for example, a costly supplemental fuel supply arrangement, then when the region's gas pipelines are tightly constrained and renewables' output is low, high real-time wholesale energy market prices will prevail – prices that cost consumers dearly, but from which the generator does not immediately profit (if it lacks fuel to operate). Those price signals normally motivate widespread investment to profit in such circumstances. And yet, if the generator *does* invest in a supplemental fuel supply arrangement – at least, to a level that meaningfully reduces the system's energy supply risk – then the investment may obviate the market's high energy price, undermining the expected return on the investment. Given that nearly any investment in supplemental supply arrangements tends to entail significant sunk costs up-front, it is no surprise that few generation owners perceive adequate incentives to undertake them.

To explore this problem in detail, this paper provides a series of numerical examples. These are intended to help make the nature of the problem, and the conditions on which is rests, readily apparent. The bottom line is that investing in more robust energy supply (*e.g.,* fuel) arrangements may often be beneficial and cost-effective for the system, but not financially viable for individual generators in a today's energy market construct.

That analysis also reveals a collateral problem. If generation owners undertake more robust energy supply arrangements, it commonly yields a greater regional supply of energy 'in storage' (in the form of LNG at regional facilities, additional fuel oil at power plants, or other forms of energy storage for power generation). However, the current wholesale energy market format may deplete that limited pool of stored energy prematurely. At issue here is that New England, like other organized wholesale electricity market regions in the US, employs a single (that is, one day in advance) day-ahead wholesale electric energy market. That one-day-ahead market – which we might better term a *one-day-at-a-time* market – is inherently myopic: it does not evaluate outcomes, nor produce prices, for days beyond tomorrow. That may not efficiently coordinate the use of limited stored energy to meet both the demand for power tomorrow *and* the demand for power after tomorrow.

At its root, the problem is one of information coordination through markets. Each generation owner knows *its* limited energy inventory (*e.g.*, the fuel in its tanks), but not the state of energy supply limitations across the system. And a single day-ahead market that produces no prices for days *after* tomorrow cannot provide a generation owner with a market signal of the value of preserving its limited energy inventory for use later (*viz*., the day after tomorrow, or the day after that, and so on). The ISO does review considerable information on the energy supply limitations of generators for coming days (due to resource-level fuel surveys and gas-electric information sharing with pipeline operators and certain LNG facilities). But there is presently no market mechanism to aggregate that information into the market clearing process, or to coordinate generators' multi-day production schedules, in the least-cost, most-reliable way.

Cognizant of that shortcoming, the ISO has made strides in recent winters to encourage generators with limited stored energy to incorporate their intertemporal opportunity costs into their day-ahead market energy offer prices.[3] Here again there is a problem, however. A generator that delays using its limited stored fuel, in the hope of receiving a greater price using that inventory at a later date, earns no revenue for doing so today. Yet that incentive commonly arises at the same time for many similarly-situated generators in New England. Though acting individually, if many delay their use of their limited fuel stocks to the date with the highest *expected* future price (via opportunity-cost-based offers in today's day-ahead market), their combined effort will result in a *low* price on that very future date.

In simple terms, when generators delay their use of limited fuel stocks to a more critical period several days later, that action generally reduces the total cost of operating the power system and its reliability risk. So far, so good. But in return, the existing one-day-ahead market design may provide little compensation for those salutary actions. As such, it should come as no surprise that generation owners have not undertaken this efficiency-enhancing activity frequently.

Deconstructing these problems in detail, as we do in this paper, has a useful summary implication – the suite of products in the ISO-administered energy markets is incomplete. Their current form and associated ancillary service products were designed more than fifteen years ago, well before just-in-time energy powered the majority of New England's generation. In that earlier era, capacity supply was a constraining reliability concern. Specifically, as long as the system had sufficient operable capacity committed each day, another increment of energy demand could be satisfied by dispatching up the next generator. In today's environment, however, we do not face a capacity shortfall problem (indeed, the system is awash in capacity). We, instead, face an *energy security* problem due to the constraints – and uncertainties – on energy for power production.

1.2 Solutions

The second portion of this paper proposes a set of market design improvements to address these problems. To facilitate stakeholder discussions, the overall design is based on a familiar set of energy and ancillary service concepts. Broadly, we recommend expanding the existing suite of energy and ancillary service products in the ISO-administered markets, in order to address – reliably

---

[3] ISO New England Inc., *Energy Market Opportunity Costs for Oil and Dual-Fuel Resources with Intertemporal Production Limitations – Revised Edition* (October 9, 2018), at https://www.iso-ne.com/static-assets/documents/2018/10/a7_memo_re_energy_market_opp_costs_for_oil_and_dual_fuel_revised_edition.pdf.

and cost-effectively – the uncertainties and supply limitations inherent to a power system evermore reliant on just-in-time energy technologies.

Based on the analysis of problems and root causes, our efforts are directed at three meaningful measures: (a) strengthening generation owners' financial incentives to undertake more robust supply arrangements, when cost-effective, while not proscribing what form those supply arrangements may take; (b) rewarding resources' flexibility that helps manage, and prepare for, energy supply uncertainties during the operating day, given the increasingly just-in-time nature of the power system; and (c) efficiently allocating electricity production across multiple days from resources that have limited stored (*non*-just-in-time) energy sources. These measures serve as milestones on our path toward long-term market solutions.

Building upon the region's competitive wholesale electricity structure, we envision three core components to help achieve those milestones. These are:

- **Multi-day ahead market.** Expand the current one-day-ahead market into a multi-day ahead market, optimizing energy (including stored fuel energy) over a multi-day timeframe and producing multi-day clearing prices for market participants' energy obligations.

- **New ancillary services in the day-ahead market.** Create several new, voluntary ancillary services in the day-ahead market that provide, and compensate for, the flexibility of energy 'on demand' to manage uncertainties each operating day.

- **Seasonal forward market.** Conduct a voluntary, competitive forward auction that provides asset owners with both the incentive, and necessary compensation, to invest in supplemental supply arrangements for the coming winter.

Taken together, these three interrelated elements will help signal, through transparent market prices, the costs of operating a reliable power system as the profile of resources comprising the New England fleet continues to evolve.

Multi-day Ahead Markets. The first of these three core components is a voluntary market for forward energy transactions extending several days in advance of the delivery (operating) day. It effectively extends the existing single day-ahead market to operate on a rolling, multi-day horizon, with the multi-day market successively re-cleared each day. That enables suppliers to refine their energy positions, consistent with any changes to their fuel supplies, costs, and capabilities, prior to the delivery day. It also enables wholesale buyers to potentially acquire energy for their customers at less volatile forward prices than are likely to occur in one-day-ahead (and in real-time) transactions in a power system reliant on the just-in-time energy sources.

The rationales for a multi-day ahead market are several. This multi-day market design provides new, resource-level forward energy schedule visibility to market participants and the ISO alike. It also provides forward price signals to replenish inventories (or to preserve current inventories to a later date) when supplies are tight. Importantly, by optimizing the region's limited energy supplies cost-effectively over a multi-day horizon, it provides an economically sound method to reduce (or eliminate) the need for costly, out-of-market posturing of resources to preserve generators' scarce fuel supplies when the region experiences extended cold weather conditions.

Viewed from another perspective, a multi-day ahead market is a more comprehensive means to advance the goals of several recent initiatives. The ISO's 21-day winter energy supply forecasts provide participants with forward estimates of potential physical supply constraints, and the ISO's intertemporal opportunity cost revisions help encourage suppliers to preserve scarce fuel supplies for a later date. However, both have limitations relative to a multi-day ahead market. The ISO's physical energy supply forecasts provide participants with no price signal of the *value* – in hard dollar-per-MWh terms – of preserving (or acquiring additional) energy supply. Moreover, the existing alternative of incorporating anticipated opportunity costs into supply offers may not remunerate a supplier for valuable fuel preservation efforts at all. That problem, discussed previously, arises when the supplier has no ability to lock-in revenue today for preserving its fuel to future days. By contrast, a true multi-day ahead market can optimize the entire system's limited energy supplies over time and directly compensate each supplier – at competitively-determined prices – for their roles in doing so.

New Ancillary Services in the Day-Ahead Markets. The second core component is a set of new ancillary service products procured in the day-ahead markets. These services are intended to help ensure that the system is prepared for, and has the capabilities to manage, a range of uncertainties in a power system increasingly reliant on just-in-time technologies.

For context, most resources in New England successfully operate during the hours for which they receive an energy supply award in the day-ahead market. Consider, however, the situation when a large day-ahead cleared resource is subsequently unable to operate for an extended (multi-hour or multi-day) duration. This creates an unanticipated 'energy gap' in the day's operating plan. The replacement energy to fill that gap must come from other resources operating above, or that did not receive, a day-ahead award. With the region's growing dependence on just-in-time energy sources and its constrained fuel delivery infrastructure, however, those replacement resources – which did not expect to run – may not be able to operate unless they made costly supplemental fuel supply arrangements *a priori.* Yet in today's market construct, it is generally unprofitable to incur the costs of procuring fuel to cover days for which a resource does not expect to operate or to be paid.

In practice, the ISO relies upon much of the generation fleet's capabilities, above and beyond their day-ahead energy awards, to fill such energy gaps. In concrete terms, these capabilities fall in three operational categories:

- Operating reserves for fast-start/fast-ramping generation contingency response, which enable the system to promptly restore power balance (consistent with the timeframes established in applicable reliability standards);

- 'Replacement' energy needed to fill the bulk of an energy gap through the balance of the day and, if necessary, the next operating day as well; and

- Resources to supply energy that covers the gap when forecast electricity demand exceeds the demand cleared in the day-ahead market.

As discussed in detail in later sections, we distinguish these three categories insofar as they address potential energy gaps that arise on, and persist for, different timeframes.

At present, the ISO does not procure or compensate for these types of ancillary service capabilities on a day-ahead timeframe. That may have been reasonable in the past, when generators without day-ahead energy schedules characteristically had large, ready stockpiles with which to fuel an unexpected, extended run whenever an energy gap arose. But, as noted previously, those generators are retiring and many that remain are at risk of future retirement. Thus, it is important to improve today's energy market construct so that the future resource mix will invest in energy supply arrangements and technologies that ensure these ancillary service capabilities remain available to the power system each operating day.

To achieve that objective, we propose to formalize the foregoing three categories of operational needs into specific ancillary service capabilities, and allow resources to compete to provide those capabilities in the day-ahead markets. Offers to provide those ancillary services will be voluntary, and awards compensated at uniform, transparent, product-specific market prices. At a high level, a day-ahead seller of those ancillary services is providing the ISO with a "call option" on its resource's energy during the operating day, with different time-related parameters applicable to the different products.

To procure these services cost-effectively, we recommend that the award of these ancillary services be co-optimized (*i.e.*, simultaneously cleared) with all participants' energy supply and demand awards in the multi-day ahead market. That co-optimization process ensures, by design, that the clearing prices for energy and each ancillary service incorporate the (marginal) suppliers' opportunity costs of not receiving an award for a different day-ahead product. It also means that the day-ahead LMPs for energy will commonly incorporate the clearing prices for the ancillary services as well.

An important feature of these new ancillary services' designs is their settlement. Consistent with their value as a call option on energy during the operating day, we propose to settle a day-ahead ancillary obligation like a call option on real-time energy. That is a familiar multi-settlement rule used in a wide variety of markets, and will function well in concert with the existing day-ahead energy market's two-settlement design. However, the second (real-time) settlement would be slightly different for day-ahead energy and for ancillary service positions, naturally reflecting that the former is selling (or purchasing) forward energy and the latter is selling a call option on energy.

Importantly, an option settlement design creates new financial incentives for sellers of these ancillary services to ensure they have the physical wherewithal (including fuel) to cover their obligations the next day. This is because a resource that commits to providing an ancillary service will face a steep financial consequence in real-time settlement if the real-time energy price is high and the resource does not perform. Using a series of numerical examples, we explain how this provides stronger incentives than the existing market design for resources to incur the costs of supplemental fuel supply arrangements. At the same time, resources will receive day-ahead compensation to cover the fixed costs of such supplemental fuel supply arrangements, even if it turns out that they are not needed to operate the next day.

These product design and settlement features fundamentally change the incentives present today. From a commercial standpoint, it becomes profitable for the resources that the ISO relies on for these ancillary services to incur the costs of maintaining reliable fuel arrangements, when such arrangements are cost effective from the standpoint of the system overall – helping ensure they *could* perform if asked to fill an energy gap, even on a day they did not expect to operate.

Taken together, these two core components of the ISO's proposal – the multi-day ahead market and its co-optimized ancillary services – have a number of notable properties and benefits. We highlight four of these notable properties here:

- **Clear market incentives.** Incents and rewards, through a market-based mechanism, more efficient investments in energy supply arrangements than today's markets – investments that seek to reduce reliability risks in New England's increasingly energy-constrained power system;

- **Mirrors existing market structures.** This design extends the current concepts that underlie the ISO's longstanding energy and reserve markets, providing a transparent, conceptually familiar framework ;

- **Creates level playing fields and fosters innovation.** Enables all technologies capable of providing energy or any of the new 'energy on call' ancillary services to be compensated for those services;

- **A risk-responsive design.** If the region's energy security risks do not emerge in future years – perhaps because they are meaningfully reduced through different policies outside the ISO-administered markets (say, through much greater renewable energy production in future years during the winter) – then these new products and services would become low-cost to procure. That 'risk-responsive' aspect of the overall design prevents locking-in consumers to new multi-year obligations that might prove both expensive and unnecessary as New England's power system continues to evolve.

Seasonal Forward Market. The third core component is a voluntary forward auction intended to facilitate similar investments in costly supplemental energy supply arrangements in advance of each season (*i.e.*, many months prior to the operating period).

As with well-designed forward markets generally, we view their value as arising from two distinct forward market roles: first, information, in the form of a price signal for investment in energy supply arrangements well in advance of each the delivery period; and second, risk-reduction, relative to the revenue streams provided by a multi-day ahead market alone. The second of these roles may reduce financial risk (*i.e.*, revenue and expenditure uncertainty) for both buyers and sellers in the ISO-administered markets, relative to (potentially volatile) day-ahead ancillary service prices over the course of the winter.

The ISO remains in the early, conceptual stages of evaluating designs for such a forward market. One means of doing so may be to substantially re-vamp the existing Forward Reserve Market so that, in effect, it becomes a forward market for the same suite of new ancillary services discussed above and that are subsequently transacted in the day-ahead market.

As a general rule, well-designed forward markets require a corresponding spot market for the same good or service. That close linkage is necessary in order to apply standard two-settlement logic to a forward market – a crucial standard in market design that recognizes that what participants ultimately deliver may deviate (for any number of reasons) from their forward obligations. That design standard, already embodied in the ISO's energy market, also avoids "money for nothing" problems that can be expected to arise whenever a forward market does not follow a two-settlement design.

This paper does not address the design of such a forward market in any detail. Our immediate focus is to first work with regional stakeholders to develop the foregoing multi-day ahead markets and their integrated new ancillary-services. With these more fully developed, the ISO and stakeholders will be better positioned to understand their interactions with any forward market construct for the same or similar services. We anticipate developing the key elements – and, ultimately, design details – for a forward market after further discussion of its potential scope, and precise objectives, with stakeholders.

On a related note, and with regard to the ISO's proposed market design improvements as a whole, we emphasize that many elements are under ongoing development – and thus this proposal remains a work in progress. Importantly, we anticipate revising this paper and the ISO's proposal in response to stakeholder input and feedback throughout the course of the year. The latest version of this paper will be available on the Energy Security Improvements Key Project page on the ISO New England website, at https://www.iso-ne.com/committees/key-projects/energy-security-improvements.

The balance of this paper provides further perspective on problems and causes, the specific goals of this proposal, and explains in greater detail how these proposed energy market improvements would work. In Section 2, we examine the problems and their root causes in detail, and illustrate the challenges they present with a series of numerical examples. In Section 3, we summarize the design objectives and principles that guide the ISO's development of market improvements to address these challenges. In Section 4, we delve into recommended solutions, including concepts, their rationale, and additional numerical examples. Section 5 concludes with next steps.

We look forward to discussing these market enhancements with stakeholders.

2. Problems and Causes

In this section, we provide a deeper diagnosis of the problems and causes underlying the ISO's energy security concerns. To lend clarity to these concerns, we identify several specific adverse consequences under the status quo market design – consequences that may become more significant in the future, as the system transitions to a greater mix of resources with just-in-time energy sources. The analysis of these consequences, and their root causes, guides our development of recommended market-based solutions in Section 4.

2.1 Focusing Deeper: Three Specific Problems

Energy security is a broad term subject to a myriad of competing interpretations. To provide focus, we constructively frame the power system's emerging energy supply risks in terms of three specific problems, enumerated below. These have interrelated market and operational components, and adversely affect both the efficiency and reliability of New England's power system.

P1. **Incentives and Compensation.** Market participants whose resources face production uncertainty may have inefficiently low incentives to invest in additional energy supply arrangements, even though such arrangements would be cost-effective from society's standpoint as a means to reduce reliability risks.

P2. **Operational Uncertainties.** There may be insufficient energy available to the power system to withstand an unexpected, extended (multi-hour to multi-day) large generation or supply loss, particularly during cold weather conditions.

P3. **Inefficient Schedules.** The power system may experience premature (inefficient) depletion of energy inventories for electric generation, absent a mechanism to coordinate and reward efficient preservation of limited-energy supplies over multiple days.

The first of these problems is one of misaligned incentives. Investing in more robust energy supply (*e.g.,* fuel) arrangements may not be financially viable for individual generators in a today's market construct, yet can be beneficial and cost-effective for the system. This has both efficiency and potential reliability consequences. We address this problem in detail first.

The second of these problems relates to operational 'energy gap' situations, as introduced in Section 1. With the region's growing dependence on just-in-time energy sources and its constrained fuel delivery infrastructure, generating resources that do not expect to run the next day (*i.e.,* that do not receive an award in the day-ahead market) may not have sufficient energy to operate – unless they made costly supplemental fuel supply arrangements *a priori*. These concerns are heightened by the fact that the ISO relies upon much of the generation fleet's capabilities, above and beyond their day-ahead energy awards, to fill any energy gaps that arise during the operating day.

The third of these problems is a coordination problem. It arises because the existing day-ahead wholesale electric energy market does not schedule resources, nor produce prices, for days beyond tomorrow. That may not efficiently coordinate the use of generators' remaining energy inventories to meet both the demand for power tomorrow *and* the demand for power after tomorrow (and the day after that, and so on). As we note, various informational and economic functions available to market participants (such as opportunity-cost-based offers) can help with this coordination problem, but do not fully resolve it. Again, this has both efficiency and reliability adverse consequences.

It is important to note that these three specific problems are interrelated. In particular, problem P1 is a contributing factor to both problems P2 and P3, though both of the latter have other root causes as well. We discuss each of these three problems in succession below, in order to provide clear explanations of their distinct causes and consequences.

2.2 Problem P1: Misaligned Incentives for Energy Supply Arrangements

This section examines problem P1, focusing on existing market incentives. Specifically, we address why the ISO-administered wholesale electricity markets, in their current form, may not provide sufficient incentives for resource owners to make additional investments in energy supply arrangements – even when such investments would be cost-effective and reduce potential reliability risks.

At a high level, investing in a costly supplemental fuel arrangement that meaningfully reduces the risk of supply shortages (and therefore the risk of high prices) entails up-front costs to the generator, yet reduces the energy market price the generator receives. The value that society places on making the supplemental fuel arrangement is based on the high price it *avoids* with the investment. However, the value the generator places on the same arrangement is based on the lower price it *receives* in the energy market with the investment. This value difference, in turn, results in a divergence between the social and private benefit of the investment – a situation we call a *misaligned incentives* problem.

Misaligned incentive problems tend to occur in particular economic situations. As we illustrate below, these conditions are likely to prevail in the context of costly supplemental energy supply arrangements for competitive power generators. The misaligned incentive problem results in too little private investment in energy supply arrangements that would be cost-effective and desirable from society's standpoint under the existing market incentives.

To explain this problem and its root causes more precisely, a simple numerical example is helpful.

2.2.1 Example 1 – One generator

This example considers a single generator without a day-ahead market award. It faces an unlikely possibility that demand may be high enough for it to operate the next day, and must decide now whether or not to incur the cost of arranging fuel.

We simplify as much as possible here to focus on the essentials: A case where the fixed cost of arranging fuel is lower than society's benefit (*i.e.*, the system's expected cost savings) from it, but the fixed cost nonetheless exceeds the generator's expected profit. As a result, the competitive generator's rational decision is not to arrange fuel in advance of the operating day, even though society would be better off if it did.

► **Assumptions.** Consider a generator with 1 MW of capacity that faces uncertainty over whether or not it will operate the next day. The generator will be dispatched (if available) if demand is high, and not dispatched if demand is low. Assume there is a 20% chance of high demand, so the generator knows, most likely, it will not operate. To simplify this example, we will reduce the time period in which the generator may operate (or not) to a single future hour and assume that the generator does not clear (*i.e.,* is not scheduled) in the day-ahead market.

The generator's costs depend upon whether or not it arranges fuel in advance of the operating day. Arranging fuel entails both fixed and variable costs. We assume that if the generator arranges fuel in advance of the operating day then it incurs up-front fixed cost of $40. By 'up-front fixed cost', we mean that *if* the generator decides to arrange fuel in advance, it would incur the $40 cost *regardless* of whether or not it operates the next day. And then, *in addition*, it would incur a marginal cost of $70/MWh to operate – but that marginal cost is incurred only if it does indeed operate. We summarize these cost and demand assumptions in Table 2-1 below.

Table 2-1. Cost and Price Assumptions for Example 1

	With Advance Fuel			No Advance Fuel	
	High Demand	Low Demand		High Demand	Low Demand
Fixed Cost of Advance Fuel	$ 40	$ 40		$ -	$ -
Marginal Cost	$ 70	n/a		$ 400	n/a
Energy Price (LMP)	$ 120	$ 50		$ 400	$ 50
Demand Probability	20%	80%		20%	80%

If the generator arranges fuel in advance of the operating day and the demand is high, then it can produce at a marginal cost of $70/MWh and would be paid (in real-time) an LMP of $120/MWh. That's the 'good' scenario in this example. Importantly, we will assume that if the generator does <u>not</u> arrange fuel in advance, then it will not be able to acquire fuel the next day and will not be able to operate. In that scenario, if demand turns out to be high, the ISO would have to operate another, high-cost resource (at the margin) that would set a real-time LMP of $400/MWh. That's the 'bad' scenario.[4] Finally, assume if demand is low then the LMP would be $50/MWh so the generator would not operate, as it would be out-of-merit.

These assumptions are intended to capture the practical realities that there are fixed costs of acquiring energy supplies in advance of an operating day, *in addition* to the (marginal) cost of using the fuel itself. For example, consider the $40 fixed cost as the retainer (per MWh) for an intraday-

---

[4] Alternatively, one can interpret the $400/MWh cost as this generator's marginal cost if it must buy spot fuel intraday (on a really bad day), if it does not make advance fuel arrangements. Either interpretation will suffice for this example.

notice gas supply contract with an LNG terminal, and the $70/MWh cost as the incremental cost of calling for gas in order to run the next day if the generator is dispatched. Or, more broadly, consider the up-front cost as the generator's fixed cost to pre-arrange oil transportation service that would enable prompt replenishment of oil inventories, without which a generator would be out of fuel and not be able to run at all (for the relevant operating day). The point here is simply that there are fixed costs to arranging fuel in advance, and if a generator decides not to incur them, then (with *some* probability) the generator may not be have fuel to operate.

Last, a note on timing: When we say 'arrange fuel in advance' in this example, we mean however far in advance of the operating day as is necessary (a day, a week, a month, or a season). Though such timing issues matter in practice, in this simplified example *how far* in advance is not pertinent to the point of the example.

Rather, the fact that there are fixed costs to arranging fuel in advance of the operating day raises two key questions. First, would a competitive generator choose to incur them? Second, would its decision produce the best outcome for the system as well?

► **Society's preferred outcome.** First, let's examine what would be the most cost-effective outcome for the system. Arranging fuel in advance has a fixed cost of $40, and 80% of the time those arrangements will not turn out to be used or useful. That might suggest arranging fuel in advance is not worthwhile, from the standpoint of a cost-effective system.

But consider the benefits. Although high demand is unlikely, it occurs 20% of the time. When it does, arranging fuel in advance means incurring a marginal cost of $70/MWh and being able to avoid dispatching an expensive resource that costs $400/MWh. The expected value of the benefit to the system (*i.e.,* the *expected* cost saving) from avoiding that 'bad' scenario is:

$$20\% \times (\$400/\text{MWh} - \$70/\text{MWh}) = \$66/\text{MWh}.$$

On net, that means arranging fuel is indeed worthwhile: The cost of arranging fuel in advance of the operating day is $40 for the MWh, and the expected value of the benefit is $66/MWh, so the expected value of the net benefit is $66 – $40 = $26/MWh. When the decision to arrange fuel must be made in advance of the operating day, society would be better off (*i.e.,* there is positive expected cost saving, on net) if the generator invests the $40 – even though it may not be used.

Note that in coming to that conclusion, we have not introduced any reliability considerations. We simply have concluded that from the standpoint of minimizing the system's expected cost, it is efficient to incur the up-front cost of arranging fuel in advance even when it may not be needed.

► **The generator's decision.** From a commercial standpoint, it is prudent for the generator to incur the costs of arranging fuel in advance of the operating day only if its expected net revenue is greater doing so than if it does not. In this example, if the generator does not arrange fuel, its earnings are straightforward: it will not operate the next day and so will earn nothing. Let's now consider the alterative decision to arrange fuel, which, as illustrated below, entails an expected loss.

The cost of the 'investment' in fuel arrangements prior to the operating day is $40 up front. As before, 80% of the time demand will be low and those arrangements will not to be used or useful.

The other 20% of the time, demand is high and the generator is dispatched. In that high-demand scenario, the generator is paid the real-time LMP of $120/MWh and incurs a marginal cost of $70/MWh, earning a gross margin of $120 – $70 = $50/MWh.

That $50/MWh gross margin is more than enough to cover – on a high-demand day – the up-front cost of arranging fuel. However, the generator does not expect to operate every day. After all, demand is high only 20% of the time. That risk changes the generator's profit and loss calculus entirely.

Accounting for that uncertainty, the generator's expected net revenue is a loss. The up-front $40 investment in fuel arrangements (a cost for sure) has only a 20% chance of earning a gross margin with which to cover it, but the margin is not enough. The generator's expected profit, if it arranges fuel, is 20% × $50/MWh – $40/MWh = – $30, a net loss. In other words, arranging fuel in advance is not financially prudent for the generation owner.

Table 2-2. Generator's Expected Net Revenue for Example 1

			Advance Fuel		No Advance Fuel	
Generator's Market Settlement		Calculation	High Demand	Low Demand	High Demand	Low Demand
[1]	Day Ahead		$ -	$ -	$ -	$ -
[2]	Real Time	*RT LMP*	$ 120.00	$ -	$ -	$ -
[3]	Total Settlement	[1]+[2]	$ 120.00	$ -	$ -	$ -
Generator's Costs						
[4]	Advance Fuel	*F*	$ (40.00)	$ (40.00)	$ -	$ -
[5]	Marginal Cost	*MC*	$ (70.00)	$ -	$ -	$ -
[6]	Total Cost	[4]+[5]	$ (110.00)	$ (40.00)	$ -	$ -
Generator's Net Revenue						
[7]	Scenario Net Revenue	[3]+[6]	$ 10.00	$ (40.00)	$ -	$ -
[8]	Scenario Likelihood	*p* or (1-*p*)	20%	80%	20%	80%
[9]	**Expected Net Revenue**	SumProd [7]*[8]	**($30)**		**$0**	

Since there are many numbers to track in these calculations, and because we will extend this example later, Table 2-2 provides the ISO settlements and net revenue calculations for this generator for four situations: high and low demand with and without arrangements for fuel in advance of the operating day. The bottom right-hand cell shows that if the generator does not arrange fuel, its expected profit is zero (since it does not operate). The bottom left-hand cell shows that if the generator does arrange fuel, it is indeed in the red – it incurs an expected loss of $30.

The point of Example 1 is an important one. The market, in its current form, may not provide sufficient incentives for resource owners to invest proactively in energy supply arrangements – even when such investments would be cost-effective and yield expected net benefits to the system.

2.2.2 Reliability Risks and Problem P1

In the prior example, the adverse consequences of the generator's decision are that the outcomes are not cost-effective. There is a 'market failure' to incent efficient outcomes, causing higher expected costs to society as a result. However, that is not the only potential problem.

Let's now modify the prior example just slightly: Assume next that if the generator does not arrange fuel in advance of the operating day and demand is high, then the system will not have sufficient resources to avoid a shortage of reserves. In that 'bad' scenario, the real-time price for reserves (in shortage) would be $1,000/MWh and the LMP (we'll assume) would be higher, at, say $1,400/MWh. If the generator does arrange fuel, however, we'll assume (in that 'good' scenario) the same outcomes as before with no reserve shortage.

Before proceeding, we note that for simplicity, we first analyze this scenario based on the energy market's incentives and outcomes. There is also an impact (in this new scenario) to consider from capacity market performance incentives, which involves additional calculations, and that we will subsequently address further below.

From the generator's standpoint, nothing changes due to the now higher LMP that may occur if it does *not* arrange fuel. It was never paid the LMP that prevails in the 'bad' scenario (since it does not operate when that occurs) – and therefore that high LMP scenario does not incent it to invest in fuel supply arrangements. In terms of the numbers in Table 2-2 above, the value of the real-time LMP in the cases with 'No Advance Fuel' (the right-hand columns) produces no revenue for the generator. And, if it does arrange fuel in advance of the operating day (the left-hand columns), and by so doing prevents the reserve shortage, it would still have an expected loss of $30.

Things are *not* the same from society's standpoint in this new situation, however. From that perspective, the benefit of arranging fuel in advance is now much larger. Here, we will assume that the costs to society of the reserve shortage are the sum of the marginal alternative generator's cost (again assumed to be $400/MWh) and the 'cost' at which the market values (at the margin) a reserve shortage, which is presently $1,000/MWh.[5] That means arranging fuel in advance avoids incurring, if high demand occurs, a cost of $1,400/MWh and instead using $70/MWh energy (at the margin) to meet demand. That 'good' scenario has an expected benefit (expected cost saving) to the system of 20% × ($1,400/MWh – $70/MWh) = $266/MWh. On net, that means arranging fuel is indeed worthwhile for the system: The cost of arranging fuel in advance of the operating day is $40 for the MWh, and the expected value of the benefit of doing so is $266/MWh, for a net expected benefit of $266 – $40 = $226/MWh.

The point here is simple. If a generator's decision to arrange fuel in advance is material enough to impact – with *some* probability – whether or not the system experiences a reserve or energy shortage, then the divergence between society's and its private incentives *gets worse.* That is, the problem of the misaligned incentives does not have only adverse efficiency consequences. It can

---

[5] This $1,000/MWh value is an administrative shortage price, defined in the ISO Tariff as the Reserve Constraint Penalty Factor (RCPF), associated with the system's real-time minimum total operating reserve requirement. *See* ISO Tariff § III.2.7A.

also have adverse reliability consequences. As this case shows, the competitive generator's rational decision is again not to arrange fuel, but society would be even better off – and the system's reliability risk lower – if it did. The more severe the consequences of the generator's decision, the more misaligned are its incentives.

We should note again here that this case with the potential reserve shortage does not incorporate other market incentives that are important in New England. Since June 1, 2018, resources that supply energy face stronger marginal incentives to perform under the ISO's Pay for Performance (PFP) market rules. We address how PFP affects these observations in greater detail, further extending the example, below.

2.2.3 Insights: Problem P1's Consequences

Although these examples are meant to be simple illustrations of a market design problem, they identify several key points that hold generally. First, investments in energy supply arrangements are often characterized as insurance, in the sense of paying more to achieve more reliable outcomes. That's logical enough, as the second example suggests. However, that isn't the whole story.

In these examples, investments in energy supply arrangements *lower* the system's expected total cost – paying *less* overall – to achieve equally (or more) reliable outcomes. That's a far more sweeping observation. It says the system would meet demand more cost-effectively overall if the generator made the up-front investment to arrange fuel, even though it may not be used. However, under the current market design (though very simplified in this example), making such fuel supply arrangements is not financially prudent from the generator's standpoint of maximizing its expected net revenue. And the generator is acting perfectly rationally and competitively (offering at its marginal cost) throughout.

What is the crucial insight here? The key is simply that the market price for energy – *i.e.,* what consumers value consuming – is impacted by the supplier's investment in fuel arrangements (at least, with positive probability). Thus, in taking a costly action (incurring the fixed cost of arranging fuel), society benefits more than the generator does. In the first example, when demand is high, the generator realizes only $50 (its gross margin) of the $330 in cost savings that its investment helps consumers achieve. The difference in those benefits (to the generators) and cost savings (to society) results in the misaligned incentive problem, and higher expected costs to society as a result.

Another insight to note is that the cost society should incur to motivate greater energy supply arrangements by generators is not unbounded, and depends on the likelihood those energy supplies would be used. In Example 1, for instance, suppose instead that the chance of high demand – when the advance arrangements are useful – is lower by half (say), at only 10 percent. In that case, the up-front cost remains $40 but the expected benefit (expected cost savings) to society is lower than before: 10% × ($400/MWh – $70/MWh) = $33/MWh. The expected costs now exceed the expected benefits. In this case, incurring the fixed costs of advance fuel arrangements is not worthwhile for either the generator, or for society as a whole.

The point here is straightforward. The social benefits of arranging advance fuel depend – crucially – on the magnitude of the costs society *avoids* by doing so (the $400/MWh, in this example), and on the likelihood it would matter. That should be logical enough. We should be willing to spend more up-front to mitigate a risk when its likelihood is greater, and when its consequences (if it does occur) are more severe.

However, those conditions are precisely when the misalignment problem is at its worst. Since the generator does not internalize (that is, is not the beneficiary of) the high price that is avoided by its investment in fuel arrangements, the divergence between its private incentives and society's preferred outcome is greatest precisely when society values such advance fuel arrangements the most.

The implication here is that the misalignment problem will not solve itself. When the underlying risk likelihood is growing over time, there will be an ever-expanding range of up-front costs of advance fuel arrangements for which generators will not find it profitable to undertake, but for which society will be better off if they did.

2.2.4 Why Doesn't Pay-for-Performance Solve This?

Resources in the New England power system are incented, in real-time, not only by the real-time LMP, but also by a performance incentive created by the ISO's PFP market rules. The PFP market rules will impact a resource's incentives to arrange fuel in advance of the operating day whenever the conditions demarking a PFP event may occur (known as Capacity Scarcity Conditions).

As an initial observation, note that in the initial analysis of Example 1, there is no reserve deficiency, and the PFP market rules would not change any of the calculations in Table 2-2. Yet that example shows how the energy market may not provide sufficient incentives for resource owners to invest proactively in energy supply arrangements – even when such investments would be cost-effective and yield expected net benefits to the system. This market failure to produce efficient, socially beneficial investment decisions would not be altered by PFP, as it stands wholly apart from the circumstances when PFP would apply.

In more extreme situations where there are shortages of energy or reserves (with positive probability), the impact of PFP on these incentives is more nuanced. In sum, even when there may be a reserve shortage, the short answer to the prompting question is that PFP helps, but it does not fully solve, problem P1. To illustrate why both parts of that answer are true, we extend the prior numerical examples next.

► **Example 1 with PFP: Additional assumptions.** We'll use the same assumptions as in the Reliability Risks extension of Example 1 discussed previously, where there is a potential reserve shortage, and now layer-in the additional settlements associated with PFP. To capture the impact of PFP, assume the generator has a Capacity Supply Obligation of 1 MW (its capacity). During the hour considered in this analysis, assume the system's balancing ratio (BR) is 80% (that exact value is not critical to what follows), and that the performance payment rate (PPR) is equal to its current Tariff value of $3,500/MWh.

We again assume that if the generator does not arrange fuel in advance of the operating day and demand is high, the system will not have sufficient resources to avoid a shortage of reserves. In that 'bad' scenario, as before, the LMP would be \$1,400/MWh. If the generator does arrange fuel, however, then we have the same outcomes (in that 'good' scenario) as before with no reserve shortage. All other assumptions are the same as those summarized in Table 2-1, above.

► **The generator's decision.** In this example, if the generator does arrange fuel in advance of the operating day, nothing changes from the outcomes summarized in the left-side columns of Table 2-2 previously. There is no reserve shortage: if demand is low the generator does not operate, and if demand is high it operates and earns a gross margin of \$120 – \$70 = \$50/MWh. As before, the expected value of its gross margin is 20% × \$50/MWh = \$10/MWh, which is not enough to cover its \$40 up-front fixed cost. Thus, as shown in the bottom row of Table 2-2, if the generator does arrange fuel in advance, it expects to incur a net loss of \$30.

Let's now consider the alterative decision to not arrange fuel in advance of the operating day. In this case, the generator no longer has an expected profit of zero when it does not run. Instead, it will incur a non-performance charge in PFP settlements.

Table 2-3 below summarizes the relevant calculations. Here's the explanation. The general PFP settlement formula (in simple terms) is

$$\textit{Performance Payment} = \text{PPR} \times (\text{A} - \text{BR}) \times \text{CSO} \times \textit{event duration}$$

where A is the resource's output (in MWh). In this example, A is zero if the generator does not make arrangements for fuel, its CSO is 1 MW, and the event duration is assumed to be one hour. Therefore, the performance payment would be a charge of:

$$\$3{,}500/\text{MWh} \times (0 - 80\%) \text{ x } 1 \text{ MW CSO} \times 1 \text{ hour} = -\$2{,}800.$$

This is shown in row [4] of Table 2-3, for the scenario (column) with high demand and no advance fuel.

Of course, whether or not that occurs depends if demand is high or not. As before, if demand is low, there is no reserve shortage, the generator is not called to operate, and its net revenue in the low-demand scenario is zero. However, there's a 20% chance of high demand and, without the fuel to operate, it would then incur the PFP non-performance charge of \$2,800. The expected value of the generator's net revenue if it does not make arrangements for fuel in advance of the operating day is therefore

$$20\% \times (-\$2{,}800) + 80\% \times \$0 = -\$560.$$

See the bottom row of Table 2-3. Viewed this way, the profit maximizing decision – which, in this case, is a *loss minimizing* decision – is laid bare. Arranging fuel involves an expected net loss of \$30, but not arranging fuel involves an expected net loss of \$560. Given these stark alternatives, the generator's prudent course of action is to incur the up-front cost of arranging fuel.

Table 2-3. Generator's Expected Net Revenue for Example 1 with PFP						
			Advance Fuel		No Advance Fuel	
Generator's Market Settlement		Calculation	High Demand	Low Demand	High Demand	Low Demand
[1]	Day Ahead		$ -	$ -	$ -	$ -
[2]	Real Time	*RT LMP*	$ 120	$ -	$ -	$ -
[3]	PFP Performance Pmt	*PPR * (A - BR)*	$ -	$ -	$ (2,800)	$ -
[4]	Total Settlement	[1]+[2]+[3]	$ 120	$ -	$ (2,800)	$ -
Generator's Costs						
[5]	Advance Fuel	*F*	$ (40)	$ (40)	$ -	$ -
[6]	Marginal Cost	*MC*	$ (70)	$ -	$ -	$ -
[7]	Total Cost	[5]+[6]	$ (110)	$ (40)	$ -	$ -
Generator's Net Revenue						
[8]	Scenario Net Revenue	[4]+[7]	$ 10	$ (40)	$ (2,800)	$ -
[9]	Demand Probability	*p* or (1-*p*)	20%	80%	20%	80%
[10]	**Expected Net Revenue**	SumProd [8]*[9]	**($30)**		**($560)**	

In that sense, PFP helps to solve problem P1, as suggested above. Society is better off (as before) if the generator arranges fuel in advance of the operating day, as it helps to avoid the high costs and reliability risks of a reserve shortage. And the generator is incented to do so, because of the high financial price to pay if it is unable to perform when a reserve shortage occurs.

► **So why doesn't PFP fully solve the problem?** There is much more to the PFP question we started with. As noted at the outset, PFP helps, but does not fully solve, problem P1. Let's now consider a minor change to the preceding scenarios that will reverse the foregoing result – and show how PFP does not fully solve the problem with (arguably) more "realistic" risks.

The preceding PFP example had a number of simplifying assumptions (many just to keep the calculations simple). One seemingly unrealistic assumption is that there would be a 20% chance of a reserve shortage (absent the fuel arrangements). Capacity Scarcity Conditions, in practice, are rare events. Let's now see what happens if we re-do the preceding calculations assuming that there is only a 1% chance of a reserve shortage (absent the fuel arrangements). That lower risk level will change things significantly, and show that PFP does not fully solve problem P1 – either in theory, or in many cases, in practice.

To expedite the narrative, Table 2-4, below, shows the full settlements and expected net revenue for the generator with PFP under the same assumptions as before, but with only a 1% chance of high demand. The bottom row corresponding to 'with fuel arrangements' now produces an expected loss of $39.50, which is close to the $40 up-front fixed cost of arranging fuel. That $40 is now a total loss 99% of the time, offset by a slim 1% chance that demand is high, the unit runs, and makes its $50 gross margin. The generator's net expected revenue, if it arranges for fuel in advance of the operating day, is thus 1% × $50/MWh – $40 = – $39.50/MWh, a net loss.

Table 2-4. Generator Expected Net Revenue for Example 1 with PFP						
			Advance Fuel		No Advance Fuel	
Generator's Market Settlement		Calculation	High Demand	Low Demand	High Demand	Low Demand
[1]	Day Ahead		$ -	$ -	$ -	$ -
[2]	Real Time	*RT LMP*	$ 120	$ -	$ -	$ -
[3]	PFP Performance Pmt	*PPR * (A - BR)*	$ -	$ -	$ (2,800)	$ -
[4]	Total Settlement	[1]+[2]+[3]	$ 120	$ -	$ (2,800)	$ -
Generator's Costs						
[5]	Advance Fuel	*F*	$ (40)	$ (40)	$ -	$ -
[6]	Marginal Cost	*MC*	$ (70)	$ -	$ -	$ -
[7]	Total Cost	[5]+[6]	$ (110)	$ (40)	$ -	$ -
Generator's Net Revenue						
[8]	Scenario Net Revenue	[4]+[7]	$ 10	$ (40)	$ (2,800)	$ -
[9]	Demand Probability	*p* or (1-*p*)	1%	99%	1%	99%
[10]	**Expected Net Revenue**	SumProd [8]*[9]	**($39.50)**		**($28.00)**	

What if the generator does not arrange fuel in advance of the operating day? As before, if demand is high there is a reserve shortage and the generator would incur the PFP performance charge of $2,800. However, that has only a 1% chance, so the expected value of the generator's PFP performance charge is now comparatively trivial: 1% × $2,800 = $28. See the bottom row of Table 2-4.

Comparing the two cases, the generator's prudent financial decision is again the loss-minimizing one. Arranging fuel involves an expected net loss of $39.50, but not arranging fuel – which now is very unlikely to be used – involves an expected net loss of $28. The generator is financially better off, given these alternatives, if it does *not* incur the up-front cost of arranging fuel. Yet, as before, society would be better off if it did: The system's outcomes would be more cost-effective, and the reliability risk of a reserve shortage would be reduced. The misalignment problem remains.

In that sense, PFP does not fully solve problem P1. Society faces a lower reliability risk (as before) if the generator arranges fuel in advance of the operating day, but those arrangements are not consistent with the generator's commercial interest.

The reason PFP does not fully solve problem P1 is not complicated. As this example illustrates, when the risks of a reserve shortage are low, the incentive PFP creates for performance are significantly muted – too muted for it to be privately financially beneficial for the generator to incur the up-front fixed costs of arranging fuel, the costs of which will be a total loss most of the time.

Of course, one might note – accurately – that if lack of fuel to operate was widespread in the generation fleet when those generators are 'in demand', then reserve shortages may no longer be such infrequent events. And indeed, as the preceding example shows (when there was a 20% risk of a reserve shortage), if the likelihood of a reserve shortage is higher than the impact of PFP is far more powerful – it then becomes a better decision to arrange fuel proactively to ensure the generator can operate if called. So, in a sense, this problem is 'self-correcting' because PFP would tend to induce resources to arrange fuel to ensure they can perform *if* the frequency of reserve shortages becomes high enough.

That, however, should be viewed as a Pyrrhic victory from reliability standpoint.

► **Root causes.** Example 1 shows that even when up-front investments in supply arrangements would be cost-effective from society's standpoint, the current market design may not provide sufficient financial incentives for competitive generation owners to undertake them. The example's simplicity belies this observation being unique to the numbers involved; instead, it should be viewed as broader problem with the current market construct.

Stated generally, problem P1 has three requisite root causes, all of which are at work in the mechanics of Example 1 (and variants thereof) discussed previously. These three root causes are:

RC1. **Uncertainty** over whether the generating unit will be in demand, or not;

RC2. **Fixed costs** of making arrangements for fuel, which must be incurred in advance of learning whether the generator will be in demand (asked to operate) or not; and

RC3. **Energy supply arrangements that matter**, in the precise sense that if the generator does not make arrangements for fuel in advance, then with some probability (*i.e.,* if the generating unit is in demand), the real-time price for energy will be higher, or reliability will be worse, than if it did.

The first two of these root causes, uncertainty (RC1) and fixed costs (RC2), are conditions commonly affecting fuel supply arrangements for much of the generation fleet – such as the oil-fired and higher-cost (higher heat-rate) gas-fired resources that do not often clear in the day-ahead energy market during the winter. As a result, their owners face uncertainty over whether those resources will be called to operate in real-time (and, more broadly, uncertainty over how often they may run over the course of the season).

The third root cause (RC3) merits note. This condition, substantively speaking, is a logical necessity for there to be an energy security problem. To see why, note that if – counterfactually – condition RC3 is false, then there must be *another* resource available at the same time, and at the same (or lower) offer price, that could serve the same increment of demand. Thus, if RC3 is false, there is neither a market efficiency problem nor a reliability problem. Stated in more intuitive terms, for problem P1 to arise, generators' energy supply arrangements must impact potential outcomes in the precise sense that without them there is a chance of a higher real-time price and higher likelihood of a shortage (of reserves or of energy).

► **Which root cause is different today than in the past?** A natural question to ask is: Which of these root causes has changed to make energy security now a larger concern? That is, has the problem illustrated by Example 1 always been a concern, or is this misaligned incentives problem more pressing today, relative to years past?

One way to think about the evolution of the system's resource mix is in terms of its impact on the third root cause, RC3. In the past, generators commonly had large, ready stockpiles with which to fuel a run whenever committed or dispatched. If some day-ahead scheduled resource wasn't able to operate unexpectedly (for any reason), there were always sufficient energy stocks to dispatch up another generator in its place. This was true provided the system was committed (or could be supplementally committed) to have sufficient capacity for the peak hour, a main operating focus in the past.

Now, with constrained fuel infrastructures, retirements of generators with ample fuel storage, and evermore just-in-time generation from renewable technologies, it is no longer assured that if a scheduled resource isn't able to operate unexpectedly (for any reason), there will always be sufficient energy to dispatch up another generator – at a similar or small change in incremental cost – in its place. Instead, if a generator has no fuel to operate during cold weather conditions, then there is an increasing likelihood that the LMP will be set by either (*a*) an expensive next resource 'in the stack', or (*b*) scarcity pricing that signals a deficiency in the system's supply of energy and reserves.

For these reasons, as New England's resource mix has evolved toward technologies with predominantly just-in-time energy sources, the third of these root causes (RC3) is a more significant potential concern than in the past. Moreover, we do not expect that root cause to abate in the future, given the generation fleet's dramatic shift to more and more just-in-time resources: natural gas-fired generating units served by increasingly constrained interstate pipelines, and the rapid growth renewable power technologies as the New England states work steadily to advance de-carbonization goals.

2.3 Problem P2: Uncertainties and Practicalities

We now turn to problem P2, concerning operational uncertainties. As noted at the outset of Section 2, with the region's growing dependence on just-in-time energy sources and its constrained fuel delivery infrastructure, generating resources that do not expect to run the next day (*e.g.,* that do not receive an award in the day-ahead market) may not have sufficient incentives to make costly energy supply arrangements *a priori*. This, in turn, precipitates the concern we identified as problem P2: there may be insufficient energy available to the power system to withstand an unexpected, extended (multi-hour to multi-day) large generation or supply loss, particularly during cold weather conditions.

Our analysis here builds on problem P1 and the insights from the prior section, but now brings to the discussion a 'reality check' regarding the power system's operational needs. These lend focus to the resources and capabilities that problem P1 may adversely impact the most in practice: those we rely upon to cover an 'energy gap' when the system's conditions during the operating day significantly differ from the ISO's day-before operating plan and the outcomes of the day-ahead market. We explain these capabilities, and how they are potentially impacted by problem P1, next.

2.3.1 Operational Requirements and Uncertainty

For context, consider again the energy gap concept introduced in Section 1. Specifically, suppose a large resource that received an energy supply award in the day-ahead market is subsequently unable to operate (for any reason). This creates an unanticipated energy gap in the day's operating plan. The replacement energy to fill that gap must come from other resources operating above, or that did not receive, a day-ahead award.

In practice, the ISO relies upon much of the generation fleet's capabilities, above and beyond their day-ahead energy awards, to fill such energy gaps. In concrete terms, these capabilities fall in three operational categories:

A. **Operating reserves for fast-start and fast-ramping generation contingency response.** Energy from resources that the ISO relies upon for real-time operating reserves. These include both off-line (fast-start) generation and the unloaded 'upper blocks' of on-line generation ('spinning' reserves). The ISO relies upon these capabilities to ensure the system is prepared to promptly restore power balance (consistent with the timeframes established in applicable reliability standards) in response to a sudden, unanticipated power supply loss during the operating day.

B. **Replacement energy.** Energy supply needed to replace a day-ahead cleared resource that is unexpectedly unable to operate for an extended (multi-hour to multi-day) duration. In that situation, the ISO must dispatch online resources above their day-ahead schedules, or supplementally commit offline resources without day-ahead schedules, to supply sufficient energy to cover the energy gap through the balance of the day (and, if applicable, the day following).

   Importantly, in preparing a reliable next-day operating plan, the ISO does not plan for the resources that meet the needs in category A – that is, those expected to provide operating reserves – to sustain their energy production all day after a contingency. A primary reason is that, under applicable reliability standards, operating reserves must be restored within proscribed time limits (*e.g.,* approximately 105 minutes for ten-minute operating reserves, generally). In addition, resources designated to provide operating reserves may be able to sustain their power production for as little as one hour.

   Taken together, these considerations mean that in preparing a reliable next-day operating plan under applicable standards, the ISO can rely on the resources in category A for energy for only a limited amount of time (*e.g.,* a few hours or less – a duration that may vary with resource availability, demand, and other operating day conditions). After that point, the replacement energy to cover a contingency's balance-of-day energy gap must come from the dispatch and commitment of other resources operating above, or that did not receive, a day-ahead award.

C. **Load-balance reserves.** Energy to supply the load-balance 'gap' when the total day-ahead energy supply awards are less than the ISO's load forecast, in one or more hours, during the next (operating) day. Under applicable reliability standards, ISO's operating plan for

> the next day is intended to ensure there is sufficient energy to cover the forecast load each hour – not simply the level of demand cleared in the day-ahead energy market. This is in addition to, and distinct from, ensuring that the system is prepared to handle supply loss contingencies addressed with operating reserves and replacement energy capability.
>
> When the day-ahead market's cleared demand (less any net virtual supply) is less than the ISO's forecast load for one or more hours the next day, the energy to cover that load-balance gap is supplied through the dispatch and post-market commitment of other resources operating above, or that did not receive, a day-ahead market award.

Our present point is that the ISO relies upon much of the generation fleet's capabilities – above and beyond their day-ahead market awards – to satisfy the next-day operating plan's requirements and maintain a reliable power system. The region, however, does not currently compensate resources in the day-ahead market for these capabilities. Instead, the ISO employs (unpriced) constraints in its day-ahead market unit commitment process to help ensure that there will be sufficient operating reserves each hour of the next day (category A); and it employs post-market procedures and optimization tools to evaluate and ensure there will be sufficient replacement energy and load-balancing capabilities the next day (categories B and C, respectively).

Given the region's growing dependence on just-in-time energy sources and its constrained fuel delivery infrastructure, the ISO is increasingly concerned that the resources the system relies upon for these capabilities may not have energy supply arrangements that will enable them to operate on days when they have no reason to expect to run (or to run above or longer than their day-ahead market award, if any). In that event, if the system experiences an unexpected, extended large generation or supply loss during cold weather conditions – particularly, if it occurs when renewable resources' production capability is low (when the sun is down or the winds are calm) – the region may not have the energy needed to reliably fill the ensuing energy gap.

► **Quantities.** The New England system has over 30 GW of capacity resources that supply power, and experiences net power demand of approximately 21 GW or so during cold weather conditions. For context, it is useful to clarify how much of that supply capability the ISO typically relies upon for the three purposes described above.

The short answer is that the total quantity of power and energy the ISO relies upon to satisfy these three operational purposes varies from day to day. The first category, operating reserves, have formulaic requirements. Total operating reserves for prompt contingency response are typically in the range of 2-to-2.5 GW, and are based on the projected size of the largest and next-largest source-loss contingencies each day. The specific amount required (for the peak hour of the operating day) is reported daily in the ISO's *Morning Report.*[6]

Replacement energy is more complex, as it depends on the scheduled energy profile of the system's largest contingency over the course of the day. This may vary if the largest contingency is, for example, an external interface with an hourly-varying import energy schedule for the next day, or it may be constant over the course of the day, if the largest contingency is a non-dispatchable

[6] At https://www.iso-ne.com/markets-operations/system-forecast-status/morning-report.

resource with constant scheduled power output over time (such as a nuclear unit). Importantly, replacement energy is an energy (GWh) not a power (GW) concept.

The load-balance gap varies from day to day. In recent years, day-ahead cleared demand, after subtracting net virtuals (*i.e.*, cleared virtual supply less virtual demand), is often within a few percent of the load forecast in most hours. However, that gap can amount to many hundreds of MWh (per hour) and occasionally over a GWh in some hours. When that load-balance gap is large, the ISO relies on resources' capabilities above their day-ahead awards to cover the load forecast, and may supplementally commit (after the day-ahead market) additional generation for this purpose. Load-balancing reserves are an hourly energy concept (*i.e.,* GWh in each hour).

The summary point here is that, on some days, the total capability that the ISO relies upon to satisfy the foregoing three operational purposes can be substantial – often several GWh (per hour) or more of generation capability. The total quantities required to provide a reliable next-day operating plan vary from day-to-day, and these quantities are objectively based on the forecast demand profile and the system's largest potential single-source energy losses during the course of the operating day. These are capabilities that are not remunerated in the day-ahead market today, however, as they are provided by resources' capabilities above and beyond the level of their day-ahead market awards.

► **Not a static set of resources.** Importantly, the most cost-effective set of resources to meet the operational needs in categories A, B, and C can (and does) vary daily. It depends on the day-ahead cleared generation pattern, the cleared and forecast demand profile over the course of the day, available resources' lead-times and capabilities, weather and intermittent-resource energy production (actual and forecast), constraints on natural-gas pipelines supplying electric generation, and so on.

As examples, the types of existing resources the system may rely upon to meet the foregoing three operational purposes includes:

a) Off-line fast-start dispatchable generators (generally, hydro-electric and distillate-fueled combustion turbines and internal-combustion units), which infrequently receive day-ahead energy market awards and are dispatched during the operating day as circumstances require;

b) Higher-cost 'blocks' of combined-cycle generators that receive day-ahead awards below their maximum output (or possibly for a lower-output configuration), which the ISO may dispatch higher or schedule longer than their day-ahead market awards for the operating day;

c) Higher heat-rate combined-cycle generators that do not clear in the day-ahead market and may be committed (after the day-ahead market or, if necessary, intra-day) to satisfy the load forecast or for replacement energy; and

d) Long lead-time oil-steam units, in certain situations (*e.g.* cold weather conditions) when these resources can be lower-cost than gas-fired alternatives or gas-pipeline constraints preclude gas-fired resources from serving the system's load-balancing and potential replacement energy needs.

Our immediate point here is to note that there aren't a static set of resources, or a specific set of technologies, that are most cost-effective in meeting the system's operational needs for operating reserves, replacement energy, and load-balancing capability. It varies from day to day. Moreover, the specific resources the ISO may rely upon for these purposes as part of a reliable next-day operating plan depend on the generation commitments and energy schedules awarded in the day-ahead energy market.

Looking forward, the ongoing evolution in New England's resource mix will also change the set of resources the system potentially relies upon for these same operational purposes. Many of the resources in category (d) the ISO considers at risk for retirement, which may subsequently leave the combined-cycle generators in categories (b) and (c) the predominant resource types to satisfy the system's load-balance and replacement energy needs. Moreover, with time, new technologies may change this mix further. For example, as new storage-based technologies become more prevalent, and their economics and energy sustainability improves, the resources that prove most cost-effective to satisfy these same operational purposes may shift to make use of those technologies.

The broader point here is that in considering how to ensure sufficient revenue so that the resources that satisfy these operational purposes each day invest in reliable energy supply arrangements (to operate above their day-ahead awards), there isn't a specific resource 'type' or technology at issue. Rather, it is important that compensation be sufficiently dynamic to reward the resources that are the most cost-effective on any given day.

2.3.2 Implications

The preceding discussion of operational requirements and uncertainty, and how the ISO addresses them, highlights the resources and capabilities that problem P1 may impact the most in practice. As illustrated in Example 1 above, and in the next Example 2 below, generating resources that do not expect to run the next day (*e.g.,* that do not receive, and do not expect to receive, an award in the day-ahead market) may not find it financially prudent to make costly energy supply arrangements *a priori,* as they may often not be used. This, in turn, prompts concern with problem P2: there may be insufficient energy available to the power system to withstand an unexpected, extended (multi-hour to multi-day) large generation or supply loss.

In contrast, consider resources that face much less uncertainty over their energy production. As an empirical matter, since the implementation of the Energy Market Offer Flexibility market design improvements,[7] the ISO has not observed significant problems with gas-fired resources that clear in

---

[7] *See ISO New England Inc., et al.,* Energy Market Offer Flexibility Changes, Docket No. ER13-1877-000 (July 1, 2013) (incorporating in the Tariff energy market offer-flexibility enhancements to allow market participants to modify their offers to supply electricity on an hourly basis within the operating day to better reflect changing fuel costs and opportunity costs in offers). *See also ISO New England Inc., et al.,* 147 FERC ¶ 61.073 (2014) (accepting, subject to conditions, the energy market offer-flexibility enhancements).

the day-ahead market failing to have sufficient fuel to meet their day-ahead market awards. In the winter, the gas-fired resources that clear in the day-ahead market tend to be among the system's more efficient (lower heat-rate) resources. As a general rule, more efficient resources face less uncertainty over whether (and for how many hours) they will clear each day. Moreover, generating units that have a superior heat rate (lower marginal cost of production) will be willing to spend the most each day to acquire whatever natural gas is available for electricity generation in New England.[8]

At the opposite end of the spectrum, the situation can be quite different for the resources and capabilities that the ISO relies upon to manage uncertainty – *viz.,* for the operational purposes discussed earlier in Section 2.3.1. For a combination of reasons, we are concerned there are inefficiently low incentives to invest proactively in energy supply arrangements necessary to provide these capabilities reliably – even when such arrangements would be a cost-effective means to reduce reliability risks.

Specifically, the combination of reasons is that the three root causes of inefficiently low incentives under the current market construct (that is, production uncertainty (RC1), fixed costs of arranging energy supply (RC2), and their materiality (RC3)), are most likely to apply to whatever resources provide these operational capabilities. As noted earlier, resources do not receive a day-ahead market award, nor any day-ahead revenue, for the capabilities that the ISO relies upon for operating reserves, replacement energy, and load balancing. Moreover, the need to call upon these capabilities is inherently difficult for resource owners to predict – thus these resources face considerable production uncertainty (*i.e.,* RC1) over whether and how often they will be called to operate, both day-to-day and over the season as a whole.

Second, there are fixed costs (RC2) to proactively arrange energy supplies that will ensure a resource can operate if called unexpectedly during, or just prior to, the operating day. Examples include arrangements by natural gas-fired generators to procure and maintain (LNG inventories at existing LNG facilities in the Northeast (for use when the inter-state gas pipelines from the west are constrained during winter), and making advance transportation arrangements to enable fuel oil supplies to be promptly replenished during winter at the region's dual-fuel (oil and gas), distillate, and heavy-oil power plants. These types of arrangements entail up-front costs to acquire fuel (or contractual rights thereto) that can then be used by the generator 'on demand'. Yet, for the reasons shown in Section 2.2, in today's market construct it is generally unprofitable to incur the costs of arranging energy supplies that a resource does not expect to use.

Third, the beneficial impact to the system from those types of energy supply arrangements (RC3) is likely to be particularly pronounced for the resources that the ISO relies upon for operating reserves, replacement energy, and load balancing. The reason is that the system tends to rely upon those capabilities the most when it experiences adverse conditions: when gas pipelines are highly constrained, when renewable resources experience adverse weather, or any other reason that lead

[8] Moreover, owners of efficient gas-fired generators that face relatively little uncertainty over their daily production during the winter commonly follow business strategies that hedge (financially) much or most of their generators' output in advance of the winter, which makes the owner relatively insensitive to (that is, not adversely impacted by) an unexpectedly high spot price of natural-gas when scheduling fuel for their resources each day.

system conditions to change markedly from those anticipated day-ahead. During such conditions, whether or not a generator providing these capabilities has the energy to operate may have a more significant impact on market prices – and in extreme conditions, impact a potential reserve or energy shortage – than during normal operating conditions. That is, the resources providing these operational capabilities are most likely to be called upon during periods when their energy supply arrangements (or absence thereof) matter to market outcomes and to system reliability.

The bottom line is that the resources that the system relies upon for the three operational purposes discussed in detail above are those we expect to be most adversely affected by problem P1. However, while it may be a cost-effective means to reduce reliability risks for these resources to invest in additional energy supply arrangements, the current market construct provides inefficiently low incentives to do so. As a result, the ISO is increasingly concerned that the system is relying upon resources for operating reserves, replacement energy, and load balancing capabilities that have no day-ahead financial obligation and may not have sufficient energy supply to operate if called.

These observations suggest it would be desirable to improve today's market construct so that the future resource mix will invest in energy supply arrangements (*e.g.*, fuel) that ensure these ancillary service capabilities remain reliable and available to the power system each operating day.

2.3.3 Example 2 – Multiple generators, with energy and real-time reserves

To illustrate these points, we next provide an example involving operating reserves. The point is to show that the incentive misalignment problem discussed previously can adversely impact the resources that the ISO relies upon for reserves. Specifically, the wholesale market, in its current form, may not provide sufficient incentives for the owners of such resources to invest in costly energy supply arrangements, *even when* such investments would be a cost-effective means to reduce reliability risks.

Though we develop this next example in the context of energy and operating reserves, the same conclusions would hold similarly if the example instead focused on the other needs detailed in Section 2.3.1 above (namely, replacement energy and load balance reserves). The situation with operating reserves is more intricate, however, because of how operating reserves are co-optimized with energy during the operating day – and because a failure of these resources to arrange for sufficient energy to operate could magnify a reserve shortage.

In this example, there are four generators that can provide both energy and operating reserve. To capture many of the factors identified in the prior section, real-time demand is uncertain, and the higher-cost generators do not receive day-ahead market awards. We consider a situation where one generator faces the possibility that real-time demand may be high enough for it to operate the next day, and it must decide whether or not to incur the cost of arranging fuel in advance of the operating day.

As with our earlier examples, the real-life timing interpretation of 'arranging fuel in advance of the operating day' is flexible. It should be viewed as however far in advance as is necessary for the generator in question (a day, a week, a month, a season). In other words, *how far* in advance does not impact the calculations or implications of Example 2.

► **Assumptions.** The capacity and offer price parameters of the four generators are shown in Table 2-5. We assume there is a single operating reserve product procured in the real-time market, and each generator's maximum capability to provide that reserve product (due to its ramp rate) is shown in Table 2-5.

Table 2-5. Assumptions for Example 2			
Generator	**Capacity (MW)**	**Offer Price ($/MWh)**	**Reserve Capability (MW)**
Gen 1	100	$25	10
Gen 2	100	$30	20
Gen 3	50	$40	30
Gen 4	50	$90	40
Additional Cost Assumptions for Generator 3			
		Marginal Cost	**Fixed Cost**
With Advance Fuel Arrangements		$40	$150
No Advance Fuel Arrangements		N/A	N/A
Real-Time Demand Scenarios			
	Low Demand	**Medium Demand**	**High Demand**
Energy Demand (MWh)	170	190	210
Scenario Probability	33%	33%	33%

In the second panel of Table 2-5, we show additional cost assumptions for Generator 3. Its costs depend upon whether or not it arranges fuel in advance of the operating day. If Generator 3 arranges fuel in advance of the operating day, then it must incur an up-front fixed cost of $150. By 'up-front fixed cost', we mean that if the generator decides to arrange fuel in advance, it would incur the $150 cost regardless of whether or not it operates the next day. And then, in addition, it would incur a marginal cost of $40/MWh to operate – but that marginal cost is incurred only if it does indeed operate.

The additional, market level assumption are:

- Day-ahead energy demand is 190 MWh for the hour.
- Real-time energy demand is uncertain: It can be low (170 MWh), medium (190 MWh), or high (210 MWh), each equally likely, as shown in the bottom panel of Table 2-5.
- The real-time reserve requirement is 30 MWh/h.

Last, some simplifications: the time period considered is a single delivery hour; there are no transmission constraints; the generators have no commitment variables (*e.g.,* no startup costs or lead-times); and demand in the day-ahead market matches the ISO's load forecast, so there is no load-balance gap. These simplifications are intended to help focus on the essentials of energy supply incentives, and do not alter the insights of this numerical example.

► **Market awards and clearing prices.** With the above setup, we first evaluate the day-ahead and real-time market outcomes in two cases: Case A, where Generator 3 decides to make advance arrangements for energy supply in advance of the operating day; and Case B, where Generator 3 does not.

We will then examine which decision maximizes Generator 3's expected net revenue, as well as which decision would produce a superior net expected benefit to the system (*i.e.,* minimize the expected value of the system's total production cost). The day-ahead market outcome is the same in Case A and Case B, and is shown in Figure 2-1. The clearing price for energy is $30/MWh, set by Generator 2's offer price. The two lower-cost generators (Generators 1 and 2) receive day-ahead market awards, and neither of the two higher cost generators (Generator 3 and 4) receives a day-ahead market award.

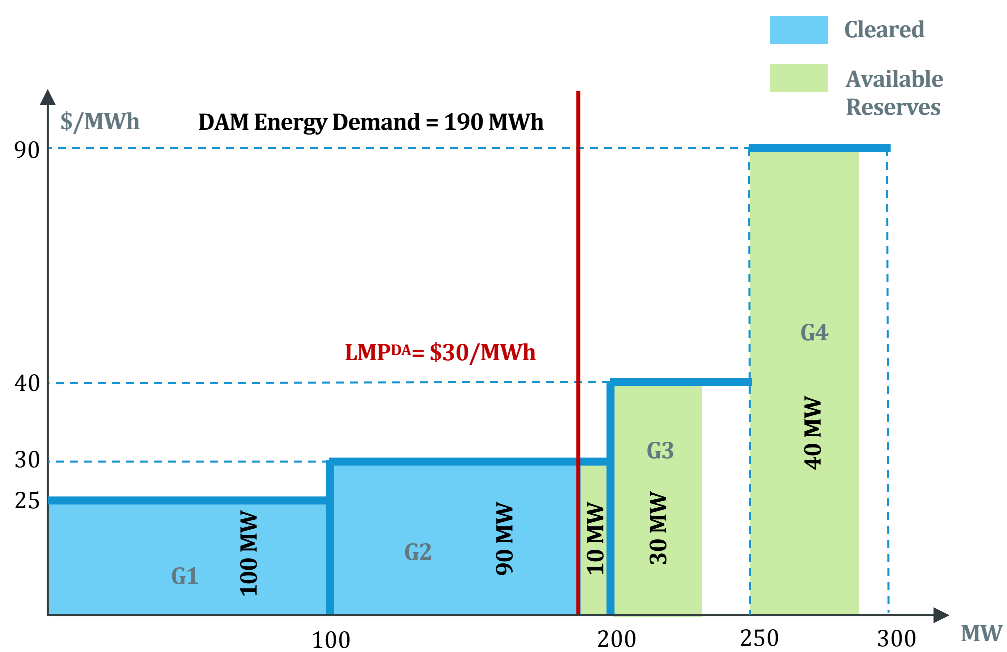

Figure 2-1. Day-ahead market outcomes for Example 2

Note that there are no day-ahead market awards for reserves in this example, which mirrors the current day-ahead market design in New England. However, the ISO can observe from the day-ahead clearing outcomes that if there are no changes in system conditions, the system would have 80 MWh of operating reserves in real-time spread across three resources (Generators 2, 3 and 4), as shown in the bars shaded green in Figure 2-1.

Case A. In real-time, demand can take one of three levels. Figures 2-2, 2-3, and 2-4 show the real-time market's energy and reserve co-optimization results for the low, medium, and high real-time demand levels, respectively.

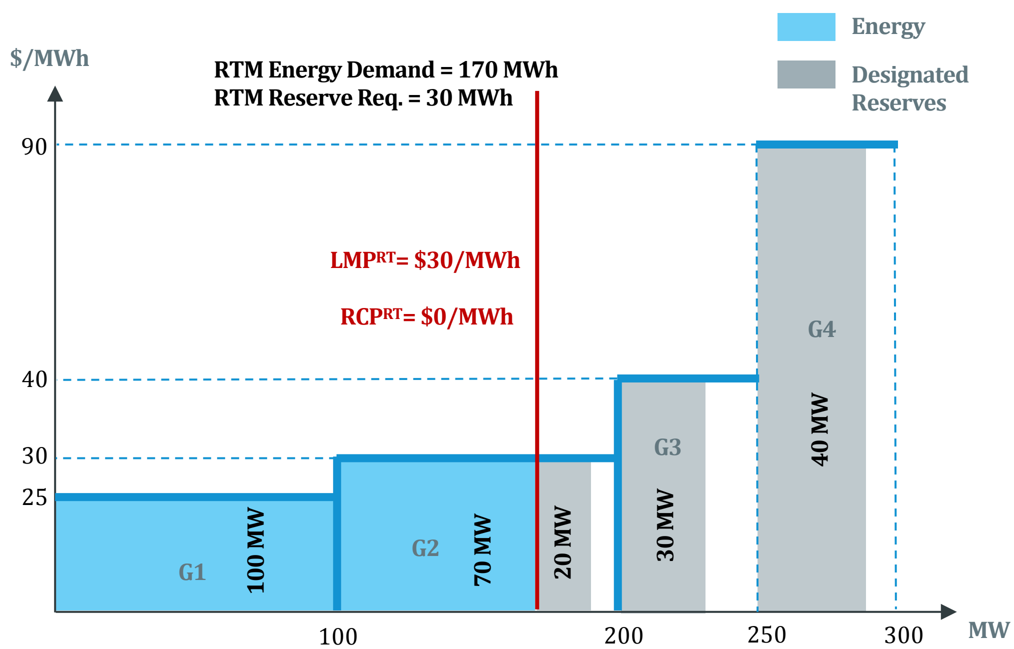

Figure 2-2. Low demand scenario real-time market outcomes for Example 2, Case A

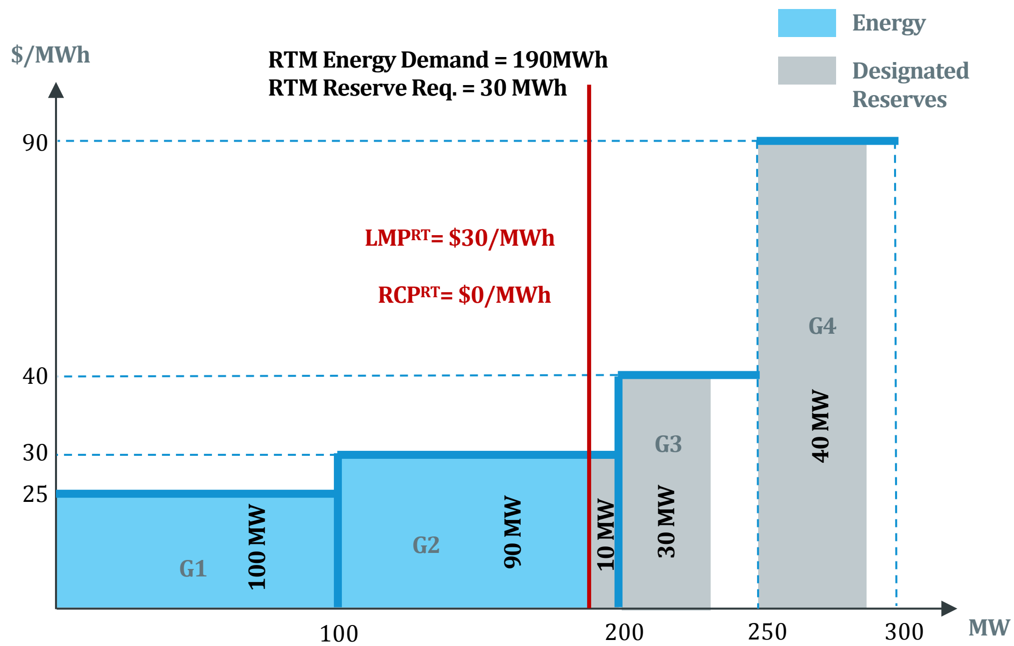

Figure 2-3. Medium demand scenario real-time market outcomes for Example 2, Case A

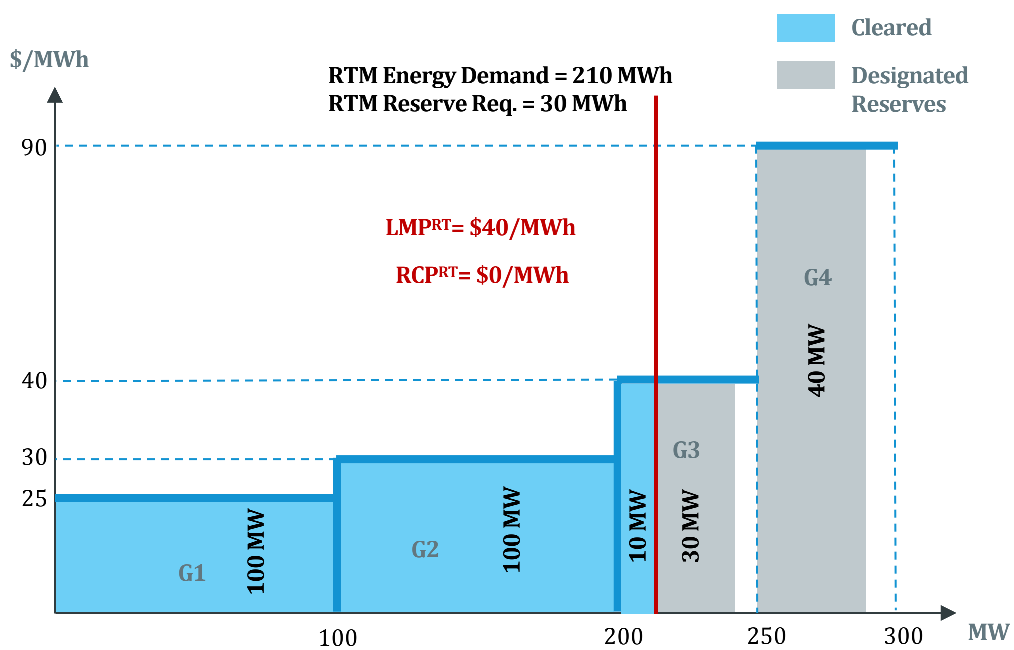

Figure 2-4. High demand scenario real-time market outcomes for Example 2, Case A

These results are summarized in Table 2-6. Figures 2-2 and 2-3, and Row [6] in Table 2-6, show that in the low and medium demand scenarios, Generator 2 remains marginal for energy and the real-time LMP is the same as day-ahead, at $30/MWh. In the high-demand scenario in Figure 2-4, Generator 3 is marginal for energy and sets the real-time LMP at $40/MWh. Row [5] indicates that in all three demand scenarios, the total supply of real-time operating reserves exceeds the reserve requirement of 30 MWh, so the real-time price for reserves is zero (as summarized in row [6]).

Table 2-6. Market Outcomes for Example 2 , Case A: Generator 3 With Fuel

		Day Ahead		Real-Time Market Outcomes					
		Market Awards		Low Demand		Medium Demand		High Demand	
	Generator	Energy	Reserve	Energy	Reserve	Energy	Reserve	Energy	Reserve
[1]	Gen 1	100	-	100	0	100	0	100	0
[2]	Gen 2	90	-	70	20	90	10	100	0
[3]	Gen 3	0	-	0	30	0	30	10	30
[4]	Gen 4	0	-	0	40	0	40	0	40
[5]	**Totals**	190	-	170	90	190	80	210	70
[6]	**Clearing Price**	$30	-	$30	$0	$30	$0	$40	$0
[7]	Scenario Total Production Cost			$4,600		$5,200		$5,900	
[8]	Demand Probability			33%		33%		33%	
[9]	**Expected Total System Production Cost**					**$5,233**			
[10]	Scenario Market Payments (incl. DAM)			$5,100		$5,700		$6,500	
[11]	**Expected Total Market Payments**					**$5,767**			

For purposes of evaluating cost-effective outcomes, we note here the system's total production costs in rows [7] and [9] of Table 2-6. Row [7] summarizes the system's total production cost in each demand scenario, under the maintained assumption that each resource offers competitively at its marginal cost. Total production costs increase with real-time demand, naturally. Importantly, in these calculations we *exclude* the $150 fixed cost of Generator 3 to arrange fuel in advance of the operating day. We will bring that into the calculations in a moment, below.

Row [9] takes the probability-weighted average of the three scenarios' total production costs, which shows that the expected (value of the) system's total production cost is $5,233 (rounding to the nearest dollar). We will compare that outcome to the expected total production cost that prevails if Generator 3 does not have advance fuel arrangements momentarily, which will identify whether the $150 up-front cost of advance fuel arrangements would be cost effective from the standpoint of the system's total production costs.

Of additional interest are total market settlements. Row [10] provides the total market settlements (including both the day-ahead market, and the deviation-based real-time settlements) for all resources in each demand scenario. Row [11] takes their probability-weighted average to obtain the expected (value of the) total market settlements of $5,767 (again rounding to the nearest dollar). In this example, the expected total market settlements are both the expected total market revenue to the generators, and the expected total payments by buyers. The expected total market settlements are greater than the expected total production costs, both here and generally, because the low-cost generators earn infra-marginal rents – the usual economic reward for superior cost efficiency in a competitive marketplace.

Since there are many numbers involved in a multi-unit market with multiple products (*i.e.,* energy and reserves), we note here the key numbers to keep in mind from Case A:

- The high-demand scenario real-time LMP is **$40/MWh**, and real-time reserve price is **$0/MWh**.
- The expected system total production cost is **$5,233** for the hour, *excluding* Generator 3's $150 fixed cost of arranging fuel (which it will incur in this Case A); and
- The expected total market settlement is **$5,767** for the hour.

Case B. Now consider the market outcomes if Generator 3 does not make arrangements for fuel in advance of the operating day. For this example, we will assume the ISO treats each generator as available unless informed otherwise by the generator, consistent with current ISO practice. If Generator 3 does not arrange fuel in advance, then we assume that it would seek to acquire fuel on short notice if instructed to operate the next day. In that situation, we assume that Generator 3 is physically unable to obtain fuel (and indicates to the ISO it not available), and the ISO dispatches the system at least-cost without Generator 3.

In Case B, the real-time market outcomes are unchanged from before in the low and medium demand scenarios; Generator 3 is not instructed to provide energy in real-time in those scenarios. The outcome is different from before in the high-demand scenario, however. In the high-demand scenario in Case B, Generator 3 would not be able to obtain fuel to operate on short notice (by assumption) and would not be available. Therefore, the real-time dispatch would turn to the next higher-cost resource in the supply stack, Generator 4. Figure 2-5 shows the real-time market outcomes in Case B's high-demand scenario.

Table 2-7 summarizes the market outcomes in Case B. To facilitate comparisons, we have shaded cells in light orange to highlight the (only) outcomes that differ from the outcomes in Case A in Table 2-6, when Generator 3 has arranged fuel.

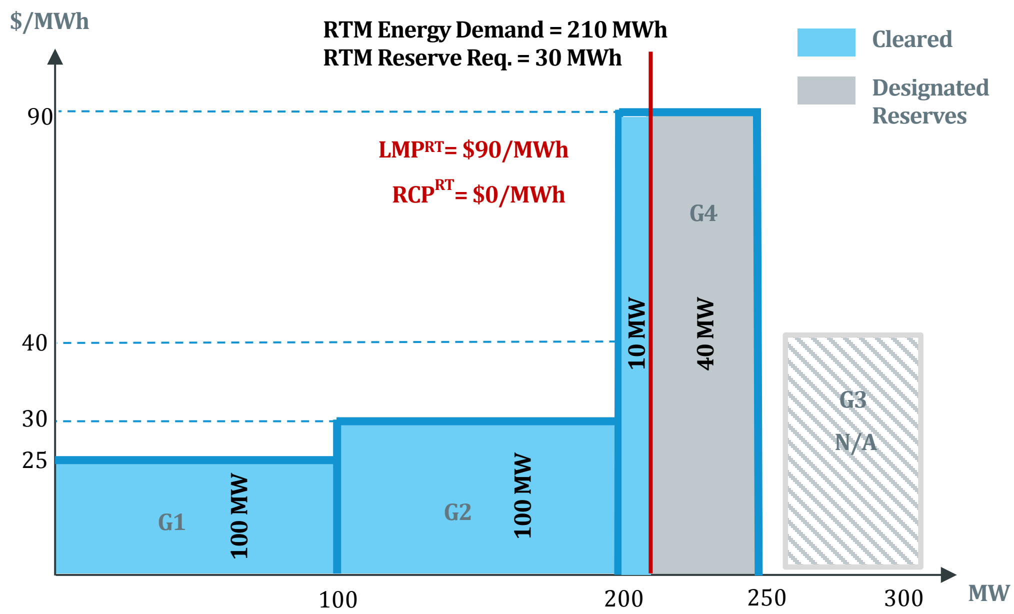

Figure 2-5. High demand scenario real-time market outcomes for Example 2, Case **B**

Table 2-7. Market Outcomes for Example 2 , Case B: Generator 3 Without Fuel

		Day Ahead		Real-Time Market Outcomes					
		Market Awards		Low Demand		Medium Demand		High Demand	
	Generator	Energy	Reserve	Energy	Reserve	Energy	Reserve	Energy	Reserve
[1]	Gen 1	100	-	100	0	100	0	100	0
[2]	Gen 2	90	-	70	20	90	10	100	0
[3]	Gen 3	0	-	0	30	0	30	0	0
[4]	Gen 4	0	-	0	40	0	40	10	40
[5]	**Totals**	190	-	170	90	190	80	210	40
[6]	**Clearing Price**	$30	-	$30	$0	$30	$0	$90	$0
[7]	Scenario Total Production Cost			$4,600		$5,200		$6,400	
[8]	Demand Probability			33%		33%		33%	
[9]	**Expected Total System Production Cos**					**$5,400**			
[10]	DAM)			$5,100		$5,700		$7,500	
[11]	**Expected Total Market Payments**					**$6,100**			

In the day-ahead market and the low and medium real-time demand scenarios, Generator 2 remains marginal for energy and the real-time LMP is the same as day-ahead, at $30/MWh. See row [6]. In the high-demand scenario, Generator 4 now is marginal for energy and sets the real-time LMP at $90/MWh.

Even with Generator 3's unavailability in the high-demand scenario, however, row [5] indicates that total supply of real-time operating reserves exceeds the reserve requirement of 30 MWh, so the real-time price for reserves is zero. Thus, in this example, Generator 3's unavailability impacts the real-time LMP (significantly, in the high-demand scenario), but its unavailability does not impact the price of reserves.

Last, for comparison purposes, we have re-calculated in Table 2-7 the system's total production costs and the total market settlements for each demand scenario in Case B, where Generator 3 does not have arrangements for fuel in advance of the operating day. Row [9] shows the expected (value of the) system's total production costs are $5,400. Row [11] shows the expected (value of the) total market settlements is now $6,100.

The key numbers to keep in mind from Case B for what comes next are:

- The high-demand scenario real-time LMP is **$90/MWh**, and real-time reserve price is still **$0/MWh**.
- The expected system total production cost is **$5,400** for the hour, where Generator 3 does not incur the $150 fixed cost of arranging fuel (in this Case B); and
- The expected total market settlement is **$6,100** for the hour.

► **Cost-effective outcome.** Now let's compare the outcomes when Generator 3 has the fuel to operate, versus when it does not. From the standpoint of operating a power system at minimum cost – in terms of the costs incurred by the suppliers to meet demand – the preferred outcome is if Generator 3 arranges fuel in advance. Even though that costs $150 up front and may not be used, it is a cost-effective investment. The system's expected total production cost without it is $5,400 and with it is $5,233, a difference of $167 – more than enough to cover the $150 up-front fixed cost of the fuel arrangements. Thus, the most efficient, cost-effective outcome for the system is if Generator 3 arranges fuel in advance of the operating day.

This same conclusion applies from the perspective of buyer's total payments, which fall from $6,100 (in Case B) to $5,767 (in Case A, with fuel). Although changes in buyers' total payments (also called *consumer surplus)* are not a measure of market efficiency, nor a measure of the minimum (most cost-effective) use of society's resources to meet demand, the reduction in total payments is logical: It avoids the scenario where high-cost Generator 4 must be used to meet demand, instead of the lower-cost Generator 3.

Before concluding that all is well, however, we need to bring a bit of the dismal science to bear on the situation. The flip-side of consumer's payments being lower in Case A (when Generator 3 has arranged for fuel) is that total market revenue to the generators is also lower in that case. The question then arises, specifically, whether – given the way the day-ahead and real-time markets currently operate – Generator 3 would find it profitable to invest in arranging fuel in advance of the operating day.

► **The generator's decision.** We now compare Generator 3's expected net revenue in each case, and whether its incentive to arrange fuel in advance is consistent with the efficient, most cost-effective outcome for the system.

The full settlement outcomes for Generator 3 in each case are detailed in Table 2-8. In brief, if Generator 3 does not arrange fuel, it produces zero energy in real-time (in any demand scenario) and its expected net revenue is $0. See the bottom right row of Table 2-8. If it does arrange fuel, Generator 3 produces energy only in the high-demand scenario. In that scenario, it is the marginal unit so it makes no profit in the real-time market (it set the real-time LMP at its marginal cost). However, it incurs the $150 up-front cost to acquire fuel. Thus, Generator 3's expected net revenue if it arranges fuel in advance of the operating day is a net financial loss, of $150.

Table 2-8. Generator 3's Expected Net Revenue in Example 2

			Case A: With Advance Fuel			**Case B:** No Advance Fuel		
Generator's Market Settlements		Calculation	Low Dmd	Med Dmd	High Dmd	Low Dmd	Med Dmd	High Dmd
[1]	Day Ahead Energy	*DA LMP * Qe_DA*	$ -	$ -	$ -	$ -	$ -	$ -
[2]	Real-Time Energy Deviation	*RT LMP*(Qe_RT - Qe_DA)*	$ -	$ -	$ 400	$ -	$ -	$ -
[3]	Real Time Reserves	*RT RCP * Qr_RT*	$ -	$ -	$ -	$ -	$ -	$ -
[4]	**Total Settlement**	[1]+[2]+[3]	$ -	$ -	$ 400	$ -	$ -	$ -
Generator's Costs								
[5]	Advance Fuel	*F*	$ (150)	$ (150)	$ (150)	$ -	$ -	$ -
[6]	Variable Cost	*MC*	$ -	$ -	$ (400)	$ -	$ -	N/A
[7]	**Total Cost**	[5]+[6]	$ (150)	$ (150)	$ (550)	$ -	$ -	$ -
Generator's Expected Profit								
[8]	Scenario Net Revenue	[4]+[7]	$ (150)	$ (150)	$ (150)	$ -	$ -	$ -
[9]	Demand Probability	*p* or (1-*p*)	0.333	0.333	0.333	0.333	0.333	0.333
[10]	**Expected Net Revenue**	SumProd [8]*[9]		**($150)**			**$0**	

The bottom line here is an important one. The energy market, in its current form, may not provide sufficient incentives for resource owners to invest proactively in energy supply arrangements – even when such investments would be cost-effective and yield expected net benefits to the system. In this example, all of the generators are acting perfectly competitively (offering at their marginal cost) throughout. Thus, the right conclusion to draw from Example 2 is *not* that there is a problem with the generators' behavior or their business acumen; rather, the right conclusion to draw from Example 2 is that there is a problem with the current energy market design.

► **Example 2-R: Reliability risks.** In Example 2, we have not (yet) raised any reliability considerations. We simply have concluded that from the standpoint of minimizing the system's expected cost, it may be efficient to arrange fuel in advance even when it may not be needed. However, that is not the only potential problem.

We now consider and extension, Example 2-R, where Generator 3's decision to arrange fuel may impact whether or not there is a reserve shortage in real-time. We make the same assumptions in in Example 2 before, but now 'scale up' two prior assumptions:

- *Higher real-time reserve requirement.* The reserve requirement is now 80 MWh; and
- *Higher fixed cost of arranging fuel.* Generator 3's up-front fixed cost to arrange fuel in advance of the operating day is now $1,200.

The first assumption will make a reserve shortage *possible* in the context of Example 2, and produce higher market prices even if the generator arranges fuel in advance of the operating day. The second assumption is related to the first: if a market may produce higher prices even in 'good' cases when the generator has fuel, then the misalignment problem tends to arise when there are higher up-front costs to arrange that fuel. In other words, the second assumption helps better reveal the misalignment problem, given the first assumption.

Stated simply, with these two revised assumptions, we have created a more "stressed system" situation in the scenario when real-time demand is high. We will see again that Generator 3 would make the same decision as before to not arrange fuel in advance, yet the system would be better off if it did. The outcomes would be more cost-effective and would reduce reliability risk.[9]

Case A. First consider the market outcomes if Generator 3 does arrange for fuel in advance of the operating day. The day-ahead market outcomes are unchanged from Figure 2-1 previously; the day-ahead LMP is again $30/MWh, set by marginal Generator 2. The total reserves available in the day-ahead market solution is 80 MWh, as before, which (just) satisfies the expected real-time reserve requirement. (Note that, under the current energy market design, the generators are not compensated for that reserve capability in the day-ahead market).

Table 2-9 summarizes the day-ahead and real-time market outcomes when Generator 3 arranges fuel in advance of the operating day. Row [6] shows that the higher reserve requirement (of 80 MWh) leads to positive real-time reserve cases in the medium and the high-demand scenarios, and therefore high production costs and energy prices in those scenarios. Row [9] shows the system's

---

[9] To simplify the analysis, in Example 2-R, we will ignore the additional settlements associated with the Pay for Performance market rules that apply during a reserve shortage, which would come into play in the high-demand scenario below. Under the assumptions for Example 2-R, the PFP performance incentives would likely change the generator's decision discussed next (we omit the supporting calculations here). As noted in Section 2.2.4, PFP therefore helps address this reliability risk. However, as illustrated earlier (see Table 2-4), if we tweaked the present example's probabilities so that there is a sufficiently lower chance of a reserve shortage, then, as illustrated, PFP would not fully resolve the misaligned incentive problem. The reasons are the same as those discussed following Table 2-4 in Section 2.2.4.

expected total production cost is $5,267 (rounding to the nearest dollar). Row [11] reports the expected total market settlement of $7,967 (rounding to the nearest dollar).

Table 2-9. Market Outcomes for Example 2-R, Case A: Generator 3 With Fuel

		Day Ahead		Real-Time Market Outcomes					
		Market Awards		Low Demand		Medium Demand		High Demand	
	Generator	Energy	Reserve	Energy	Reserve	Energy	Reserve	Energy	Reserve
[1]	Gen 1	100	-	100	0	100	0	100	0
[2]	Gen 2	90	-	70	20	90	10	90	10
[3]	Gen 3	0	-	0	30	0	30	20	30
[4]	Gen 4	0	-	0	40	0	40	0	40
[5]	**Totals**	190	-	170	90	190	80	210	80
[6]	**Clearing Price**	$30	-	$30	$0	$40	$10	$90	$60
[7]	Scenario Total Production Cost			$4,600		$5,200		$6,000	
[8]	Demand Probability			33%		33%		33%	
[9]	**Expected Total System Production Cost**					**$5,267**			
[10]	Scenario Market Payments (incl. DAM)			$5,100		$6,500		$12,300	
[11]	**Expected Total Market Payments**					**$7,967**			

Case B. Now consider the market outcomes if Generator 3 does not make arrangements for fuel in advance of the operating day. In this situation, the real-time market outcomes are unchanged in the low and medium demand scenarios. The high-demand scenario is different from Case A, however. In this scenario, Generator 3 would not be able to obtain fuel to operate (by assumption) and would be unavailable. Therefore, the real-time dispatch would instead use the next higher-price resource, Generator 4. Figure 2-6 shows the real-time market outcomes in this high-demand scenario. Here the remaining capability on the system (of Generator 1, 2, and 4) is not enough to cover the total energy and reserve requirement, and there is a reserve shortage (or reserve deficiency) of 40 MWh.

Table 2-10 below summarizes the day-ahead and real-time market outcomes when Generator 3 does not arrange fuel in advance of the operating day. The outcomes that differ from Case A in Table 2-9, when Generator 3 has arranged fuel, are in the cells shaded in light orange in Table 2-10. We assume here that the reserve clearing price is $1,000 per MWh in the high-demand scenario with the reserve shortage. Row [9] shows the system's expected total production cost is $18,733 (rounding to the nearest dollar), which incorporates the cost of the reserve shortage at its market price signal of $1,000. Row [11] reports the expected total market settlement of $26,367 (rounding to the nearest dollar).

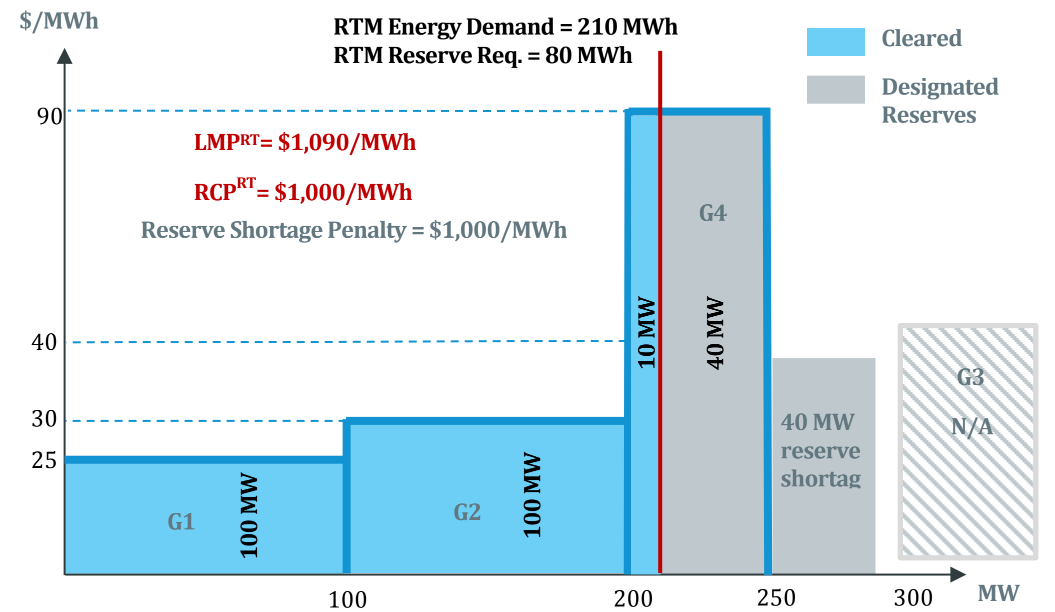

Figure 2-6. High demand scenario real-time market outcomes for Example 2, Case **B**

Table 2-10. Market Outcomes for Example 2-R, Case B: Generator 3 Without Fuel

		Day Ahead		Real-Time Market Outcomes					
		Market Awards		Low Demand		Medium Demand		High Demand	
	Generator	Energy	Reserve	Energy	Reserve	Energy	Reserve	Energy	Reserve
[1]	Gen 1	100	-	100	0	100	0	100	0
[2]	Gen 2	90	-	70	20	90	10	100	0
[3]	Gen 3	0	-	0	30	0	30	0	0
[4]	Gen 4	0	-	0	40	0	40	10	40
[5]	**Totals**	190	-	170	90	190	80	210	40
[6]	**Clearing Price**	$30	-	$30	$0	$40	$10	$1,090	$1,000
[7]	Scenario Total Production Cost			$4,600		$5,200		$46,400	
[8]	Demand Probability			33%		33%		33%	
[9]	**Expected Total System Production Cost**					**$18,733**			
[10]	Scenario Market Payments (incl. DAM)			$5,100		$6,500		$67,500	
[11]	**Expected Total Market Payments**					**$26,367**			

► **Cost-effective outcome.** Now let's compare the outcomes when Generator 3 has the fuel to operate, versus when it does not, for Example 2-R. The system's expected total production cost without it, including the reserve deficiency of 40 MWh at its shortage price, is $18,733; and with fuel arrangements, which prevent the reserve shortage, is $5,267 – a vast difference that is more than enough to cover the $1,200 up-front fixed cost of the fuel arrangements. Thus, the most efficient outcome is if Generator 3 arranges fuel in advance of the operating day.

► **The generator's decision.** Now compare Generator 3's expected net revenue in each case, and whether its incentive to arrange fuel in advance is consistent with the efficient, most cost-effective outcome for the system.

The full settlement outcomes for Generator 3 are detailed in Table 2-11. In brief, if Generator 3 does not arrange fuel, its expected net revenue is $100. See the bottom-right cell in the last row of Table 2-11. If it does arrange fuel, it is once again in the red: incurs a $167 net loss, as shown in the bottom-left cell of the last row in Table 2-11.

The point here is simple. The current energy and real-time-only reserve market design does not provide proper incentives for Generator 3 to incur the high up-front $1,200 cost of arranging energy supplies in advance, but it would both be cost-effective from society's standpoint and reduce the system's reliability risk if it did.

Table 2-11. Generator 3's Expected Net Revenue for Example 2-R, Under Status Quo/Existing Rules

			Case A: Advance Fuel			**Case B:** No Advance Fuel		
Generator's Market Settlements		Calculation	Low Dmd	Med Dmd	High Dmd	Low Dmd	Med Dmd	High Dmd
[1]	Day Ahead Energy	*DA LMP * Qe_DA*	$ -	$ -	$ -	$ -	$ -	$ -
[2]	Real-Time Energy Deviation	*RT LMP*(Qe_RT - Qe_DA)*	$ -	$ -	$ 1,800	$ -	$ -	$ -
[3]	Real-Time Reserves	*RT RCP * Qr_RT*	$ -	$ 300	$ 1,800	$ -	$ 300	$ -
[4]	**Total Settlement**	[1]+[2]+[3]	$ -	$ 300	$ 3,600	$ -	$ 300	$ -
Generator's Costs								
[5]	Advance Fuel	*F*	$ (1,200)	$ (1,200)	$ (1,200)	$ -	$ -	$ -
[6]	Variable Cost	*MC*	$ -	$ -	$ (800)	$ -	$ -	NA
[7]	**Total Cost**	[5]+[6]	$ (1,200)	$ (1,200)	$ (2,000)	$ -	$ -	$ -
Generator's Expected Profit								
[8]	Scenario Net Revenue	[4]+[7]	$ (1,200)	$ (900)	$ 1,600	$ -	$ 300	$ -
[9]	Demand Probability	*p* or (1-*p*)	0.333	0.333	0.333	0.333	0.333	0.333
[10]	**Expected Net Revenue**	SumProd [8]*[9]		**($167)**			**$100**	

2.3.4 Implications and Summary

We now consider some of the broader insights illustrated in Example 2, and connect those insights back to the ISO's concerns with operational requirements and capabilities as discussed in Sections 2.3.1 and 2.3.2 above.

First, Example 2 illustrates a market inefficiency with the current energy market design: Generator 3's market incentives are to *not* incur the fixed costs of arranging fuel in advance of the operating day, but society would be better off if it did. In plain terms, the current market construct does not incent cost-effective outcomes.

As in the earlier examples, the key insight is that there are three root causes of this market inefficiency:

RC1. Generator 3 faces significant production uncertainty – after all, there is only a 33% chance it will be in demand the next day;

RC2. Generator 3 faces significant fixed costs (of $150), relative to its expected gross margin (which, in this example, is zero) – leaving it no infra-marginal revenue with which to cover that fixed cost;

RC3. Generator 3's decision to invest in fuel arrangements (or not) impacts the resulting market price for energy – *i.e.,* what consumers value consuming (at least, with positive probability). Doing so enables it to produce in the high-demand scenario, rather than forcing the system to use the next higher-cost resource in the supply stack (Generator 4 – and, in the more stressed case, with a reserve shortage).

Thus, in incurring the private cost of investing in energy supply arrangements in advance of the operating day, the generator cannot recoup that investment in this energy market design. However, that same investment would produce more than enough savings in expected total system production costs to make it efficient and cost-effective for the system overall. This difference between the private benefits of the investment (to Generator 3) and the expected total production cost savings (to society) results in the misaligned incentive problem, and higher expected costs to society as a result.

Second, Example 2 is structured to illustrate why the root causes in RC1 through RC3 are of potential concern for the resources and capabilities that the system relies upon to satisfy the three operational purposes itemized earlier in Section 2.3.1. Generator 3 is extra-marginal in the day-ahead market and does not have a day-ahead award, which is characteristic of the resources the ISO relies up for those three operational capabilities each day. Generator 3 has slim energy market gross margins (infra-marginal revenue) on the occasions when it is dispatched for energy, providing little revenue with which to recoup – therefore little financial incentive to incur – the up-front fixed cost of arranging fuel in advance. And yet, during stressed system's conditions, Generator 3's operation is essential to prevent the system from needing to turn to much higher-cost generators to meet demand (and, in extreme cases, to avoid a reserve shortage).

Third, Example 2 also shows why these root causes may not apply (or not apply to the same extent) to some of the system's lower cost, more efficient resources that clear in the day-ahead market. Imagine, for example, that the low-cost Generator 1 and higher-cost Generator 3 in Example 2 both faced similar $150 fixed cost of arranging fuel in advance of the operating day. The efficient outcome would also be for Generator 1 to incur that cost. Would it be financially incented to do so, under the current energy market? Yes. In Example 2, Generator 1 makes a $500 gross margin *in the day-ahead market*, easily enough to motivate – and recoup the cost of – a $150 up-front fixed cost of arranging its fuel in advance of the operating day. This logic, though simplified in the context of Example 2, mirrors the real-world economic rationale for why the ISO has not observed significant problems with gas-fired resources that clear *in the day-ahead market* failing to have sufficient fuel to meet their day-ahead energy market awards (*see* Section 2.3.2).

Last, from a reliability perspective, Example 2 illustrates that the system is potentially relying upon resources for reserves that may not be able to obtain fuel if dispatched for energy during the operating day. In Example 2, Generator 3 does not *expect* to operate and it *plans* to acquire fuel if dispatched (if it does not arrange fuel in advance – see Case B on page 34). From Generator's 3 perspective, it is not financially prudent to incur the costs of arranging fuel in advance, knowing that the fuel most likely will not be used. However, as a result, based on the day-ahead market outcome Example 2-R (see Figure 2-5), the ISO would anticipate having 80 MWh of reserves when preparing the next-day operating plan – even though Generator 3 may not be able to produce. As Example 2 illustrates, for the resources and capabilities to which the three root causes above reasonably apply, it is logical to be concerned that the region may find these resources do not have sufficient energy supply arrangements if called during the operating day.

Our broad conclusion from these observations is that those resource owners' are acting rationally given the operating uncertainties and difficult economic circumstances they face – and that the problem lies in the existing energy market design. The ISO's energy and ancillary services markets have not changed, in their fundamental product suite, for about fifteen years; and they were not designed in anticipation that the three root causes identified above would present a material issue for much of the generation fleet. The misaligned incentive problem that these three root causes precipitate is not likely to solve itself, given the evolving resource mix in New England's power system and the greater operational uncertainties with evermore just-in-time energy sources. Therefore, we conclude that it is important to develop market improvements that will better align the incentives, so that generators will find it in their private interest to invest in additional energy supply arrangements whenever those arrangements would be a cost-effective means to reduce the system's reliability risk.

2.4 Problem P3: Inefficient Schedules

We now turn to problem P3, concerning inefficient schedules. As noted at the outset of Section 2, with the region's growing dependence on just-in-time energy sources and its constrained fuel delivery infrastructure, assets with a supply of energy 'in storage' that can easily be used for electricity production today or tomorrow (or beyond) provide a valuable 'buffer stock' of energy to the power system. The buffer stock of these *non*-just-in-time generation technologies enables the system to reliably serve demand when gas pipelines are most constrained, renewable resources experience adverse weather, or both.

The ISO has growing concerns that the current energy market's scheduling horizon may deplete the power system's limited pool of stored energy prematurely, particularly during extended cold weather conditions. Like other organized wholesale electricity market regions in the US, the ISO employs a single (that is, one day in advance) day-ahead wholesale electric energy market. As a result, the day-ahead market develops schedules for power production that do not evaluate outcomes, or produce prices, for days beyond tomorrow. In effect, the market-clearing and scheduling process formally treats the value of preserving a generator's stored energy beyond tomorrow as zero.

This current market design leads to the concern identified at the outset of Section 2 as problem P3: The power system may experience premature (inefficient) depletion of energy inventories for electric generation, absent a mechanism to better schedule and reward efficient preservation of limited-energy supplies over multiple days. That, in turn, may magnify the challenge of problem P2: If the system's limited pool of stored energy is depleted prematurely, there may be insufficient energy available to withstand an extended (multi-day) large generation or supply loss during cold weather conditions.

► **Information and coordination.** At its core, scheduling limited energy resources over time is a coordination problem. These types of coordination problems are normally overcome by markets when there is ample information available to participants about the value of preserving their individual buffer stock of energy to a future date, rather than using it today. However, the current one-day-ahead market – in effect, a *one-day-at-a-time* market – is inherently myopic: it does not provide participants with any price signal of the value, in hard dollar-per-MWh terms, of preserving (or acquiring additional) energy supply.

Presently, suppliers may account, in their energy supply offer prices, for estimated intertemporal opportunity costs of preserving fuel as a means to help preserve scarce energy supplies for a later date. However, the price forecasts these intertemporal opportunity costs are based on are not derived from an efficient, multi-day scheduling process, nor informed by the state of the system's stored energy buffer stock as a whole. From an informational standpoint, each generation owner knows *its* limited energy inventory (*e.g.*, the fuel in its tanks), but not the state of energy supply limitations across the system. And a single day-ahead market that produces no prices for days *after* tomorrow cannot provide a generation owner with a market signal of the value of preserving its limited energy inventory for use later (*viz*., the day after tomorrow, or the day after that, and so on).

The ISO does review considerable information on the energy supply limitations of generators for coming days (due to resource-level fuel surveys and gas-electric information sharing with pipeline operators and certain LNG facilities). In principle, resources with limited pools of stored energy could have those energy limitations accounted for in a market clearing process spanning multiple days, so that they are scheduled more efficiently over time and the (shadow) prices of their limited energy stocks would be incorporated into forward-looking market prices for future days. But without a multi-day market clearing mechanism, there is presently no means to coordinate generators' multi-day production schedules and aggregate the information on resources' limited stored energy in a least-cost, most-reliable way.

Moreover, there is an additional problem when suppliers in New England's system seek to preserve limited stored energy supplies, without a forward price signal at which to sell that stored energy for future delivery. A generator in New England that delays using its limited stored energy earns no revenue for doing so today, in the hope of receiving a greater price using that inventory at a later date. Yet that incentive commonly arises at the same time for many similarly-situated generators in New England. Though acting individually, if many delay their use of their limited energy stocks to the date with the highest *expected* future price (via opportunity-cost-based offers in today's one-day-ahead market), their combined effort will result in a *low* price on that very future date.

That is effectively another manifestation of problem P1: the very act of incurring an up-front cost of deferring production (thus foregoing revenue today) changes the hoped-for price in the future, eliminating the economic return from deferring production in the first place. In simple terms, when generators delay their use of limited fuel stocks to a more critical forecasted period several days later, that action generally reduces the total cost of operating the power system and its reliability risk. But, in return, the existing one-day-ahead market design may provide little compensation for those salutary actions.

The general solution to this problem is to provide additional forward market opportunities. That serves two functions: one, it provides a clear price signal to all of the value of preserving limited energy to a (specific) later date; and two, it enables participants to capture that value by selling energy forward and "locking in" the price from preserving it. To do that within the ISO-administered marketplace would amount to extending the existing single day-ahead market into a true multi-day ahead market that can optimize the entire system's limited energy supplies over time and directly compensate each supplier – at competitively-determined prices – for their roles in doing so.

Ultimately, such a multi-day market design would provide resource-level forward energy schedule visibility to market participants and the ISO alike. It also provides an economically sound method to further reduce (or eliminate) the need for costly, out-of-market posturing of resources to preserve generators' scarce energy supplies when the region experiences extended cold weather conditions.

In the fall of 2018, the ISO provided stakeholders with a number of numerical examples of these issues at the NEPOOL Markets Committee meetings.[10] In a future update of this paper, we anticipate extending those examples to further explore the causes and consequences of premature depletions, how opportunity cost-based offers can help (but not fully resolve) this problem, and the

[10] ISO New England Presentation to NEPOOL Markets Committee, *Winter Energy Security Improvements: Market-Based Approaches* (Dec. 11-12, 2018), https://www.iso-ne.com/static-assets/documents/2018/12/a2a_iso_presentation_winter_energy_security_improvements.pptx; ISO New England Presentation to NEPOOL Markets Committee, *Winter Energy Security Improvements: Market-Based Approaches* (Nov. 7-8, 2018), https://www.iso-ne.com/static-assets/documents/2018/11/a2_presentation_winter_energy_security_improvements.pptx; ISO New England Presentation to NEPOOL Markets Committee, *Winter Energy Security Improvements: Market-Based Approaches* (Oct. 10, 2018), https://www.iso-ne.com/static-assets/documents/2018/10/a9_presentation_winter_energy_security_improvements.pptx.

"catch 22" generators face when preserving limited energy supplies without a forward price to sell against. There are also interdependencies between a multi-day ahead market and participation in the new ancillary services markets discussed in the next section of this paper, and those interactions warrant more detailed attention. We anticipate that exploring these issues with additional examples may help facilitate further stakeholder discussions and feedback on the design, settlement, and functioning of a multi-day ahead market.

3. Objectives and Design Principles

To make tangible progress toward effective longer-term market solutions, it is useful to proceed from a concise statement of design objectives and principles. Although the set of possible objectives is large, and stakeholders may have varying perspectives on their relative importance, the ISO's approach to developing sound solutions reflects the following objectives and design principles.

3.1 Three Broad Objectives

In concise terms, the ISO suggests three broad objectives that define the desired outcomes for both near-term and longer-term regional energy security improvements. These are:

1. **Risk Reduction.** Minimize the heightened risk of unserved electricity demand during New England's cold winter conditions by solving problems P1, P2, and P3.

2. **Cost Effectiveness**. Efficiently use the region's existing assets and infrastructure to achieve this risk reduction in the most cost-effective way possible.

3. **Innovation.** Provide clear incentives for all capable resources, including new resources and technologies, that can reduce this risk effectively over the long-term.

We anticipate this third objective will become increasingly important over time, as the region's older (non-gas-fired) generation facilities reach the end of their service lives and as the New England states work steadily to advance de-carbonization goals.

Based on the analysis of problems and root causes, our tangible recommendations to achieve these broad objectives entail several meaningful measures. One is to strengthen generation owners' financial incentives to undertake more robust energy supply arrangements, when cost-effective, while not proscribing what form those supply arrangements may take. We view that tangible goal as consistent with all three broad objectives above.

Another is to improve the today's electricity market construct to better reward resource flexibility that helps manage, and prepare for, energy supply uncertainties during the operating day. These uncertainties may well become more challenging over time, given the increasingly just-in-time nature of New England's power system. Making these improvements through competitive, transparent market mechanisms that reward all capable resources – regardless of technology – serves all three broad objectives above.

A third measure is to efficiently allocate *over time* (*i.e.,* across multiple days) electricity production by resources that have stored – *non*-just-in-time – energy sources. Resources that have stored energy sources –whether traditional hydro-electric or oil-fired generators, or new grid-scale storage technologies (*e.g.,* batteries) – can provide a valuable "buffer stock" of energy supply to the power system. This may become increasingly important when New England's gas pipelines are most constrained, renewable resources experience adverse weather, or both. A market means to provide participants with a price signal of the *value* – in hard dollar-per-MWh terms – of preserving (or quickly arranging additional) energy supplies is consistent with all three broad objectives above.

These measures serve as milestones on our path toward long-term market solutions.

3.2 Design Principles

In developing market improvement recommendations to achieve these three objectives, we are guided by several core design principles. These design principles usefully circumscribe the means through which we achieve the foregoing objectives.

DP1. **Product definitions should be specific, simple, and uniform.** The same well-defined product or service should be rewarded, regardless of the technology used to deliver it. Simplicity in product definitions enhances competition and participants' understanding of their obligations.

DP2. **Transparently price the desired service.** A resource providing an essential reliability service (for instance, a call on its energy on short notice) should be compensated at a transparent price for that service.

DP3. **Reward outputs, do not specify inputs.** Compensating for obligations to deliver the output that a reliable system requires creates a level playing field for competitors that deliver energy reliably through cold-weather conditions. This rewards suppliers that reduce risk in the most cost-effective ways – and fosters innovation in new solution technologies.

DP4. **Sound forward markets require sound spot markets.** Forward-market procurements work well when they settle against a transparent, spot price for the underlying service.

DP5. **Compensate all resources that provide the desired service similarly.** Paying similar rates for similar service is non-discriminatory by fuel-type or technology, and sends the broadest-possible market signal for the desired attribute.

These are not novel but rather 'tried and true' market design principles, rooted more in practicality than philosophy. They help to guide concepts toward tangible solutions that will be robust and will continue to function properly as market fundamentals change over time.

Indeed, as the economic environment evolves, a good solution will not need to be continually revisited, and its market rules will not need to be successively perturbed. Achieving that requires a solution approach that employs sound economic principles, integrates well with the existing wholesale market structure (both from a technical standpoint for the ISO, and from a commercial standpoint for Market Participants), and minimizes administrative rules, restrictions, and parameters whose appropriateness may not persist as the system evolves.

Ultimately, the market design improvements we discuss next seek to help allay the region's tensions that have emerged over New England's energy security challenges in recent years, achieve a workable solution that the ISO can implement, and provide sustainable market enhancements by adhering to familiar market design principles.

4. Solution Concepts

In this section, we discuss a set of recommended market design improvements to address the problems detailed in Section 2, and to achieve the objectives summarized in Section 3. At a high-level, the overall design approach builds upon familiar energy and ancillary service concepts in the wholesale electricity markets. Broadly, we propose expanding the existing suite of energy and ancillary service products in the ISO-administered markets, in order to address – reliably and cost-effectively – the uncertainties and supply limitations inherent to a power system evermore reliant on just-in-time energy technologies.

As introduced in Section 1, we envision three interrelated market design improvements to help achieve these objectives. These are:

- **Multi-day ahead market.** Expand the current one-day-ahead market into a multi-day ahead market, optimizing energy (including stored fuel energy) over a multi-day timeframe and producing multi-day clearing prices for market participants' energy obligations.

- **New ancillary services in the day-ahead market.** Create several new, voluntary ancillary services in the day-ahead market that provide, and compensate for, the flexibility of energy 'on demand' to manage uncertainties each operating day.

- **Seasonal forward market.** Conduct a voluntary, competitive forward auction that provides asset owners with both the incentive, and necessary compensation, to invest in supplemental supply arrangements for the coming winter.

In this version of this paper, we focus below on the rationale, design, and properties of the second of these enhancements: new ancillary services in the day-ahead market. The second design element entails new products, and the framework developed below is intended to facilitate stakeholder discussion and feedback as the detailed design and technical elements for these products is developed.

A discussion of the rationale, properties, and illustrative examples for the first enhancement, a multi-day ahead market, will be addressed in a future release of this paper. As noted in Section 2.4, the ISO provided initial examples and discussion materials related to the multi-day ahead market at stakeholder meetings in November 2018.[11]

With respect to the third enhancement, the ISO remains in the early, conceptual stages of evaluating designs for such a forward market. One means of doing so may be to substantially re-vamp the existing Forward Reserve Market so that, in effect, it becomes a forward market for the same suite of new ancillary services discussed below and that would be subsequently transacted in the new day-ahead market. We anticipate developing the key elements – and, ultimately, design details – for a forward market after further discussion of its potential scope, and precise objectives, with stakeholders.

4.1 New Ancillary Services in the Day-Ahead Markets

Section 2.3 described three categories of resource capabilities that the system relies upon, above and beyond resources' day-ahead energy market awards, to prepare for and to help manage the potential energy 'gaps' that can arise in the next-day operating plan. We propose to formalize these three categories of operational requirements into specific ancillary service capabilities, and allow resources to compete to provide those capabilities in an expanded day-ahead energy and ancillary services (E&AS) market. Broadly, the purpose is to improve today's market construct so that the future resource mix will undertake energy supply arrangements and/or pursue technologies that ensure these ancillary service capabilities remain available to the power system each operating day.

In concrete terms, these three categories of day-ahead ancillary services are:

A. **Generation Contingency Reserves (GCR).** Generation contingency reserves refer to the three products that the ISO currently designates and maintains in the real-time market for operating reserves. These are ten-minute spinning reserves (TMSR), ten-minute non-spinning reserves (TMNSR), and thirty-minute operating reserves (TMOR). In simpler terms, all three are forms of fast-start or fast-ramping generation capability, and the ISO proposes to procure, and compensate for, these same three capabilities in the day-ahead E&AS market.

   The ISO relies upon these capabilities to ensure the system is prepared to promptly restore power balance (consistent with the timeframes established in applicable reliability standards) in response to a sudden, unanticipated power supply loss during the operating day.

B. **Replacement Energy Reserves (RER).** Replacement energy reserves are a new product to be procured in the day-ahead E&AS market. At a high-level, this provides the ISO with the option to "call" on the energy of a RER resource, above and beyond its day-ahead energy market award (if any), over a timeframe of more than an hour. RER would not be

[11] *See n.* 10, *supra.*

procured in the real-time market, but is obtained on a day-ahead basis as it serves to ensure the system has a reliable next-day operating plan.

The ISO relies on RER capabilities to provide energy supply needed to replace a day-ahead cleared resource that is unexpectedly unable to operate for an extended (multi-hour to multi-day) duration. As noted earlier (*see* Section 2.3.1(B)), the ISO can rely on the GCR resources in category A for energy for only a limited amount of time after a contingency (*e.g.,* a few hours or less – a duration that may vary with resource availability, demand, and other operating day conditions). After that point, the replacement energy to cover a contingency's balance-of-day energy gap must come from the dispatch and commitment of other resources operating above day-ahead award (if any).

C. **Energy Imbalance Reserves (EIR).** Energy imbalance reserves are a new product to be procured in the day-ahead E&AS market. Like RER, the EIR provides the ISO with the option to "call" on the energy of an EIR resource, above and beyond its day-ahead energy market award, if any. Unlike RER resources, the response time (ramp-up or startup timeframe) for EIR resources may be considerably longer, as EIR serves a different purpose: Energy to supply the load-balance 'gap' when generators' total day-ahead energy supply awards are less than the ISO's load forecast, in one or more hours, for the next (operating) day. EIR would not be procured in the real-time market, but is obtained on a day-ahead basis as it serves to ensure the system has a reliable next-day operating plan.

The ISO relies upon EIR capabilities to help ensure there is sufficient energy to cover the forecast load each hour of the next operating day – which frequently (but not always) exceeds the total generation cleared in the day-ahead energy market.

At the outset of Section 2 on problems and causes, we identified three specific problems that help explain the region's energy security risks and the need for market design improvements: Inefficiently low incentives (problem P1), operational uncertainties (problem P2), and inefficient scheduling of resources (problem P3). The three new day-ahead ancillary services are intended to address the first two of those problems, P1 and P2. The multi-day ahead market's focus is on addressing problem P3 and, to a different extent, also helps with problem P1.

Specifically, for the resources and capabilities that the ISO relies upon for these three operational purposes, the proposed design of these day-ahead products directly addresses the misaligned incentives problem explored previously (in Section 2). That is, the proposed design seeks to better align incentives so that generators will choose to invest in additional energy supply arrangements when those arrangements are a cost-effective means to reduce the system's reliability risk.

4.1.1 Practicalities: Key Components and Rationales

As an initial matter, developing new products in the wholesale electricity markets requires careful attention to details such as resources' offer formats, requisite capabilities, settlement processes, and so forth. We expect to provide a more detailed discussion of these elements in later versions of this paper, as these design details are developed.

Here, we provide a higher-level overview of the ISO's current thinking on the key concepts and properties for the new day-ahead ancillary services. Our immediate purpose is to provide sufficient conceptual clarity to enable productive stakeholder discussions on these concepts and to facilitate subsequent work to refine the detailed design elements.

► **Product definitions: concept.** At a high level, a day-ahead seller of these ancillary services is providing the ISO with a "call option" on its resource's energy during the operating day. We expect that different time-related parameters (*e.g.,* generator ramp or startup times) will be relevant to the awards of different ancillary service products. A generator that provides these services would have both a day-ahead and a real-time settlement for the hour(s) for which it acquires a day-ahead ancillary services obligation.

This approach to the product definition is intended to be specific, simple, and uniform, thereby satisfying design principles DP1 (specificity and simplicity in product definitions), DP3 (rewarding the output capabilities the system requires, not suppliers' inputs), and DP4 (the day-ahead forward sales of ancillary services have a delivery-dependent real-time market settlement). This combination ultimately helps achieve all three broad objectives in Section 3.

► **Pricing and compensation.** The market clearing optimization formulation in design for the day-ahead E&AS markets would compensate both energy and ancillary service awards at uniform, transparent, product-specific market prices. This is intended to satisfy design principles DP2 (price transparency) and DP5 (non-discrimination).

As noted at the end of Section 2.3.1 (p. 26), the clearing prices of these day-ahead ancillary services would vary over time (*i.e.,* each day), as their supply and demand dictate. In this way, the pricing serves to reward the resources that are the most cost-effective suppliers of each product on any given day. Importantly, we are cognizant that the clearing prices of each day-ahead product should account for the *inter-product opportunity cost* that can arise if a seller is awarded one particular day-ahead product instead of a different day-ahead product for the same delivery hour(s) (more about this below).

► **Participation.** We envision that offers to provide these ancillary services will be voluntary. This would enable resource owners that expect their only profitable opportunity in the ISO's energy and ancillary services is to continue to sell just energy (*i.e.,* and not day-ahead ancillary services) to be free to continue to participate that way. The ISO's regulation ancillary service market, for example, operates on a similar premise.

At the opposite end of the participation spectrum, a resource owner that wishes to submit day-ahead offers to sell energy and (one or more of) the day-ahead ancillary services would be free to do so.[12] The clearing process would be designed to select the most valuable assignment of offers to

[12] We expect the resource requirements to provide each day-ahead ancillary service product (GCR, RER, and EIR) will differ by product (*e.g.,* with regard to ramp-up rates or startup timeframes), as befits each product's purpose. The specific capabilities necessary to supply a product will be addressed as the detailed design work proceeds for these day-ahead ancillary services.

awards, so as not to award multiple obligations to the same MWh of energy capability in the same delivery hours.

► **Day-ahead co-optimized clearing.** To procure these services cost-effectively, we recommend that the award of these ancillary services be co-optimized (*i.e.*, simultaneously cleared) with all participants' energy supply and demand awards in the multi-day ahead market. That enables the market-clearing process to determine the most cost-effective assignment of resource offers to awards for all products.

That co-optimization process ensures, by design, that the clearing prices for energy and each ancillary service incorporate the (respective marginal) suppliers' opportunity costs of not receiving an award for a different day-ahead product. It also means that, whenever these inter-product opportunity costs are non-zero (as determined in the clearing process), the day-ahead LMPs for energy will incorporate the clearing prices for the ancillary services as well. The ISO's real-time energy and operating reserve co-optimization logic has similar pricing properties presently.

► **Quantities.** We anticipate the quantities to be cleared for each day-ahead ancillary service product would reflect, at a minimum, the procedures applied by the ISO in developing a reliable next-day operating plan summarized earlier in Section 2.3.1 (see *Quantities*, p. 24). Those quantity requirements are neither *ad hoc* nor novel. As discussed in that section, those requirements are inherently dynamic, varying day-by-day based (in large part) on the forecast demand profile, the generation cleared for energy in the day-ahead market, and the system's largest anticipated potential single-source energy losses during the course of the operating day.

► **Product substitution.** Conceptually speaking, there is some 'overlap' between the different ancillary service products and the capabilities of resources to supply them. For instance, a unit that is capable of providing GCR, but that is not cleared as GCR (economically) for a particular delivery hour, might instead be economically cleared as EIR or RER. The reverse is not necessarily true (*e.g.,* depending on a resource's startup-time or ramp-rates, a resources that can provide EIR may not be capable of supplying CGR).

At a different level, the capabilities procured for GCR can, for a period of time after a contingency, potentially serve to meet some of the system's replacement energy requirements during the same hours. The details of these relationships, which can be quite technical, are formally known as the *product substitution structure*. These will be addressed in detail as the design of these ancillary services products is further developed.

► **Real-time remains least-cost dispatch.** Because suppliers in the New England markets have the ability to update their energy supply offer prices ("reoffer") during the operating day when their costs change, the least-cost solution to the system's total real-time energy and reserve requirements each hour of the operating day may be different than the day-ahead solution. Thus, the dispatch of energy, and the economic evaluation of any additional commitments (whether fast-start or otherwise) after the day-ahead market, will continue to be performed based on the real-time resource offers in effect at the time.

The real-time dispatch of the system would continue to perform co-optimization of energy and operating reserves only, as is the case today. As noted above, there is no explicit real-time analog to the EIR and RER products. These products provide the ISO with a call option on energy, and energy is their corresponding (or "underlying") real-time product.

Importantly, our current thinking is that the day-ahead GCR products (*viz.,* day-ahead TMSR, TMNSR, and TMOR) be settled similarly to the settlement of all other day-ahead ancillary services (RER and EIR). The GCR products do have real-time analogs (namely, real-time dispatch operating reserve designations and have real-time pricing). However, the incentive properties created by a day-ahead GCR ancillary service obligation can be superior – *i.e.,* more efficient – when day-ahead GCR obligations are settled similarly to the RER and EIR obligations. We address this issue further below.

Stated more generally, the overall day-ahead and real-time market structure means that, like today, all obligations acquired in the day-ahead market – whether for energy or for an ancillary service – will be financially-settled in the real-time market based on the real-time locational marginal price (LMP) for energy. We discuss that settlement logic, and the strong incentive properties it provides, next.

4.2 Settlement of Day-Ahead Ancillary Service Awards as Real Options

An important feature of these new ancillary services' design is their settlement. Since the same settlements logic and rationale applies to each of the proposed day-ahead ancillary service products, our discussion below is applicable to the proposed day-ahead GCR, RER, and EIR products.

Consistent with the concept and value of a call option on their energy during the operating day, we propose to settle each day-ahead ancillary obligation like a call option on real-time energy. That provides a familiar settlement rule used in a wide variety of markets, and will function well in concert with the existing day-ahead energy market's two-settlement design. However, the second (real-time) settlement would be slightly different for day-ahead energy and for ancillary service positions, reflecting that the former is selling (or purchasing) *forward energy* and the latter is selling an *option on energy*.

Although day-ahead ancillary service obligations are financially settled, these products are real options on energy – in the formal sense that the settlements are expressly based on what the resource physically produces in real time. This real-option design creates new financial incentives for sellers of these ancillary services to ensure they have the physical wherewithal (including fuel) to cover their obligations the next day. This is because a resource that commits to providing an ancillary service will face a financial consequence in real-time settlement if the real-time energy price is high and the resource does not perform. At the same time, resources will receive day-ahead compensation to cover the fixed costs of such supplemental energy (*e.g.,* fuel) supply arrangements, even if it turns out that they are not needed to operate the next day.

These product design and settlement features fundamentally change the incentives present today. From a commercial standpoint, it becomes profitable for the resources that the ISO relies on for these ancillary services to incur the costs of maintaining reliable fuel arrangements, when such arrangements are cost effective from the standpoint of the system overall – helping ensure they can perform if dispatched to fill an energy gap, even on a day they did not expect to operate.

To illustrate and explain these beneficial incentive properties, below we provide a series of numerical examples. For clarity, we start next with a summary of how option settlements work, oriented to the context of the proposed day-ahead E&AS market settlement. We then examine the prior numerical examples from Section 2, and show how this design overcomes the concrete problems (P1 and P2) identified earlier – and helps solves the incentive misalignment problem present in the existing energy market construct.

4.2.1 Energy Option Settlements: Logic and Simple Examples

In today's day-ahead energy market, all forward energy sales (and purchases) are financially-settled. That settlement is based on the energy produced in real-time, and the real-time energy price. In the proposed day-ahead E&AS market, all day-ahead energy *and* all ancillary services will also be financially-settled, based on those two elements as well. However, the real-time settlement for energy and for an ancillary service is slightly different.

In this section, we summarize the settlement mechanics for day-ahead ancillary services and provide several simple examples. These examples mirror the established ways call options are settled in markets that clear both forward sales (and purchases) and call options for future delivery.

► **Components and mechanics.** First, the settlement components. A call option involves three elements: (1) the sale of the option, which occurs at the *option price,* and we denote by *V*; (2) the option *strike price*, which we denote by *K*; and (3) the real-time price of whatever good the seller is providing an option on, which in our context is the real-time LMP.

The strike price is a pre-defined value, set *before* sellers specify their offer prices in a day-ahead E&AS market. The strike price, and the level at which it is set, plays a key role in the beneficial incentives an E&AS design creates. (We'll explain that in greater detail below.) For the moment, think of it as simply a threshold price, known to all before the option is offered and cleared.

Like a forward sale of energy, a call option involves both a day-ahead and a real-time settlement. The day-ahead settlement is a payment to the seller at the day-ahead option clearing price, *V,* for each MWh of the option sold. The real-time settlement is based on what the seller delivers in real-time (if anything), and has two parts. The first part is a charge, for each MWh of the option sold day-ahead, equal to the real-time LMP *minus* the strike price *K,* if that difference is positive. In mathematical terms, this is written as:[13]

$$-\max\{0, RT\ LMP - K\}.$$

[13] Note settlement sign conventions here: A negative number is a charge to the participant; a positive number is a credit. This convention is followed in all the settlement tables throughout this paper.

The second part of the real-time settlement is a credit, at the real-time LMP, for the MWh the resource actually produces.

This standard settlement rule has the effect that if a resource sells 1 MWh of day-ahead ancillary service and is dispatched for (and produces) 1 MWh of energy in real-time, the real-time credit to the seller is, at most *K.* If the resource does not produce in real time, it incurs a real-time charge to cover the ISO's real-time 'replacement cost' of its undelivered 1 MWh (at the RT LMP) – but only to the extent that replacement cost exceeds the threshold price *K.*

Here are a few simple examples.

► **Simple settlement examples.** In each case below, assume a resource sells 1 MWh of a day-ahead ancillary service with a strike price of K = \$50 (per MWh) at an option clearing price of V = \$5 (per MWh).

First, consider three cases where the resource produces energy in real-time:

a) The resource produces 1 MWh in real-time and the real-time LMP is **\$60**/MWh. Its net settlement of *V* – max{0, *RT LMP – K*} + RT LMP is

$$\$5 - \max\{0, \$60 - \$50\} + \$60 = \$55.$$

In this case, the resource's net settlement simplifies to *V* + *K.* The real-time settlement result in a payment at the strike price, *K*, rather than the real-time LMP, since the real-time price is higher than the strike. In that situation, standard terminology is to say the option is "in the money."[14]

b) The resource produces 1 MWh in real-time and the real-time LMP is **\$40**/MWh. Its net settlement of *V* – max{0, *RT LMP – K*} + RT LMP is

$$\$5 - \max\{0, \$40 - \$50\} + \$40 = \$45.$$

In this case, the resource's net settlement simplifies to *V* + RT LMP*.* The real-time settlement component results in a payment at the real-time LMP, since the real-time price is higher than the strike. In that situation, standard terminology is to say the option is "out of the money."

c) The resource produces **2 MWh** in real-time and the real-time LMP is **\$60**/MWh (*i.e.,* in the money again). Its net settlement is now *V* – max{0, *RT LMP – K*} + **2 ×** RT LMP, which is

$$\$5 - \max\{0, \$60 - \$50\} + \mathbf{2} \times \$60 = \$115.$$

In this example, the first MWh produced in real-time 'covered' it day-ahead option position, and the second is just an additional MWh sale at the real-time price. In the calculation

[14] This standard terminology conventionally reflects a *buyer's* perspective on a call option's value at the time of delivery.

above, the real-time settlement therefore provides credit for the higher quantity (i.e., 2 MWh) delivered in real-time at the real-time LMP.

Cases (a) and (c) have a useful economic implication and interpretation. In general, a seller of a call option is 'giving up' its potential gain from alternatively selling in the real-time market, if the real-time price is higher than the strike, *K.* This is in contrast to its real-time settlement if it sells its day-ahead energy *forward* and delivers the same amount in real time, which would have a real-time credit of zero. In essence, selling energy forward is giving up the entire potential gain from selling in real-time, but selling a (call) option on energy is not giving up the entire potential gain from selling in real-time. For that reason, call options have lower clearing prices than forward prices for the same delivered product.

Next, consider two cases where the resource sells 1 MWh of a day-ahead ancillary service and does not produce energy in real-time. As before, assume the resource sells 1 MWh of a day-ahead ancillary service with a strike price of K = \$50 (per MWh) at an option clearing price of V = \$5 (per MWh).

d) The resource produces 0 MWh in real-time and the real-time LMP is **\$60**/MWh (*i.e.,* in the money). Its net settlement of V – max{0, *RT LMP* – K} + RT LMP is

$$\$5 - \max\{0, \$60 - \$50\} + \$0 = -\$5,$$

for a net charge of \$5. In this case, the resource's net settlement simplifies to V – (*RT LMP* – K). The real-time settlement component is a charge to the resource, at real-time LMP *less* the strike price.

In effect, in the real-time settlement in this situation, the option seller is paying for the ISO's cost of replacing its day-ahead MWh, but only to the extent that cost exceeds the pre-specified threshold price *K.*

e) The resource produces 0 MWh in real-time and the real-time LMP is **\$40**/MWh (*i.e.,* out of the money). Its net settlement of V – max{0, *RT LMP* – K} + RT LMP is

$$\$5 - \max\{0, \$40 - \$50\} + \$0 = \$5.$$

In this case, the resource keeps the day-ahead clearing price of \$5 for accepting the day-ahead ancillary service obligation, and no money changes hands in real-time settlement. In this case, the resource's net settlement simplifies to just *V.*

Case (d) has an important economic (and incentive) implication. In this case, the resource is compensating the buyer (the ISO, in the immediate instance, and load, ultimately) for the replacement cost of energy that the system must incur in real-time, at the margin, if the seller does not produce – but only to the extent that replacement cost is above the pre-specified threshold (strike) price *K.* This is why a higher threshold (strike) price lowers the up-front cost of acquiring an option (that is, the option clearing price), other things equal.

More importantly, this obligation to cover a resource's replacement cost at the prevailing real-time market price (if it exceeds the strike) changes a seller's performance incentives significantly, relative to today's energy market design. In the current market, if a resource does not clear in the day-ahead market and does not produce in real-time, is net settlement is zero. By contrast, if the resource receives a day-ahead ancillary service award and does not produce in real time, it incurs a potentially steep financial consequence if the real-time cost to replace its MWh is high (e.g., during a stressed operating day).

In this way, the resources that the ISO relies upon for ancillary services in planning the next-day operating plan will have stronger incentives to ensure they can "cover" their positions and operate if and when dispatched during the operating day in the future, relative to the case today when these resources (typically) have no day-ahead energy market award. We will explore these incentives, and how that serves to resolve the misalignment problem generally, in the context of more detailed numerical examples below.

Before doing so, it is useful to touch on how these settlement mechanics relate to real-time dispatch incentives and outcomes.

► **Settlement and co-optimized real-time dispatch: simple examples.** The purpose of the next two simple examples is to show that the day-ahead ancillary services market outcomes do not change the existing co-optimized real-time market's incentives for resources to follow their assigned real-time dispatch.

The next case examines how the settlements work when a resource is dispatched to provide reserves in real-time, rather than energy. Specifically, consider a resource that is awarded a day-ahead ancillary service (of any type), designated for operating reserves in real-time, and its energy dispatch in real-time is zero. As before, assume the resource sells 1 MWh of a day-ahead ancillary service with a strike price of K = \$50 (per MWh) at an option clearing price of V = \$5 (per MWh).

f) Assume the resource's offer cost is \$45/MWh, the real-time LMP is \$60/MWh (in the money), and the real-time reserve clearing price (RT RCP) is \$20/MWh. The resource produces 0 MWh of energy and provides 1 MWh of reserves in real time.

   In this case, its option will be settled as before in example (d), and it will be credited for its 1 MWh of real-time reserves at the RCP. Its total net settlement is:

   1 MWh × *V* – **1 MWh** × max{0, *RT LMP* – *K*} + **0 MWh** × RT LMP + **1 MWh** × RT RCP

   which is \$5 – max{0, \$60 – \$50} + \$0 + \$20 = \$15.

   In this situation, it is sensible that the real-time co-optimization assigned this resource's capability to real-time reserves rather than energy. This is because the real-time reserve price (\$20) exceeds the resource's potential energy margin (\$60 - \$45 = \$15). Regardless of what it sold day-ahead, the resource is better off providing reserves instead of energy in real-time.

This case (f) illustrates an important point. Despite selling the day-ahead ancillary service (which settles as an option on *energy*), in real-time the resource is better off providing reserves than producing energy, given the real-time prices. That is true, generally. Regardless of whether or not a resource has sold a day-ahead ancillary service, the resource is better off following the assigned real-time dispatch (to energy or reserves), given the real-time energy and reserve prices, than doing anything else. The day-ahead ancillary services market outcomes do not change the existing co-optimized real-time market's incentives for resources to follow their assigned real-time dispatch.[15]

Along a similar line, this next simple example considers a situation where the resource re-offers in real-time. As before, assume the resource sells 1 MWh of a day-ahead ancillary service with a strike price of K = \$50 (per MWh) at an option clearing price of V = \$5 (per MWh).

g) The resource's day-ahead energy offer cost is \$45/MWh, and in real-time the resource re-offers energy at \$75/MWh. The real-time LMP is \$60/MWh, the real-time RCP is \$0/MWh, and the resource produces 0 MWh of energy in real-time (as its real-time energy offer price exceeds its RT LMP).

   In this case, its option will be settled as before in example (d). Its net settlement is:

   $$\mathbf{1\ MWh} \times V - \mathbf{1\ MWh} \times \max\{0, \mathit{RT\ LMP} - K\} + \mathbf{0\ MWh} \times \text{RT LMP}$$

   which is \$5 – max{0, \$60 – \$50} + \$0 = – \$5. In this situation, the real-time dispatch produces an efficient outcome (given the RT LMP) in which the resource is 'buying out' its day-ahead ancillary service position in real-time, at a cost of *RT LMP – K =* – \$10. That buy-out cost is less than the resource's cost if produced energy facing the same prices, which would result in a real-time loss of *RT LMP – RT Offer* = – \$15.

This case (g) illustrates another important point. Despite selling the day-ahead ancillary service (which settles as an option on *energy*), in real-time the resource is better off following the assigned real-time dispatch and 'buying out' its day-ahead ancillary service obligation (when that is the real-time dispatch outcome), rather than self-scheduling and producing energy to match its day-ahead option award. That is true generally, whether or not the resource re-offers due to a change in its fuel (say) costs from day-ahead to real-time; we simply assumed it re-offered here to illustrate the latter can (and likely will) happen to resources that re-offer in practice.

4.2.2 Forward Energy and Ancillary Service Positions Settlement

All of the foregoing examples are special cases of a general multi-settlement method when there is more than one day-ahead product. The general method is particularly useful when a single resource receives both forward energy and ancillary services awards in the day ahead E&AS market.

---

[15] Note that, if the RCP was lower (say, \$10) in this example, the real-time co-optimized dispatch would have dispatched this resource for energy and not designated it for real time reserves, because its opportunity cost of energy (\$60 LMP - \$40 offer cost) would exceed the RCP.

To make this concrete, let us assume that in the day ahead E&AS market, a resource sells 1 MWh of energy and 2 MWh of an ancillary service (which specific ancillary service does not matter here) for a particular hour the next day. During that hour (*i.e.,* in real-time), the resource produces 5 MWh of energy. How is this settled?

The various day-ahead and real-time settlements for the two products are easiest to organize in a table, as shown below. (The reserve clearing price is abbreviated RCP below). The resource's total market settlement is the sum of the five entries in the table.

The first row shows the resource's credits (revenue) for its day-ahead market awards. In the second row, the resource's day-ahead energy position is 'closed out' at the real-time price; and, similarly, the resource's day-ahead ancillary services position is 'closed out' using the first part of a call-options real-time settlement, as illustrated above in simple settlement examples (a) thru (g). The last row shows the resource's credits for the energy actually provided, in this case, 5 MWh.

	Forward Sale of Energy	Option Sale of Ancillary Service
Day Ahead Awards (credits)	2 MWh × *DA LMP*	1 MWh × *DA RCP*
Real-Time Closeout of Day-Ahead Positions (debits)	–2 MWh x *RT LMP*	–1 MWh × max{0, *RT LMP* – *K*}
Real-Time Supply (credits)	5 MWh × *RT LMP*	

This multi-product settlement method is commonly employed in commodity markets where participants transact both forwards and options (for the same delivery time). As such, it is not novel market design by any means. In the context of a day-ahead E&AS market, it has several useful properties worth noting here.

First, if a resource does not clear any ancillary services day-ahead, its energy settlement is unchanged from the system in use today. The ISO's existing day-ahead energy market settlement logic, based on real-time energy deviations, is the special case in which the day-ahead ancillary service quantity is 0 MWh (instead of 1 MWh, as assumed above). In that special case, by setting the day-ahead ancillary award quantity (and its closed-out quantity) to 0 MWh and summing the table's five entries, we obtain

$$2 \text{ MWh} \times \mathit{DA\ LMP} + (5 \text{ MWh} - 2 \text{ MWh}) \times \mathit{RT\ LMP}$$

which is the familiar two-settlement deviation logic for energy used today. In other words, the settlement of forward energy positions remains unchanged under this multi-product settlement method.

Second, this multi-settlement method avoids the need for a market participant, or the ISO, to assign (or 'track') its real-time energy production to its distinct day-ahead forward energy obligation and ancillary service obligations. Each day-ahead position is separately closed out in the appropriate way, and then the resource is credited for whatever it provides in real time.

Third, this multi-settlement method handles real-time reserve credits gracefully. As a concrete example, suppose the resource in the table above has the same day-ahead awards, but now let's assume that it provides 5 MWh of energy in real time and, in addition, it also provides (*e.g.*) 6 MWh of real-time reserve. In that case, its total settlements would be the sum of the five entries in the table shown below. The only change from above is the addition of the real-time reserve credit in the last row (noted in red text for emphasis in the table).[16]

	Forward Sale of Energy	Option Sale of Ancillary Service
Day Ahead Awards (credits)	2 MWh × *DA LMP*	1 MWh × *DA RCP*
Real-Time Closeout of Day-Ahead Positions (debits)	–2 MWh x *RT LMP*	–1 MWh × max{0, *RT LMP* – *K*}
Real-Time Supply (credits)	5 MWh × *RT LMP* + 6 MWh × *RT RCP*	

The settlements with real-time reserves, illustrated above, has a more important implication as well. All this works smoothly if day-ahead GCR (that is, day-ahead TMSR, TMNSR, and TMOR) awards are settled as call options on energy, and then designated and priced in real-time based on the same co-optimized real-time market in use today.

Mechanically, it is also possible to settle day-ahead reserve obligations as deviations against the real-time reserve price, rather than as options on energy. That will produce different payments, in general (and is used in some other ISOs, although ISOs' day-ahead reserve market designs vary greatly). We recommend the option settlement method for day-ahead GCR because the settlement of day-ahead reserves as an option on energy can provide better incentives for the resource to invest in energy supply arrangements – that is, to be able to "cover" its day-ahead ancillary service obligation.

Last, we note again here that the actual real-time dispatch (and the economic evaluation of any additional generation commitments after the day-ahead market, if needed) would continue to be based on resources' energy supply offers in effect during the operating day. Resources would continue to be able to re-offer if, for example, fuel costs change during the operating today, consistent with existing market rules and procedures. And like today, day-ahead energy and ancillary services awards do not enter into the real-time dispatch calculations. In this way, the

[16] We simplify here in assuming only one reserve product. In practice, a resource providing operating reserves would receive credit for the quantities of TMOR, TMSR, and TMNSR it provides at their respective prices.

decisions on which resources provide energy and reserves in real-time will continue to reflect the least-cost dispatch of the system.

► **Incentives and implications.** As noted in Section 4.2, these financially-settled day-ahead ancillary service products are real options on energy – the settlement depends on what the resource physically produces in real time. Real options strengthen suppliers' incentives to invest in energy supply (or any other) arrangements that will enable them to more reliably produce on short notice, so that they can 'cover the call' – that is, produce energy during the operating day if instructed to do so. In that way, the resource owner who acquires a day-ahead ancillary service is able to avoid incurring the potentially high cost of buying out its call option position if the real-time price is high and it cannot perform.

These investments have a cost, of course, and sellers of these ancillary services need to be appropriately compensated for the settlement obligations and risks they voluntarily assume. For this reason, it is essential that the day-ahead ancillary services be biddable products. The benefit of a co-optimized day-ahead market for energy and ancillary services it that it can find the lowest clearing prices at which all awarded sellers are willing to accept their assigned obligations, and compensate them competitively for doing so.

4.3 Strike Prices

The introduction of a strike price is a new concept in the ISO's energy and ancillary service markets, presenting questions about how it would be set. Fortunately, real options are sufficiently widely used in other contexts that both economic theory and practical experience provide considerable guidance.

As a general matter, there are three important aspects that guide how strike prices should be addressed in a day-ahead E&AS market design. These are:

1. **Known before offers due**. The numerical value of the strike price must be known to participants in advance of when they must submit energy and ancillary service offers in the day-ahead market;

2. **At the money.** The most efficient outcomes with real options are obtained when the strike price is set at approximately the expected value of the energy price at which the options will settle;

3. **Close is good enough.** In practice, setting the strike price precisely 'at the money' doesn't matter much, within limits.

We explain each of these guidelines, and discuss their practical ramifications for a day-ahead E&AS market.

► **Guideline 1: Known before offers due.** The first guideline requires that the market rules for a day-ahead E&AS provide for a strike price that will be fixed (or "locked down") in its numerical value prior to the offer submission window for the day-ahead market. If *K* is not fixed before offers are due, then suppliers would have no way, in advance of submitting their offers, to anticipate how their positions would settle (*i.e.*, to estimate when they would yield a profit or a loss), since their settlement depends expressly on the strike price.

Put differently, the price of an option – and the (minimum) price a resource owner would be willing to offer to take on a day-ahead ancillary service obligation – depends explicitly on the (numerical) value of the strike price*, K.* This dependence can be seen in the multi-settlement table entries shown above (where the real-time 'close out' of day-ahead positions formula includes the strike price *K*). If resource owners do not know the value of the strike price, they would have no obvious way to formulate a competitive offer price for the call option product.

► **Guideline 2: At the money.** The second guideline specifies that the most efficient value of the strike price is the expected value of price of the underlying good, in this case the real-time LMP for the hour(s) at which the option will be settled. The term of art that goes with this guideline is setting the strike price "at the money."

There's a fairly sophisticated mathematical theory behind this guideline, but the idea is simple. Consider the extreme case where the strike price is set very high – higher than the highest value the RT LMP could go. In this case, in the option's settlement, the value of the real-time close-out of the day-ahead award (that is, the value of –max{0, RT LMP – K}) would always be zero. And if that were the case, the competitive clearing price of day-ahead ancillary services would also be zero, as there would never be any charge applied when resources with day-ahead ancillary service obligations fail to perform.

Ironically, the situation just described is effectively the same, from an incentive and compensation standpoint, as the energy-only day-ahead market construct we have today for resources without day-ahead energy market awards. If the option strike price is set too high, then suppliers with a day-ahead ancillary service obligation face no risk of having to incur the cost of "replacing their MWh" in real-time settlement if they do not perform; their only settlement would be a real-time credit for what the supply in real time. That's the same as what the current market design provides for the resources that do not clear day-ahead. Put simply, if the strike price is set at an extreme high value, then incentives for suppliers to invest in arranging energy supplies in advance – and problem P1 generally – would not be changed from today's market design at all.

Now consider the opposite extreme: A strike price set very low, at zero. Then, by inspection of the settlement table entries above, a value of K =0 makes K irrelevant in the real-time 'close out' of day-ahead positions step in settlements. In fact, in that situation, the real-time settlement boils down to charging the seller the RT LMP, offset by the RT LMP *if* it produces energy.[17] In that case, the "option" product settlement becomes identical to a *forward* sale of energy.

We diagnosed, throughout Section 2, that there are inefficient incentive misalignment problems that arise with today's energy-only day-ahead market construct. In this opposite-extreme case

[17] We simplify, inasmuch as we are ignoring the possibility of a negative RT LMP in this explanation. That simplification does not undermine the conclusion that follows, however.

where the strike is set very low (*i.e.,* zero), we would have a new day-ahead E&AS design that would be still just a *de facto* energy-only day-ahead market design. Said plainly, if the strike price is set far too low, creating a day-ahead ancillary services market achieves nothing.

Where is the just-right value between these unhelpful extremes (what we might call the 'Goldilocks' value of the strike)? Fortunately, specialists have thought much about this, as with option theory generally, over the last several decades.[18] In theory, the preferred strike price value, from the standpoint of providing efficient investment incentives for the resources providing these incentives, is at the expected value of the RT LMP (for the corresponding delivery hours). A strike price that is set materially higher or lower will tend to mute incentives (and, potentially, inefficiently increase risk), undermining the efficiency of the day-ahead E&AS design.

► **Guideline 3: Close is good enough.** This brings us to the third guideline. It might seem ironic that if theory implies that the ideal strike price is "at the money", theory also concedes it doesn't need to be exactly right. That's true, and there's a sensible reason: if the strike is set higher or lower than the 'at the money' value, the impact is non-linear. In simpler terms, a strike price that is set a little bit above the 'at the money' level doesn't change incentives materially. However, a strike price that is set far above the 'at the money' level can undermine incentives, and the benefits of a day-ahead ancillary service design, dramatically. The latter is what we found earlier for the extreme case of a strike above the highest possible RT LMP.

The bottom line is that a strike price that is a bit above, or a bit below, the 'at the money' level will be accounted for in a competitive ancillary service market's day-ahead clearing prices, and netted out in supplier's real-time close out of the day-ahead ancillary service awards, without impacting incentives. Thus, close is good enough.

► **Practicalities.** Given these guidelines, our current thinking on a reasonable approach to setting the strike price for in the day-ahead E&AS market includes the following. First, the strike price will need to be fixed ("locked down") at least a day before participants' supply offers are due for (the multi-day ahead market, to satisfy guideline 1. Second, the locked-down value might be usefully set with reference to an external mark, such as a several-day forward price, that will tend to reflect LMPs on the relevant delivery day, reflecting guidelines 2 and 3. Both the ISO and various power market exchanges could provide such external marks. Presently, we take no position on which might be the most commercially sensible approach for participants, and seek stakeholder input on how this might be usefully structured as the new day-ahead E&AS market design proceeds.

With these observations, we now return to central objectives of creating day-ahead ancillary service products. We next explore how these day-ahead ancillary services products can change resource owners' decisions take 'real' actions – such as to incur the up-front costs of arranging energy supplies in advance of the operating day, even when it may be unlikely to be used. For this purpose, we revisit the prior numerical examples from Section 2 to show how this design helps solve the incentive misalignment problem (P1) present with today's energy-only day-ahead market construct.

---

[18] A well-regarded, if technical, reference is Hull, John C., *Options, Futures, and Other Derivatives* (10th Ed., 2017).

4.4 Example 1, Revisited: Solving Problem P1

In this section, we revisit Example 1 from Section 2 and show that introducing a call option on energy product can incent generators to invest in energy supply arrangements that benefit the system overall. The main point of this example is to illustrate how and why introducing day-ahead ancillary services, when settled as call options on real-time energy, can solve the incentive misalignment problem (P1) discussed in Section 2.

At the outset, it is useful to note that in creating a market product to solve the misalignment problem, the market must achieve two distinct, but interrelated, goals. First, it must compensate the supplier sufficiently that it will be willing to incur the (fixed) costs of arranging energy supplies, whenever that would be cost-effective from the system's standpoint. Second, that compensation cannot simply be a handout. There needs to be a financial consequence tied to whether or not the resource provides energy, so that it will be induced to incur the fixed costs of actions that ultimately help supply energy. In revisiting Example 1 below, we show how both of these elements work.

► **Example 1: A recap.** In Example 1 from Section 2.2.1, a single generator, without a day-ahead market position, faces an unlikely possibility that demand may be high enough for it to operate the next day, and must decide now whether or not to incur the cost of arranging fuel. It knows there is only a 20% change that its resource will be dispatched (if available) the next day, so the fuel it arranges will, in all likelihood, not be used. The main assumptions, from Table 2-1, are reproduced below for convenience.

Table 2-1. Cost and Price Assumptions for Example 1

	With Advance Fuel		No Advance Fuel	
	High Demand	Low Demand	High Demand	Low Demand
Fixed Cost of Advance Fuel	$ 40	$ 40	$ -	$ -
Marginal Cost	$ 70	n/a	$ 400	n/a
Energy Price (LMP)	$ 120	$ 50	$ 400	$ 50
Demand Probability	20%	80%	20%	80%

In analyzing Example 1, the key results we obtained (*see* Section 2.2.1) were:

- The expected (value of the) net benefit to the system of arranging fuel (*i.e.,* expected cost savings) is 20% × ($400 – $70) = $66 by avoiding running the expensive $400 generator if demand is high, minus the $40 up-front cost of arranging fuel, for a net benefit to society of **$26.** The most cost-effective – *i.e.*, efficient – outcome for the system would be achieved if the generator makes the arrangements for fuel in advance of the operating day.

- The expected benefit to the generator of arranging fuel in advance comes from an expected gross margin of 20% × ($120 – $70) = $10 by being able to operate if demand is high, but this is not enough to cover the $40 up-front cost of arranging fuel. The generator's expected profit is therefore a net loss, **–$30**. In other words, arranging fuel in advance is not financially prudent for the generation owner.

The main point of Example 1 was that the energy market, in its *current* form, does not provide sufficient incentives for resource owners to invest proactively in energy supply arrangements – even when such investments would be cost-effective and yield expected net benefits to the system. There is a 'market failure' to incent efficient outcomes, causing higher expected costs to society as a result.

The logic underlying this conclusion is important. In essence, the value that society places on making the fuel arrangement is based on the high price it *avoids* with the investment (*i.e.,* the $400 RT LMP in Table 2-1). However, the value the generator places on the same arrangement is based on the lower price it *receives* in the energy market with the investment (*i.e.,* the $120 RT LMP in Table 2-1). This value difference results in the misaligned incentives problem: a divergence between the social and private benefit of the investment.

► **Example 1 with a day-ahead ancillary service award.** Now let's examine how the outcomes change if the generator in Example 1 has an ancillary service award, settled as a call option on energy. Our point is to illustrate that such a product serves to align the incentives properly: the generator will find it in its private interest to arrange fuel in advance of the operating day, when that action is cost-effective from the system's standpoint.

As in Section 2.2.1, the timeframe in Example 1 was a single hour. Thus, in the discussion next, we do not specify whether the specific day-ahead ancillary service product here is GCR, RER, or EIR; that is, the discussion is equally applicable to any of those products.

Example 1 is simple enough to reveal readily the change in the generator's incentive to arrange fuel in advance of the operating day. For that purpose, we make the following assumptions about the day-ahead ancillary service product prices:

- Assume the option price – that is, the day-ahead clearing price for reserves – is $50/MWh.

- Assume the strike price is $120/MWh.

The first assumption will be sufficient for this generator to be willing to accept the day-ahead ancillary services obligation, given the strike. This will be apparent after a few initial calculations, provided below.

The strike price for this example matches the LMP in the high demand scenario, when the generator has arranged fuel. This particular strike value is an assumption of convenience (to simplify calculations); the conclusion of this example would be unchanged if the strike price was higher or lower over a broad range (though there are limits, of course).[19]

Table 3-1 shows the generator's expected net revenue, for the case where it arranges fuel and the case when it does not. In row [1], we show that the generator receives the $50 day-ahead clearing price for its 1 MWh day-ahead ancillary service award. If demand is high, it is paid the RT LMP of $120/MWh, incurs its fixed cost to arrange fuel of $40 and marginal cost of $70, for a scenario net revenue of $60 in row [8]. If demand is low and it has arranged fuel, it does not operate; the $50 day-ahead award covers the $40 fixed cost of arranging fuel, for a scenario net revenue of $10 in row [8]. Its expected net revenue, if it arranges fuel, is therefore $20; see the bottom-left cell in row [10]. Arranging fuel is now a profitable endeavor, even though there is an 80% chance it would not be used.

Table 3-1. Generator Expected Net Revenue for Example 1, With Option Award

			Advance Fuel		No Advance Fuel	
Generator's Market Settlement			High Demand	Low Demand	High Demand	Low Demand
[1]	Day-Ahead Award	*OCP*	$ 50.00	$ 50.00	$ 50.00	$ 50.00
[2]	Day-Ahead Close-Out	-max{0, *RT LMP* - *K*}	$ -	$ -	$ (280.00)	$ -
[3]	Real-Time	*RT LMP*	$ 120.00	$ -	$ -	$ -
[4]	Total Settlement	[1]+[2]+[3]	$ 170.00	$ 50.00	$ (230.00)	$ 50.00
Generator's Costs						
[5]	Advance Fuel	*F*	$ (40.00)	$ (40.00)	$ -	$ -
[6]	Marginal Cost	*MC*	$ (70.00)	$ -	$ -	$ -
[7]	Total Cost	[5]+[6]	$ (110.00)	$ (40.00)	$ -	$ -
Generator's Net Revenue						
[8]	Scenario Net Revenue	[4]+[7]	$ 60.00	$ 10.00	$ (230.00)	$ 50.00
[9]	Demand Probability	*p* or (1-*p*)	20%	80%	20%	80%
[10]	**Expected Net Revenue**	SumProd [8]*[9]	**$20**		**($6)**	

Now consider the generator's revenue if it does not arrange fuel in advance. It again receives the $50 day-ahead price for its day-ahead reserve award in row [1]. If demand is low, it does not run and incurs no costs, for a scenario net revenue of $50 (as shown in the last column of row [8]). If demand is high, however, it would not be able to operate without fuel arrangements. In this

[19] In this example, any strike price below $250 would suffice to support the example's conclusions about the efficacy of the day-ahead ancillary service product in solving the misalignment problem. This exceptionally wide range of effective strike prices is an artifact of an example with only one generator.

scenario, its cost to settle (or 'buy out') of its day-ahead ancillary service position in real-time settlements, given the high $400 real-time LMP, is

$$-\max\{0, \mathit{RT\ LMP} - K\} = -\max\{0, \$400 - \$120\} = -\$280.$$

See row [2]. Its net revenue in this scenario is $50 – $280 = – $230, a net loss; see row [8]. Taking the scenario likelihood-weighted average, its expected net revenue if it does not arrange fuel in advance of the operating day is a net financial loss of $6, shown in row [10].

Of course, we haven't fully "closed the loop" on this generator's decisions yet. Specifically, would the generator be willing to accept the day-ahead option award, given a clearing price of $50, in the first place? The answer is yes, assuming it is an expected profit-maximizing generator. As Table 3-1 shows in row [10], the generator would now find it financially prudent to incur the $40 cost of arranging fuel in advance, a decision that yields an expected net revenue of $20.

In fact, in this example, an expected profit-maximizing generator would be willing to accept a day-ahead ancillary service award price down to (just above) $30, as that would yield a greater profit than the $0 net revenue it would obtain under its best alternative without a day-ahead ancillary services award (*see* Section 2.2.1, Table 2-2, bottom right cell). Indeed, if this was a broader example with many competitors, this generator's (lowest) profitable offer price for the day-ahead reserve obligation would therefore be (just over) $30.

► **Implication.** The main point here is simple, and important. Real options change behavior – in this case, the generator's willingness to undertake a costly investment in arranging fuel that may not be used. In this example, that willingness arises because the generator's valuation of the investment (the $40 up-front cost to arrange fuel) is no longer based solely on the $120 LMP it receives (at best) when it has fuel. Instead, its valuation of the investment is also based on the $400 LMP that society *avoids* if it makes the investment. Mathematically, this occurs because that $400 RT LMP enters into the generator's cash flows in row [2], when it is in demand but cannot operate. In that case, as noted above, it incurs a charge in the ancillary service market settlement in the amount of:

$$-1\text{ MWh} \times \max\{\,0, \$400\ \mathit{RT\ LMP} - \$120\ \mathit{strike}\,\} = -\$280.$$

Conceptually, the generator is now basing its decision on the same high cost that the system would incur *to replace the generator's output*, even though there is an unlikely (*i.e.,* only 20%) chance that the generator is needed to meet demand.

This function of real-option settlements is quite general, as it aligns the generator and society's incentives to focus on the same, high $400 avoided cost. Selling the option leads the generator to internalize, in its financial calculus, the high cost that may prevail if it cannot operate when its generation is in demand.

► **Cost-effectiveness.** There is a second important implication here as well. We indicated earlier (in Section 4.2) that with day-ahead ancillary services that settle as options on energy, it becomes profitable for generators to pay the up-front costs of maintaining reliable fuel arrangements *when such arrangements are cost effective* from the standpoint of the system overall. Let's focus on that cost-effective attribute now. It implies there are limits to the costs the generator would incur up-front – but those limits align with the limits on what society would find worthwhile as well.

To see why in the context of Example 1, recall with a $40 up-front cost to arrange fuel in advance, the expected cost savings to the system are $26. That means the *most* society would be willing to incur, from the standpoint of cost-effectiveness, would be a $66 up-front cost to arrange fuel in advance. In that situation, the expected benefits and expected costs would be equal.

Now consider the case if the up-front cost was even higher – let's assume, for the moment, it was $75. In that case, the expected (value of the) benefit to the system of arranging fuel (*i.e.,* expected cost savings) is 20% × ($400 – $70) = $66 by avoiding running the expensive $400 generator if demand is high, minus the now **$75** up-front cost of arranging fuel, for a net benefit to society of $66 – $75 = **–$11**. In this case, from the system's standpoint, the investment in advance fuel arrangements is not cost-effective; it would be more cost-effective to run the high-cost $400 generator in the unlikely (*i.e.,* 20%) chance that demand is high.

Does this align with the generator's incentives, in a market that provides it with the opportunity for a day-ahead ancillary service award? Yes – the generator would not find the investment cost-effective either, at a now $75 up-front cost. To see this, note that if the up-front cost is now $75, then the entries in row [5] in Table 3-1 would change from ($40) to ($75), a difference of $35, and the generator's expected net revenue in the bottom row of Table 3-1 – in the case where it arranges fuel in advance – would drop *by* $35, to become a net financial loss of $20 – $35 = **–$15.** Facing that prospect, the generator's prudent financial decision would be to not incur the $75 up-front cost of cost of arranging fuel in advance, in which case its expected net revenue would be zero (*see* Section 2.2.1, Table 2-2, bottom right cell).

One can do this same exercise with a range of possible up-front investment costs, with the same pattern of conclusions. The generator would find it financially prudent to invest if the up-front cost of arranging fuel is up to $66, but not any higher. That matches, exactly, the maximum investment that would be cost-effective as well.

At first read, this conclusion might seem to be an artifact of the particular numbers chosen in Example 1. However, that is not the case. The conclusion illustrated here is a general property. Providing generators with the opportunity to compete for day-ahead ancillary service awards that are, in real-time, settled as call options on energy will make it financially prudent for an expected profit-maximizing generator to incur the up-front, fixed cost of arranging energy supplies in advance – whenever those arrangements would be cost-effective from society's standpoint as well.[20]

[20] In Sections 2.2.2 and 2.2.4, we also provided several variations on Example 1 with different assumptions (*e.g.,* when PFP applies). It can be shown that these same conclusions about the efficacy of the day-ahead ancillary service award hold for all of those Example 1 variations as well (we omit the detailed calculations here).

4.5 Example 2, Revisited

In this section, we revisit the co-optimization problem from Example 2 in Section 2.2.1, and show that introducing a call option on energy as a day-ahead ancillary service improves the generator's incentives to invest in energy supply arrangements, to the benefit of the system overall. The central point of this example is to again illustrate, in a more complex setting with multiple generators and multiple products, how and why introducing day-ahead ancillary services, when settled as call options on real-time energy, can solve the incentive misalignment problem (P1) discussed in Section 2.

The structure of Example 2 lends itself to three additional points, which we will highlight. First, this example will show how the reserve settlements function when day-ahead GCR (*i.e.*, day-ahead operating reserves) are settled as a call option on energy, and real-time reserves are settled based on the designations and prices from the real-time cooptimized dispatch in use today. We emphasized this design approach earlier in Section 4.2.2.

Second, we noted in Section 4.1.1 that the co-optimized clearing in the day-ahead market of energy and the new ancillary services will tend to result in ancillary service prices that also impact the day-ahead energy price. Example 2 illustrates how this works.

Last, in this example the total market revenue – and therefore total E&AS payments by wholesale buyers (and, ultimately, by consumers) – is higher under the day-ahead E&AS design than under the existing energy-only day-ahead market. This is true even though the new E&AS design produces more efficient outcomes – that is, the power system operates more cost-effectively (at lower total production costs) overall. We expect these observations to hold in practice. The principal reason for the increase in total day-ahead market payments is that the market is now compensating suppliers, at transparent, competitive prices, for the ancillary services that the ISO has always relied upon in preparing the system's next-day operating plan – but simply did not compensate suppliers for providing in the past, to the detriment of a cost-effective system that minimizes reliability risk.

► **Example 2: A recap.** In Example 2 from Section 2.3.3, there are four generators that can provide both energy and operating reserve. Real-time demand is uncertain, and the higher-cost generators (Generator 3 and Generator 4) do not receive day-ahead market awards. Generator 3 faces the possibility that real-time demand may be high enough for it to operate the next day, and it must decide whether or not to incur the cost of arranging fuel in advance of the operating day. The main assumptions, from Table 2-1, are reproduced below for convenience.

The additional, market level assumptions for Example 2 are a day-ahead energy demand of 190 MWh for the hour, and a reserve requirement (now both day-ahead and real-time) of 30 MWh/h.

Table 2-5. Assumptions for Example 2			
Generator	**Capacity (MW)**	**Offer Price ($/MWh)**	**Reserve Capability (MW)**
Gen 1	100	$25	10
Gen 2	100	$30	20
Gen 3	50	$40	30
Gen 4	50	$90	40
Additional Cost Assumptions for Generator 3			
		Marginal Cost	**Fixed Cost**
With Advance Fuel Arrangements		$40	$150
No Advance Fuel Arrangements		N/A	N/A
Real-Time Demand Scenarios			
	Low Demand	**Medium Demand**	**High Demand**
Energy Demand (MWh)	170	190	210
Scenario Probability	33%	33%	33%

In analyzing Example 2 in Section 2.3.3, the key results we obtained were:

- The expected system total production cost is **$5,233** for the hour, if Generator 3 arranges fuel in advance of the operating day, and **$5,400** if it does not. See Table 2-6 and 2-7 (Section 2.3.3). This difference, or $167, is more than enough to cover the $150 up-front fixed cost of the fuel arrangements. Thus, the most efficient, cost-effective outcome for the system would be achieved if Generator 3 makes the arrangements for fuel in advance of the operating day.

- If Generator 3 does not arrange fuel, it produces zero energy in real-time (in any demand scenario) and its expected net revenue is $0. If it does arrange fuel, Generator 3's expected net revenue is a net financial loss, of $150. See Table 2-8 (Section 2.3.3). In other words, arranging fuel in advance was not financially prudent for Generator 3.

- Under the status quo energy market construct, the expected total market settlement is **$6,100** for the hour. See Table 2-7 (Section 2.3.3).

From the analysis of these results in Section 2.3.3, our principal conclusion is that that the energy market, in its *current* form with real-time cooptimized energy and reserves, would not provide sufficient incentives for Generator 3 to incur the up-front fixed cost of arranging energy supplies in advance – even though making those arrangements would be cost-effective from the system's standpoint. Under the status quo, the generator's incentives are misaligned with society's interest in operating an efficient, least-cost power system.

► **Example 2 with ancillary services as energy options.** Now let's examine how the outcomes change if the Generator 3 has an ancillary service award defined (and settled) as a call option on energy. Our point is to illustrate that such a product serves to align the incentives properly: Generator 3 will find it in its private interest to arrange fuel in advance of the operating day, and that action is cost-effective from the system's standpoint.

As in Section 2.3.3, the timeframe in Example 2 is a single delivery hour. Since Example 2 focused on co-optimization of energy and operating reserves, it will be useful to interpret the day-ahead ancillary service here as providing GCR; however, the broader conclusions below are not restricted by that interpretation. As in Example 2 earlier, we assume (for simplicity) a single day-ahead GCR product and a single corresponding real-time operating reserve product.

We make the following assumptions about day-ahead ancillary service pricing:

- The strike price is $35/MWh. This is approximately the average value of the real-time LMP in this example, consistent with the concepts discussed in Section 4.3 on strike prices.

- The day-ahead ancillary service offer price from low-cost Generator 2 is $2/MWh, medium-cost Generator 3 is $11/MWh, and high-cost Generator 4 is $17/MWh.

The assumptions about the offer prices for day-ahead ancillary services from Generators 2, 3, and 4 are consistent with profitable offers for those services from each generator, given the competition they face in the day-ahead market (with one another) for both energy and for ancillary services, under the assumption that their fixed costs of arranging fuel in advance of the operating day to cover a day-ahead ancillary service award is $100, $150, and $150, respectively. Since Generator 3 was assumed to have an up-front fixed cost of arranging fuel of $150 throughout Example 2, the substantive new assumption is that Generator 2's fixed cost is lower (at $100), and Generator 4's (at $150) is no less than Generator 3's.[21]

► **Day-ahead co-optimized energy and ancillary services.** With the above setup, we first evaluate the day-ahead E&AS market outcomes. This will differ from the day-ahead market outcome when there was no day-ahead ancillary service, shown previously in Figure 2-1 in Section 2.3.3.

The day-ahead market outcome with both energy and the day-ahead ancillary service is shown in Figure 3-1 below. The two lower-cost generators (Generator 1 and Generator 2) receive day-ahead energy awards, and Generator 2 and Generator 3 receive day-ahead ancillary service awards. The total ancillary services procured (just) satisfies the day-ahead requirement of 30 MWh/h.

The marginal ancillary services provider is Generator 3, which sets the day-ahead ancillary service clearing price at its ancillary service offer price of $11/MWh.

---

[21] The salient assumption here is the ordering of these fixed costs (*i.e.*, that Generator 2's is lower than Generator 3's, and Generator 4's is at least as high as Generator 3's). This example's conclusions would generally follow with different numerical values that respect these cost relationships.

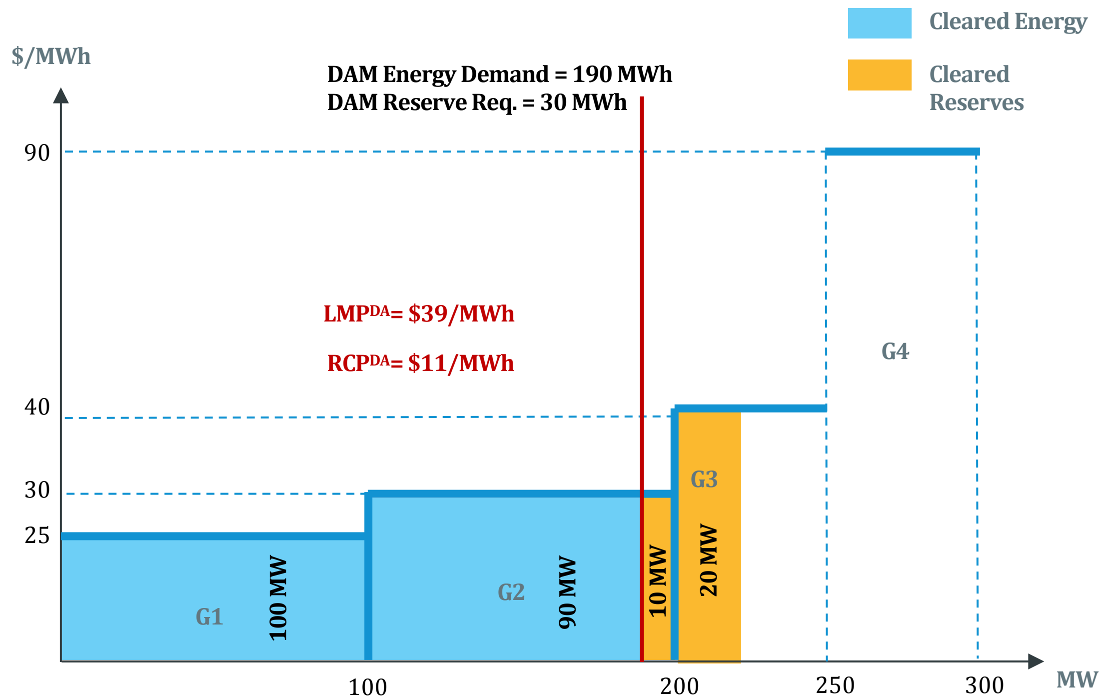

Figure 3-1. Day-ahead Market Outcomes for Example 2, with Energy and Ancillary Service

The day-ahead LMP is $39/MWh, and reflects the pricing of *both* energy and ancillary services offers. Specifically, the highest-price offer for energy cleared in the day-ahead market is that of Generator 2, at $30/MWh. The LMP is not set by just one generator's offer price, however. It is set, as LMPs are generally, by the change in the system's total production cost if there is another increment of energy demand. That, in turn, would require a "redispatch" of the ancillary services awards, at an incremental cost of another $9/MWh, which produces the $39/MWh LMP for energy.

To see this more precisely, let us step through this "redispatch" logic. Suppose day-ahead energy demand increased from 190 MWh, by 1 additional MWh. The least-cost solution would then increase the energy cleared from Generator 2, at an incremental cost of $30/MWh (its energy offer price). However, that additional cleared energy reduces Generator 2's MWh available for ancillary services by 1 MWh, from 10 MWh to 9 MWh. To replace that 1 MWh of ancillary service and still satisfy the day-ahead ancillary services requirement, the market would then clear 1 additional MWh of ancillary service from Generator 3. The net cost of this "redispatch" of 1 MW of ancillary service from Generator 2 (which offered the ancillary service at $2/MWh) to Generator 3 (which offered at $11/MWh) is therefore the difference in their offer prices, or $11 – $2 = $9/MWh. Putting it all together, the incremental cost of another 1 MWh of *energy* demand, while still maintaining the day-ahead ancillary service requirement, is $30 + $9 = $39/MWh. The DA LMP for energy is, therefore, $39/MWh.

In this way, creating a day-ahead E&AS market raises the energy price – and the revenue of all day-ahead cleared resources with energy awards – relative to a day-ahead energy market alone. In Example 2 originally, the day-ahead LMP was only $30/MWh. This is a broader and general point; while the day-ahead E&AS market clearing will not *always* produce a higher day-ahead LMP than would a day-ahead energy-only market, the co-optimization illustrated above will tend to produce that change in outcomes (depending, in practice, on resources' offers, demands, and so on).

The day-ahead E&AS clearing here also illustrates another, more subtle observation: With a co-optimized E&AS market with ancillary service *offers,* the LMP for energy may not be set by the offer price of any one generator all. Rather, it may be set by a combination of several generators' energy and ancillary service offer prices. In this way, the LMPs for energy may commonly reflect one (or more) suppliers' ancillary service offer prices.

► **Full market awards and outcomes.** The real-time market outcomes for Example 2 with the day-ahead E&AS market remain unchanged from those shown previously in Example 2 in Section 2.3.3. (*See* Figures 2-2, 2-3, and 2-4.)

For reference, the full set of market outcomes for the case when Generator 3 arranges fuel in advance (Case A) are shown in Table 3-2 below; and when Generator 3 does not arrange fuel in advance (Case B), in Table 3-3 below. Cell differences from Table 3-2 to Table 3-3 are shaded in light orange in Table 3-3 for reference.

Table 3-2. Market Outcomes for Example 2 with Day Ahead E&AS Market, Case A: Generator 3 With Fuel

		Day Ahead Market Awards		**Real-Time Market Outcomes: Low Demand**		**Medium Demand**		**High Demand**	
	Generator	Energy	Option	Energy	Reserve	Energy	Reserve	Energy	Reserve
[1]	Gen 1	100	0	100	0	100	0	100	0
[2]	Gen 2	90	10	70	20	90	10	100	0
[3]	Gen 3	0	20	0	30	0	30	10	30
[4]	Gen 4	0	0	0	40	0	40	0	40
[5]	**Totals**	190	30	170	90	190	80	210	70
[6]	**Clearing Price**	$39	$11	$30	$0	$30	$0	$40	$0
[7]	Scenario Total Production Cost			$4,600		$5,200		$5,900	
[8]	Demand Probability			33%		33%		33%	
[9]	**Expected Total System Production Cost**					$5,233			
[10]	Scenario Market Payments (incl. DAM)			$7,203		$7,803		$8,453	
[11]	**Expected Total Market Payments**					$7,819			

Table 3-3. Market Outcomes for Example 2 with Day Ahead E&AS Market, Case B: Generator 3 Without Fuel									
		Day Ahead		Real-Time Market Outcomes					
		Market Awards		Low Demand		Medium Demand		High Demand	
	Generator	Energy	Option	Energy	Reserve	Energy	Reserve	Energy	Reserve
[1]	Gen 1	100	0	100	0	100	0	100	0
[2]	Gen 2	90	10	70	20	90	10	100	0
[3]	Gen 3	0	20	0	30	0	30	0	0
[4]	Gen 4	0	0	0	40	0	40	10	40
[5]	**Totals**	190	30	170	90	190	80	210	40
[6]	**Clearing Price**	$39	$11	$30	$0	$30	$0	$90	$0
[7]	Scenario Total Production Cost			$4,600		$5,200		$6,400	
[8]	Demand Probability			33%		33%		33%	
[9]	**Expected Total System Production Cost**					**$5,400**			
[10]	Scenario Market Payments (incl. DAM)			$7,203		$7,803		$7,953	
[11]	**Expected Total Market Payments**					**$7,653**			

► **Generator 3's decision with the day-ahead ancillary service award.** Given these market outcomes, let's now examine what decision maximizes Generator 3's expected net revenue. For this, we make use of the full market outcome results (in Tables 3-2 and 3-3) with the day-ahead E&AS market and the general energy-option settlement rules, as discussed in Section 4.2.2. Table 3-4 shows Generator 3's expected net revenue, for the case where it arranges fuel (Case A) and the case when it does not (Case B).

Table 3-4. Generator 3's Expected Net Revenue for Example 2 with Day-Ahead E&AS Market								
			Advance Fuel			No Advance Fuel		
Generator's Market Settlements		Calculation	Low Dmd	Med Dmd	High Dmd	Low Dmd	Med Dmd	High Dmd
[1]	Day Ahead Energy	*DA LMP * Qe_DA*	$ -	$ -	$ -	$ -	$ -	$ -
[2]	Real-Time Energy Close-Out	*-RT LMP * Qe_DA*	$ -	$ -	$ -	$ -	$ -	$ -
[3]	Day Ahead Option	*DA OCP * Qo_DA*	$ 220	$ 220	$ 220	$ 220	$ 220	$ 220
[4]	Day Ahead Option Close-Out	*-max(RT LMP-K, 0) * Qo_DA*	$ -	$ -	($100)	$ -	$ -	($1,100)
[5]	Real-Time Energy	*RT LMP * Qe_RT*	$ -	$ -	$ 400	$ -	$ -	$ -
[6]	Real-Time Reserves	*RT RCP * Qr_RT*	$ -	$ -	$ -	$ -	$ -	$ -
[7]	**Total Settlement**	[1]+[2]+[3]+[4]+[5]+[6]	$ 220	$ 220	$ 520	$ 220	$ 220	$ (880)
Generator's Costs								
[8]	Advance Fuel	*F*	$ (150)	$ (150)	$ (150)	$ -	$ -	$ -
[9]	Variable Cost	*MC*	$ -	$ -	$ (400)	$ -	$ -	NA
[10]	**Total Cost**	[8]+[9]	$ (150)	$ (150)	$ (550)	$ -	$ -	$ -
Generator's Expected Profit								
[11]	Scenario Net Revenue	[7]+[10]	$ 70	$ 70	$ (30)	$ 220	$ 220	$ (880)
[12]	Demand Probability	*p* or (1-*p*)	0.333	0.333	0.333	0.333	0.333	0.333
[13]	**Expected Net Revenue**	SumProd [11]x[12]		$37			($147)	

In row [3] of Table 3-4, we show that Generator 3 receives a day-ahead market ancillary service credit of $220, on an award of 20 MWh of ancillary service at the $11 day-ahead clearing price for the ancillary service option (shown as 'DA OCP' in row [3]).

In row [4], we show its real-time close-out of its day-ahead option award settlement. In Case A, the LMP in the high demand scenario is $40/MWh, just above the strike price of $35, so this settlement amount (a charge) is:

$$– 20 \text{ MWh} \times \max\{0, \mathit{RT\ LMP} – K\}/\text{MWh} = – 20 \text{ MWh x } \max\{0, \$40 - \$35\}/\text{MWh} = – \$100.$$

In Case B, when the generator does not have fuel, the LMP in the high demand scenario is $90/MWh so this settlement amount is a much greater charge:

$$– 20 \text{ MWh} \times \max\{0, \mathit{RT\ LMP} – K\}/\text{MWh} = – 20 \text{ MWh x } \max\{0, \$90 - \$35\}/\text{MWh} = – \$1100.$$

Next, rows [5] and [9] show that the generator's energy revenue and variable fuel costs are a wash. This is because Generator 3 offered energy at cost and sets the RT LMP in the only scenario when it produces energy, so it has no gross margin on its real-time energy output. Last, row [8] shows Generator 3's $150 up-front cost of arranging fuel, for the Case A scenarios when it does so.

The bottom row of Table 3-4 summarizes the results. Generator 3's expected net revenue, if it arranges fuel, is now $37. Arranging fuel is now a profitable endeavor, even though there is only a 33% chance it would be used. Compare this with Generator 3's decision when there is no day-ahead ancillary services market, as shown in Table 2-8 (*see* Section 2.3.3). In that case, Generator 3 did not run at all and had expected net revenue of zero. Last, Table 3-4 shows that if Generator 3 does not arrange fuel and clears the day-ahead ancillary services position of 20 MWh, its expected net revenue is a financial loss of $147.

The immediate point to observe is that with the day-ahead ancillary service product defined as a call option on energy, and a co-optimized day-ahead E&AS market, Generator 3 would be willing to incur the $150 up front fixed cost of arranging fuel in advance of the operating day. Indeed, in this example, Generator 3 would be willing to accept a day-ahead ancillary service award at a clearing price down to (just above) $9.17/MWh, as that would yield a positive expected profit – still better than the zero expected net revenue it would obtain under its best alternative without a day-ahead ancillary services award (*see* Section 2.3.3, Table 2-8, bottom right cell).

► **Implications.** The implication of this example, and our reason for revisiting it, is important: The opportunity to sell a real option on energy, as a day-ahead ancillary service, improves the generator's willingness to undertake a costly investment in arranging fuel – even knowing that fuel may not be used.

In this example, that willingness arises because the generator's valuation of the investment is no longer based solely on the $40 LMP (at best) that it earns when it has fuel. Instead, its valuation of the investment is also based on the $90 LMP that society *avoids* if it makes the investment. This $90 LMP is accounted for in the generator's financial calculus in row [4] of Table 3-4, where it drives the steep –$1,100 charge if Generator 3 fails to cover its day-ahead ancillary service position by arranging the fuel to operate and the high-demand scenario subsequently occurs.

As noted in Example 1, this function of real-option settlements is quite general, as it aligns the generator and society's incentives to similarly account for the same, high $90 cost of "replacing" Generator 3's energy whenever it does not perform and holds an ancillary service obligation. As a result, there is no divergence between the value that society places on the investment in its energy supply arrangements, and the value that the generator places on the same investment. The real-option design of the day-ahead ancillary service product solves the misalignment problem, and would lead the generator to incur the fixed costs of making energy supply arrangements whenever they would be cost-effective for the system as a whole.

That property is a general one with this real-option design of a day-ahead E&AS market. We could create additional numerical examples *ad nauseam,* but they would all come to the same conclusion: A real-option design of a day-ahead E&AS market aligns a resource owner's incentives to invest in energy supply arrangements with the costs that society would incur, at the margin, if it fails to do so. As a result, there is no longer a divergence between the social and private benefit of the investment. Put succinctly, this market design solves problem P1.

For completeness, it is important to emphasize that this does not imply lower levels of total payments by wholesale buyers (or, ultimately, consumers). Under the status quo, when Generator 3 was not incented to arrange fuel (and the higher-cost generator must be used in its place during high-demand scenarios), the total market payments were $6,100 (*see* Table 2-7 in Section 2.3.3, bottom row). Under the day-ahead E&AS design, where Generator 3 arranges fuel in advance and the systems expected total production costs are *lower,* the total market payments are *higher*, at $7,819 (*see* Table 3-2, bottom row). The reason for this increase in total market payments is not subtle: The new day-ahead E&AS market is now compensating resources for the ancillary services capabilities that the ISO, and ultimately consumers, were relying upon previously, but were not compensating supplier for in the day-ahead timeframe. With the day-ahead E&AS market design, the market now signals, through transparent prices, the total cost to wholesale buyers of maintaining a reliable power system.

5. Continuing Efforts

The purpose of this paper is to share the ISO's current perspective on market design improvements to address regional energy security, and the FERC's directive for the ISO "to develop longer-term market solutions."[22]

To that end, this paper shares the ISO's current assessment of the underlying problems, root causes, and the existing energy market's limitations. While the power system's growing reliance on technologies with just-in-time energy sources poses new challenges, we believe these challenges have sensible solutions. Further, we readily agree with the FERC's affirmation that these challenges are appropriately addressed through market mechanisms.

---

[22] *See* n. 2, *supra.*

The ISO looks forward to discussing these concepts and recommendations with stakeholders, with the objective of modifying the energy and ancillary services market rules to address these challenges. We recognize that these changes will require a significant amount of time and effort from the region. These efforts are important to ensure the continued competitiveness and reliability of the region's electric system. We hope this paper is informative, and look forward to the opportunity to discuss these changes with stakeholders.